

# IEEE Signal Processing MAGAZINE

Volume 37 | Number 2 | March 2020

## SINGLE-PARTICLE CRYO-ELECTRON MICROSCOPY

Mathematical theory,  
computational challenges,  
and opportunities

Cooperative Wireless  
Mobile Caching

Harnessing Sparsity  
Over the Continuum

Delivering Standards to  
Industries: The MPEG Case

In Memory of  
A.H. "Steen" Gray Jr.

IEEE  
Signal  
Processing  
Society



# CALL FOR PAPERS

## 54<sup>th</sup> Annual Asilomar Conference on Signals, Systems, and Computers



Asilomar Hotel and Conference Grounds  
Pacific Grove, California  
November 1 – 4, 2020  
[www.asilomarssc.org](http://www.asilomarssc.org)



Authors are invited to submit papers in the following areas:

**A. Communications Systems:** 1. Modulation and Coding, 2. Cognitive Radios and Spectrum Sharing, 3. Ultra-Low Latency, 4. Optimization, 5. Physical Layer Security and Privacy, 6. mmWave, 7. Underwater, 8. Wireline and Optical Communications, 9. Satellite, 10. IoT, 11. V2V, 12. 5G and Beyond

**B. MIMO Communications and Signal Processing:** 1. Multiuser and Massive MIMO, 2. Channel Estimation and Equalization, 3. Full-Duplex, 4. Cooperation and Relaying, 5. Interference Management and Awareness, 6. mmWave, 7. Optimization

**C. Networks:** 1. Wireless Networks, 2. IoT, 3. Network Information Theory, 4. Optimization, 5. Graph Signal Processing, 6. Social Networks, 7. Distributed Algorithms, 8. Security, 9. Computational Offloading, 10. Self-Organizing Networks, 11. UAV and V2V Networks, 12. Smart Grid

**D. Adaptive Systems, Machine Learning, Data Analytics:** 1. Compressive Sensing, 2. Supervised Learning, 3. Unsupervised Learning, 4. Estimation and Inferences, 5. Adaptive and Cognitive Systems, 6. Adaptive Filtering, 7. Fast and Scalable Algorithms, 8. Models for High-Dimensional Large-Scale Data, 9. Distributed Computation and Storage, 10. Deep Learning, 11. Reinforcement Learning

**E. Array Processing and Multisensor Systems:** 1. Source Localization and Separation, 2. Beamforming, 3. Robust Methods,

4. MIMO and Cognitive Radars, 5. Tensor Models and Processing, 6. Sparse Sensor Arrays, 7. Optimization, 8. Applications (Imaging, Radar, Sonar, Microphone Arrays)

**F. Biomedical Signal and Image Processing:** 1. Molecular and Medical Imaging, 2. Computational Imaging, 3. Neuroengineering, 4. Processing of Physiological Signals, 5. Bioinformatics and Computational Biology, 6. Image Registration and Multimodal Imaging, 7. Functional Imaging, 8. Brain Machine Interfaces, 9. Neural Signal Processing, 10. Visualization

**G. Architectures and Implementation:** 1. Energy Efficiency, 2. Accelerators, 3. Reconfigurable Processing, 4. Non-idealities, 5. Multicore, Many-core and Distributed Systems, 6. Algorithm and Architecture Co-optimization, 7. Architectures for Big Data, 8. Architectures for Machine Learning, 9. Cyber-Physical Systems, 10. Test-beds, 11. Mixed-Signal Processors, 12. Arithmetic and Algorithms

**H. Speech, Image and Video Processing:** 1. Speech Coding, 2. Speech Recognition, 3. Audio Coding, 4. Document Processing, 5. Models for Speech and Image Processing, 6. Image and Video Coding, 7. Learning and Autonomous Systems, 8. Computer Vision, Image and Video Analysis, 9. Image and Video Forensics, 10. Biometrics and Security, 11. Hybrid Imaging Systems

Submissions must include a 50–100 word abstract, an extended summary (500–1000 words plus figures), the title of the paper, each author's name and affiliation, technical area, and topic designation from the above list. The conference website ([www.asilomarssc.org](http://www.asilomarssc.org)) provides specific information on the electronic submission and publication processes. The conference executive committee and IEEE reserve the right to exclude a final paper submission from inclusion in the conference proceedings, if the final paper submission does not meet conference or IEEE standards for scope and/or quality. Submissions will be accepted starting February 1, 2020. **All submissions must be received by May 1<sup>st</sup>, 2020.**

- **No more than FOUR submissions are allowed per contributor, as author or co-author.**
- **To have a paper published, an author must register and present the paper at the conference.**

Notifications of acceptance will be mailed by mid-July 2020. Full papers are due shortly after the conference and published in early 2021. All technical questions should be directed to the Technical Program Chair **Prof. Marco F. Duarte** ([mduarte@ecs.umass.edu](mailto:mduarte@ecs.umass.edu)) or to the General Chair **Prof. Joseph R. Cavallaro** ([cavallar@rice.edu](mailto:cavallar@rice.edu)). Prospective organizers of special sessions containing typically four papers are invited to submit proposals to the General or Technical Chair by January 31, 2020. Proposals must include title, topic, rationale, session outline, contact information, and a description of the session organization with paper titles and contact information for proposed special session authors.

### CONFERENCE COMMITTEE (\* serving in his/her capacity)

General Chair: Joseph R. Cavallaro, *Rice University*  
Technical Program Chair: Marco F. Duarte, *University of Massachusetts Amherst*  
Conference Coordinator: Monique P. Fargues\*, *Naval Postgraduate School*  
Publication Chair: Michael Matthews, *Northrop Grumman Innovation Systems*  
Publicity Chair: Linda S. DeBrunner, *Florida State University*  
Finance Chair: John D. Roth\*, *Naval Postgraduate School*  
Electronic Media Chair: Marios S. Pattichis, *University of New Mexico*


The site for the 2020 Conference is at the Asilomar Conference Grounds, in Pacific Grove, CA. The grounds border the Pacific Ocean and are close to Monterey, Carmel, and the scenic Seventeen Mile Drive in Pebble Beach. The Conference is organized by the non-profit Signals, Systems and Computers Conference Corporation.

# Contents

Volume 37 | Number 2 | March 2020

## FEATURES

- 18 COOPERATIVE WIRELESS MOBILE CACHING**  
Soheil Mohajer, Itsik Bergel, and Giuseppe Caire
- 39 HARNESSING SPARSITY OVER THE CONTINUUM**  
Yuejie Chi and Maxime Ferreira Da Costa
- 58 SINGLE-PARTICLE CRYO-ELECTRON MICROSCOPY**  
Tamir Bendory, Alberto Bartesaghi, and Amit Singer



## ON THE COVER

This issue of *IEEE Signal Processing Magazine* includes three feature articles covering cooperative wireless caching, atomic norm minimization for super-resolution, and single-particle cryo-electron microscopy (the latter on which this issue's cover design is based). Happy reading!

COVER IMAGE: EMPIAR: ELECTRON MICROSCOPY PUBLIC IMAGE ARCHIVE

- 14 Reader's Choice**  
Top Downloads in IEEE Xplore
- 77 Lecture Notes**  
The Importance of Continuity for Linear Time-Invariant Systems  
*Ming Zhang and Xiaoming Chen*
- 81 Standards in a Nutshell**  
Delivering Standards to Industries: The MPEG Case  
*Leonardo Chiariglione*
- The Challenge: From MPEG Intellectual Property Rights Ontologies to Smart Contracts and Blockchains  
*Panos Kudumakis, Thomas Wilmering, Mark Sandler, Victor Rodríguez-Doncel, Laurent Boch, and Jaime Delgado*
- 96 Reflections**  
In Memory of A.H. "Steen" Gray Jr.  
Robert M. Gray
- 104 In the Spotlight**  
Highlights From the Signal Processing Theory and Methods Technical Committee  
*José C.M. Bermudez, Mónica F. Bugallo, and Alle-Jan van der Veen*



PG. 10

## COLUMNS

- 5 Society News**  
SPS Fellows and Award Winners Recognized
- 10 Special Reports**  
Robotics Rolls Into High Gear With Signal Processing  
*John Edwards*



PG. 104

**IEEE SIGNAL PROCESSING MAGAZINE** (ISSN 1053-5888) (ISPREG) is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, 17th Floor, New York, NY 10016-5997 USA (+1 212 419 7900). Responsibility for the contents rests upon the authors and not the IEEE, the Society, or its members. Annual member subscriptions included in Society fee. Nonmember subscriptions available upon request. **Individual copies:** IEEE Members US\$20.00 (first copy only), nonmembers US\$246 per copy. Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of U.S. Copyright Law for private use of patrons: 1) those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923 USA; 2) pre-1978 articles without fee. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. **For all other copying, reprint, or republication permission,** write to IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854 USA. Copyright © 2020 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Periodicals postage paid at New York, NY, and at additional mailing offices. **Postmaster:** Send address changes to IEEE Signal Processing Magazine, IEEE, 445 Hoes Lane, Piscataway, NJ 08854 USA. Canadian GST #125634188 **Printed in the U.S.A.**

Digital Object Identifier 10.1109/MSP.2019.2960988



## 3 From the Editor

What Does an Editor-in-Chief of *IEEE Signal Processing Magazine* Do, Anyway?  
Robert W. Heath, Jr.

## 101 Dates Ahead



The 54th Annual Asilomar Conference on Signals, Systems, and Computers will be held 1–4 November in Pacific Grove, California.

ISTOCKPHOTO.COM

## EDITOR-IN-CHIEF

Robert W. Heath, Jr.—The University of Texas at Austin, U.S.A.

## AREA EDITORS

## Feature Articles

Matthew McKay—Hong Kong University of Science and Technology, Hong Kong SAR of China

## Special Issues

Namrata Vaswani—Iowa State University, U.S.A.

## Columns and Forum

Rodrigo Capobianco Guido—São Paulo State University (UNESP), Brazil

Roberto Togneri—The University of Western Australia

## e-Newsletter

Ervin Sejdic—University of Pittsburgh, U.S.A.

## Social Media and Outreach

Tiago Henrique Falk—INRS, Canada

## Special Initiatives

Andres Kwasinski—Rochester Institute of Technology, U.S.A.

Nuria Gonzalez Prelcic—Universidade de Vigo, Spain

## EDITORIAL BOARD

Daniel Bliss—Arizona State University, U.S.A.

Danijela Cabric—University of California, U.S.A.

Volkan Cevher—École polytechnique fédérale de Lausanne, Switzerland

Mrityunjoy Chakraborty—Indian Institute of Technology, Kharagpur, India

George Christikos—Qualcomm, Inc., U.S.A.

Elza Ekip—New York University, U.S.A.

Alfonso Farina—Leonardo S.p.A., Italy

Yifan Gong—Microsoft Corporation, U.S.A.

B. Vikram Gowreesunker—Texas Instruments Inc., U.S.A.

Joseph Guerri—Information Systems Laboratories, Inc., U.S.A.

Nageen Himayat—Intel, U.S.A.

Clem Karl—Boston University, U.S.A.

C.-C. Jay Kuo—University of Southern California, U.S.A.

Erik Larsson—Linköping University, Sweden

David Love—Purdue University, U.S.A.

Maria G. Martini—Kingston University, U.K.

Helen Meng—The Chinese University of Hong Kong, China

Meinard Mueller—Friedrich-Alexander Universität Erlangen-Nürnberg, Germany

Phillip A. Regalia—U.S. National Science Foundation

Alejandro Ribeiro—University of Pennsylvania, U.S.A.

Douglas O'Shaughnessy—INRS Université de Recherche, Canada

Osvaldo Simeone—Kings College London, U.K.

Milica Stojanovic—Northeastern University, U.S.A.

Ananthram Swami, Army Research Labs, U.S.A.

Jong Chul Ye—KAIST, South Korea

Qing Zhao—Cornell University, U.S.A.

Josiane Zerubia—INRIA Sophia-Antipolis Méditerranée, France

## ASSOCIATE EDITORS—COLUMNS AND FORUM

Ivan Bajic—Simon Fraser University, Canada

Balázs Bank—Budapest University of Technology and Economics, Hungary

Panayiotis (Panos) Georgiou—University of Southern California, U.S.A.

Hana Godrich—Rutgers University, U.S.A.

Yuan-Hao Huang—National Tsing Hua University, Taiwan

Euee Seon Jang—Hanyang University, South Korea

Vishal M. Patel—Rutgers University, U.S.A.

Christian Ritz—University of Wollongong, Australia

Changshui Zhang—Tsinghua University, China

H. Vicky Zhao—Tsinghua University, China

## ASSOCIATE EDITORS—e-NEWSLETTER

Tamir Bendory—Tel Aviv University, Israel

Behnaz Ghorraani—Florida Atlantic University, U.S.A.

Anubha Gupta—IIIT Delhi, India

Alessio Medda—Georgia Tech Research Institute, U.S.A.

Irena Orovic—University of Montenegro, Podgorica

Sarah Ostadabbas—Northeastern University, U.S.A.

## ASSOCIATE EDITOR—SOCIAL MEDIA/OUTREACH

Guijin Wang—Tsinghua University, China

## IEEE SIGNAL PROCESSING SOCIETY

Ahmed Tewfik—President

Athina Petropulu—President-Elect

Fernando Pereira—Vice President, Conferences

Shrikanth Narayanan—VP Education

K.V.S. Hari—Vice President, Membership

Sergio Theodoridis—Vice President, Publications

Tülay Adalı—Vice President, Technical Directions

## IEEE SIGNAL PROCESSING SOCIETY STAFF

William Colacchio—Senior Manager, Publications and Education Strategy and Services

Rebecca Wollman—Publications Administrator

## IEEE PERIODICALS MAGAZINES DEPARTMENT

Jessica Welsh, *Managing Editor*

Geraldine Krolin-Taylor,  
*Senior Managing Editor*

Janet Dudar, *Senior Art Director*

Gail A. Schnitzer, *Associate Art Director*

Theresa L. Smith, *Production Coordinator*

Mark David, *Director, Business Development - Media & Advertising*

Felicia Spagnoli, *Advertising Production Manager*

Peter M. Tuohy, *Production Director*

Kevin Lisankie, *Editorial Services Director*

Dawn M. Melley, *Staff Director, Publishing Operations*

Digital Object Identifier 10.1109/MSP.2019.2963634



IEEE prohibits discrimination, harassment, and bullying.  
For more information, visit  
<http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

**SCOPE:** *IEEE Signal Processing Magazine* publishes tutorial-style articles on signal processing research and applications as well as columns and forums on issues of interest. Its coverage ranges from fundamental principles to practical implementation, reflecting the multidimensional facets of interests and concerns of the community. Its mission is to bring up-to-date, emerging and active technical developments, issues, and events to the research, educational, and professional communities. It is also the main Society communication platform addressing important issues concerning all members.



## What Does an Editor-in-Chief of *IEEE Signal Processing Magazine* Do, Anyway?

**A**fter more than two years, my term as editor-in-chief (EIC) of *IEEE Signal Processing Magazine (SPM)* is coming to an end in December 2020. I am looking forward to finishing out this year with a bang. Soon, the search will start for my replacement. In this editorial, I will provide a summary of the EIC's job from my perspective, in case you are interested in applying.

I break down my description in three areas. The first area is the initial team-building phase. The second area is the operations phase. In parallel with the first two, the third area is strategic planning. I am grateful to *SPM*'s past EIC Prof. Min Wu for helping me to undertake this difficult task.

Typically, the EIC starts six months prior to taking over in January 2021 (as will be the case for the new EIC). During that time, the EIC-in-waiting learns about standard operating procedures, ScholarOne tricks, the functions of the different people associated with *SPM* [area editors, editorial board, guest editors, associate editors, IEEE Signal Processing Society (SPS) staff, and IEEE editorial staff]. For me, not having served *SPM* as an area editor, much of these functions were new. As I learned, *SPM* operates very differently than a typical IEEE transaction with a team structure to match its unique features.

The most important task during the EIC-in-waiting phase is to identify area editors for all of the areas: columns and

forums, feature articles, special issues, special initiatives, social media, and the eNewsletter. I was extremely lucky to have such a high-quality team. As the magazine has so many different pieces, having people who are proactive in their areas is a must. It is important to recruit a diverse group of people to bring as many different perspectives as possible.

Once a team is in place, the next main task is to set up the steady-state operations. The operations are built around bringing content to six published issues a year. The EIC has two specific action items for every issue:

**It is important to recruit a diverse group of people to bring as many different perspectives as possible.**

writing the editorial and providing input on the cover design. I find that writing an editorial can take as little or as much time as you can provide. I was always impressed with SPS Past-President Ali H. Sayed's "President's Messages." They were well researched and touched on important issues including the importance of diversity [1] and openness in science [2]. I addressed some issues in my editorials as well, though admittedly without the same level of research, including "Going Toward 6G" [3], "Vehicular Applications of Signal Processing" [4], and "Research Gems Found Digging With Industry" [5]. I also covered more pragmatic aspects of the magazine including results from the annual board meetings [6] and provided some guidance on feature articles [7] and special issue submissions [8]. If I could do it over, I would try to make my editorials a

bit longer and try to address an important topic for the Society in every one.

Reviewing and approving the cover of an issue of *SPM* is a surprising aspect of the EIC's job. This is a collaborative effort with the IEEE staff led by Managing Editor Jessica Welsh and the IEEE Magazine Department's design team. For a special issue article, we look for images connected to the overall theme of the issue. We may try to include the lead guest

editor in the process, to make sure that any images he or she provides for consideration are royalty-free and have all of

the necessary (and required) permission rights. The process involves finding an image that can be licensed by the IEEE from an online stock image marketplace such as istockphoto.com. Although it depends on the specific chosen image for a cover, the IEEE Magazine Department's design team has the ability to resize images, use specific pieces of an image, change colors, and make overall changes.

When we have a collection of feature articles, as in this issue, we base the cover design around the theme of one of the articles. In this case, we were actually able to build a cover using the molecule visualizations produced by the authors of that article [9]. If a future EIC has more design experience and/or a background in that area, he or she could potentially offer more input on the cover in the front end of the process by supplying

keywords to the managing editor of the magazine, who can then work with the design team and create a conceptual cover image that will grab the reader's attention, encouraging them to pick up and read the magazine. The editorial staff reviews the highly technical text that could be used on the cover and suggests creative enhancements in wording that may create more action-oriented text that interests the reader. This collaborative effort produces a well-crafted cover utilizing the EIC for his or her technical knowledge and magazine staff for their creative input.

The steady state of the EIC's job also involves encouraging the submission of content. This includes reaching out to potential authors based on current trends. A common source for feature article authors is tutorial presentations at ICASSP, ICIP, and other SPS conferences. Similarly, for special issues, one or more guest editors may be solicited from special session organizers at ICASSP, or the Asilomar conference, for example. We may encourage submissions of columns in a similar fashion. I want to emphasize that we welcome submissions from everyone. We would love to receive more proactive contributions to the magazine.

There are also many small tasks that an EIC handles. For example, it is common for me to review the special issue and feature article proposals. I

also spend some time monitoring an article's status in ScholarOne. I interact with the magazine's area editors, relying on them to take the lead in their areas, with no micromanagement on my part. Of course, a future EIC may be as involved as he or she wants.

The third area of contribution for the magazine is to articulate and define

**It is rewarding to see a complete issue that has an interesting cover, high-quality content, and topical themes of wide interest to academia and industry.**

a vision. This is the EIC's opportunity to put their own special touch on the magazine. In my case, I had outlined three strategic directions: industrial content, commercialization, and humor. I have been most successful

with the humor initiative, coauthoring several cartoons with Nuria González Prelcic. I am still working on improving industrial content and commercialization. In other words, I give myself a failing grade in those areas. But 2020 is going to be a great year, and there is still time. While the impacts of these efforts may not be felt during my term, I think that it will propagate through future years of the magazine.

Is the EIC job fun? Perhaps the most honest answer is that it is not a relevant question. Being an EIC, area editor, or associate editor for any publication is considered a *service*. That label is usually applied to tasks that are not necessarily fun for the server, but are of great broader benefit, in this case to the SPS. Service is a way to make a difference in

the SPS. Given that *SPM* is sent to all SPS members, it is a platform for influencing directions of the Society. Most of all, the EIC's job is satisfying. It is rewarding to see a complete issue that has an interesting cover, high-quality content, and topical themes of wide interest to academia and industry. So, with that, please consider applying to be the next *SPM* EIC!

## References

- [1] A. H. Sayed, "Science is blind," *IEEE Signal Process. Mag.*, vol. 35, no. 5, pp. 4–6, 2018. doi: 10.1109/MSP.2018.2849043.
- [2] A. H. Sayed, "Galileo, Fourier, and openness in science," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 6–8, 2018. doi: 10.1109/MSP.2018.2809224.
- [3] R. W. Heath Jr., "Going toward 6G," *IEEE Signal Process. Mag.*, vol. 36, no. 3, pp. 3–4, 2019. doi: 10.1109/MSP.2019.2899083.
- [4] R. W. Heath Jr., "Vehicular applications of signal processing," *IEEE Signal Process. Mag.*, vol. 36, no. 1, pp. 3–6, 2019. doi: 10.1109/MSP.2018.2880687.
- [5] R. W. Heath Jr., "Research gems found digging with industry," *IEEE Signal Process. Mag.*, vol. 35, no. 2, pp. 4–18, 2018. doi: 10.1109/MSP.2018.2793165.
- [6] R. W. Heath Jr., "Highlights from the IEEE *SPM*'s editorial board meeting," *IEEE Signal Process. Mag.*, vol. 35, no. 4, pp. 3–4, 2018. doi: 10.1109/MSP.2018.2834778.
- [7] R. W. Heath Jr., "Making a good feature article submission," *IEEE Signal Process. Mag.*, vol. 36, no. 5, pp. 3–4, 2019. doi: 10.1109/MSP.2019.2918010.
- [8] R. W. Heath Jr., "Organizing a special issue of *IEEE SPM*," *IEEE Signal Process. Mag.*, vol. 36, no. 6, pp. 3–4, 2019. doi: 10.1109/MSP.2019.2938152.
- [9] T. Bendory, A. Bartsaghi, and A. Singer, "Single-particle cryo-electron microscopy," *IEEE Signal Process. Mag.*, vol. 37, no. 2, pp. 58–76, Mar. 2020.



**SP**

## SPS Fellows and Award Winners Recognized

In this column of *IEEE Signal Processing Magazine*, 37 IEEE Signal Processing Society (SPS) members are recognized as IEEE Fellows and award recipients are announced.

### 37 SPS members elevated to IEEE Fellow

Each year, the IEEE Board of Directors confers the grade of Fellow on up to one-tenth of 1% of the voting Members. To qualify for consideration, an individual must have been a Member, normally for five years or more, and a Senior Member at the time of nomination to Fellow. The grade of Fellow recognizes unusual distinction in the IEEE's designated fields.

The SPS congratulates the following 37 SPS members who were recognized with the grade of Fellow as of 1 January 2020.

*Paavo Alku*, for contributions to analysis, synthesis, and quality improvement of speech signals.

*Adel Belouchrani*, for contributions to blind source separation and to non-stationary signal and array processing.

*Sen-Ching Cheung*, for contributions to multimedia data processing and their applications in autism interventions.

*Chong-Yung Chi*, for contributions to convex analysis and optimization for blind source separation.

*Wan Choi*, for contributions to the analysis and design of multicell communication systems.

*Zhiguo Ding*, for contributions to nonorthogonal multiple access and energy harvesting communication.

*Octavia Dobre*, for contributions to the theory and practice of signal intelligence and emerging wireless technologies.

*Touradj Ebrahimi*, for contributions to visual information representation and assessment of quality of experience in multimedia.

*Sina Farsiu*, for contributions to multiframe superresolution and ophthalmic image processing.

*Feifei Gao*, for contributions to channel estimation and signal processing for wireless communications.

*Guofei Gu*, for contributions to malware detection and security of next-generation networks.

*Reinhold Haeb-Umbach*, for contributions to robustness of automatic speech recognition.

*Mark Hasegawa-Johnson*, for contributions to speech processing of underresourced languages.

*Daesik Hong*, for contributions to wireless and cellular communication technologies.

*Eduard Jorswieck*, for contributions to resource allocation in wireless interference networks.

*Markku Juntti*, for contributions to multiuser and multiantenna communications.

*Jeff Krolík*, for contributions to statistical signal and sensor array processing for radar and sonar.

*Jia Li*, for contributions to real-time automatic image annotation and image retrieval.

*Richard Lippmann*, for contributions to neural networks and assessment of computer security systems.

*Meinard Mueller*, for contributions to music signal processing.

*Robert Murphy*, for contributions to machine-learning algorithms for biological images.

*Patrick Naylor*, for contributions to signal processing for speech dereverberation and analysis.

*Moshe Nazarathy*, for contributions to optical systems and analog fiber-optic transmission.

*Ana Perez-Neira*, for contributions to signal processing for satellite communications and systems.

*Markus Pueschel*, for contributions to the implementation of signal processing techniques.

*Christ Richmond*, for contributions to adaptive array processing algorithms.

*Tanja Schultz*, for contributions to multilingual speech recognition and bio-signal processing.

*Gokhan Tur*, for leadership in spoken language understanding and applications to virtual personal assistant products.

*Shahrokh Valaee*, for contributions to localization of wireless nodes.

*Dimitri Van De Ville*, for contributions to image processing for computational brain imaging.



*Ba-Ngu Vo*, for contributions to the theory of multiple-object tracking and estimation.

*Aaron Wagner*, for contributions to distributed data compression.

*Yunhong Wang*, for contributions to iris and face recognition.

*Yonggang Wen*, for contributions to cloud systems for multimedia signal processing and communications.

*Pengfei Xia*, for contributions to multiple-input, multiple-output millimeter-wave wireless communications.

*Jong Chul Ye*, for contributions to signal processing and machine learning for biomedical imaging.

*Bowen Zhou*, for leadership in human language technologies.

The following individuals were evaluated by the SPS but are not SPS members.

*Joel Tropp*, for contributions to sparse signal processing.

*Shaohua Zhou*, for contributions to image analysis for medical imaging and face recognition.

## 2019 IEEE SPS Awards presented in Barcelona, Spain

The IEEE SPS congratulates the following SPS members who will receive the Society's prestigious awards during IEEE ICASSP 2020 in Barcelona, Spain, 4–8 May.



Georgios B. Giannakis

The Norbert Wiener Society Award honors outstanding technical contributions in a field within the scope of the IEEE SPS and outstanding leadership within that field. It comprises a plaque, a certificate, and a monetary award of US\$2,500. It is the highest-level award bestowed by the IEEE SPS. This year's recipient is Georgios B. Giannakis, "for fundamental contributions to statistical signal processing, especially for networking and communications, and for outstanding mentoring of young researchers."

The Claude Shannon–Harry Nyquist Technical Achievement Award is presented this year to C.-C. Jay Kuo "for significant contributions to visual signal processing technologies and their



C.-C. Jay Kuo

technical contributions to theory and/or practice in technical areas within the scope of the Society, as demonstrated by publications, patents, or recognized impact on this field. The prize for the award is US\$1,500, a plaque, and a certificate.



Thrasyvoulos N. Pappas



Helen Meng



Peter Stoica

applications." The Claude Shannon–Harry Nyquist Technical Achievement Award honors a person who, over a period of years, has made outstanding technical contributions to theory and/or practice in technical areas within the scope of the Society, as demonstrated by publications, patents, or recognized impact on this field. The prize for the award is US\$1,500, a plaque, and a certificate.

The Leo L. Beranek Meritorious Service Award was presented this year to Thrasyvoulos N. Pappas and Helen Meng, "for exemplary service to and leadership in the Signal Processing Society." The award comprises a plaque and a certificate; judging is based on dedication, effort, and contributions to the Society.

The Carl Friedrich Gauss Education Award honors educators who have made pioneering and significant contributions to signal processing education. Judging is based on a career of meritorious achievement in signal processing education as exemplified by writing of scholarly books and texts, course materials, and papers on education; inspirational and innovative teaching; creativity in the development of new curricula and methodology. The award comprises a plaque, a monetary award of US\$1,500, and a certificate. The recipient of the Carl Friedrich Gauss Education Award is Peter Stoica, "for lifetime contributions to education and mentoring in statistical signal processing."

The Industrial Leader Award recognizes an industry business or technical



Li Deng

leader whose leadership has resulted in major and outstanding advances or new directions using signal processing technologies within the scope of the Society.

This award is for executive leadership resulting in major advances and new directions using signal processing in a business area. The prize is US\$1,500, a plaque, and a certificate. The recipient of the Industrial Leader Award is Li Deng "for leadership in pioneering research and development on large-scale deep learning that disrupted worldwide speech recognition industry and for leadership in natural language processing and finance engineering."



Gene Alan Frantz

The Industrial Innovation Award is presented this year to Gene Alan Frantz "for contributions to the development of signal processing systems." The Industrial Innovation Award recognizes an individual or team at any level who were industry employees whose technical contributions have resulted in significant advances using signal processing technologies within the scope of the Society. Selection will be based on major industrial accomplishments, standards, the deployment of important processes or products, etc., that are of substantial benefit to the public, which use signal processing technologies, and are visible beyond the company or institution where the contribution was made. The award is open to individuals at any level who were industry employees who played a significant role in the technical contribution at the time of the accomplishments being recognized. The prize is US\$1,500 per awardee (up to a maximum of US\$4,500 per award), and a plaque, and certificate.

The Meritorious Regional/Chapter Service Award honors the outstanding contributions of any member of the Society to regional activities of the SPS. Judging is based on dedication, effort, and contributions made to activities aimed at promoting the technical and educational





Mohammad Faizal  
Ahmad Fauzi



Syed Abd Rahman  
Abu-Bakar



Yuejie Chi

activities of the SPS in one specific Region/Chapter as well as its local membership participation. The award comprises a plaque and a certificate. The recipients of the Meritorious Regional/Chapter Service Award are Mohammad Faizal Ahmad Fauzi, “for exemplary leadership and strong commitment toward the growth of the SPS Chapter in Malaysia” and Syed Abd Rahman Abu-Bakar, “for outstanding service to the Signal Processing Society Malaysian Chapter and Region 10.”

The Early Career Technical Achievement Award honors an individual who, over a period of years in his/her early career, has made significant technical contributions to theory and/or practice in technical areas within the scope of the Society, as demonstrated by publications, patents, or recognized impact on the field including, but not limited to, a standard, a product, or a technology trend. Nominees shall be judged on the basis of their significant technical contributions to theory and/or practice in technical areas within the scope of the Society, as demonstrated by publications, patents, or recognized impact on the field including, but not limited to, a standard, a product, or a technology trend. The prize is US\$1,500, a plaque, and a certificate. The recipient of the Early Career Technical Achievement Award is Yuejie Chi, “for contributions to high-dimensional structured signal processing.”

The IEEE Signal Processing Magazine Best Paper Award honors the author(s) of an article of exceptional merit and broad interest on a subject related to the Society’s technical scope and appearing in the Society’s magazine. The prize comprises US\$500 per

author (up to a maximum of US\$1,500 per award) and a certificate. In the event that there are more than three authors, the maximum prize is divided equally among all authors, and each receives a certificate. This year, the IEEE Signal Processing Magazine Best Paper Award recipients are Emmanuel J. Candès and Michael B. Wakin for their article “An Introduction to Compressive Sampling: A Sensing/Sampling Paradigm That Goes Against the Common Knowledge in Data Acquisition,” published in *IEEE Signal Processing Magazine*, vol. 25, no. 2, March 2008.

The IEEE Signal Processing Magazine Best Column Award honors the author(s) of a column of exceptional merit and broad interest on a subject related to the Society’s technical scope and appearing in the Society’s magazine. The prize consists of US\$500 per author

(up to a maximum of US\$1,500 per award) and a certificate. In the event that there are more than three authors, the maximum prize is divided equally among all authors, and each receives a certificate. This year, the IEEE Signal Processing Magazine Best Column Award recipients are Emil Björnson, Mats Bengtsson, and Björn Ottersten for “Optimal Multiuser Transmit Beamforming: A Difficult Problem With a Simple Solution Structure,” published in *IEEE Signal Processing Magazine*, vol. 31, no. 4, July 2014.

The IEEE Signal Processing Letters Best Paper Award honors the author(s) of a letter article of exceptional merit and broad interest on a subject related to the Society’s technical scope and appearing in *IEEE Signal Processing Letters*. The prize consists of US\$500 per author (up to a maximum of US\$1,500 per award) and a certificate. To be eligible for consideration, an article must have appeared in *IEEE Signal Processing Letters* in an issue predating the spring awards board meeting by five years (typically held in conjunction with ICASSP). Judging shall be on the basis

of the technical novelty, the research significance of the work, and quality and effectiveness in presenting subjects in an area of high impact to the Society’s members. The recipients of the IEEE Signal Processing Letters Best Paper Award are Stelios Timotheou and Ioannis Krikidis for “Fairness for Non-Orthogonal Multiple Access in 5G Systems,” *IEEE Signal Processing Letters*, vol. 22, no. 10, October 2015.

The Sustained Impact Paper Award honors the author(s) of a journal article of broad interest that has had sustained impact over many years on a subject related to the Society’s technical scope. The prize consists of US\$500 per

author (up to a maximum of US\$1,500 per award) and a certificate. In the event that there are more than three authors, the maximum prize is divided equally among all authors, and each receives a

certificate. To be eligible for consideration, an article must have appeared in one of the IEEE SPS transactions or *IEEE Journal of Selected Topics in Signal Processing*, in an issue predating the spring awards board meeting by at least 10 years (typically held in conjunction with ICASSP). This year, the Sustained Impact Paper Award recipients are Hamid Krim and Mats Viberg for “Two Decades of Array Signal Processing Research: The Parametric Approach,” *IEEE Signal Processing Magazine*, vol. 13, no. 4, July 1996.

Six Best Paper Awards were presented, honoring the author(s) of a paper of exceptional merit dealing with a subject related to the Society’s technical scope and appearing in one of the Society’s transactions, irrespective of the author’s age. The prize is US\$500 per author (up to a maximum of US\$1,500 per award) and a certificate. Eligibility is based on a five-year window preceding the year of election, and judging is based on general quality, originality, subject matter, and timeliness. Up to six Best Paper Awards may be presented each year. This year, the awardees are

**The Carl Friedrich Gauss  
Education Award honors  
educators who have made  
pioneering and significant  
contributions to signal  
processing education.**

- Martin Sundermeyer, Hermann Ney, and Ralf Schlüter, “From Feedforward to Recurrent LSTM Neural Networks for Language Modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 3, March 2015.
- Donggeek Shin, Ahmed Kirmani, Vivek K Goyal, and Jeffrey H. Shapiro, “Photon-Efficient Computational 3-D and Reflectivity Imaging With Single-Photon Detectors,” *IEEE Transactions on Computational Imaging*, vol. 1, no. 2, June 2015.
- Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser, “Deep Convolutional Neural Network for Inverse Problems in Imaging,” *IEEE Transactions on Image Processing*, vol. 26, no. 9, September 2017.
- Yuxuan Wang, Arun Narayanan, and DeLiang Wang, “On Training Targets for Supervised Speech Separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, December 2014.
- Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi, “Phase Retrieval Using Alternating Minimization,” *IEEE Transactions on Signal Processing*, vol. 63, no. 18, September 2015.
- Chanwoo Kim and Richard M. Stern, “Power-Normalized Cepstral Coefficients (PNCC) for Robust Speech Recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 7, July 2016.

The Young Author Best Paper Award honors the author(s) of an especially meritorious paper dealing with a subject related to the Society’s technical scope and appearing in one of the Society’s transactions and who, upon date of submission of the paper, is younger than 30 years of age. The prize is US\$500 per author (up to a maximum of US\$1,500 per award) and a certificate. Eligibility is based on a three-year window preceding the year of election, and judging is based on general quality, originality, subject matter, and timeli-

ness. Three Young Author Best Paper Awards are being presented this year:

- Daichi Kitamura, for the paper coauthored with Nobutaka Ono, Hiroshi Sawada, Hirokazu Kameoka, and Hiroshi Saruwatari, “Determined Blind Source Separation Unifying Independent Vector Analysis and Nonnegative Matrix Factorization,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 9, September 2016.
- Sebastian Dörner and Sebastian Cammerer, for the paper coauthored with Jakob Hoydis and Stephan ten Brink, “Deep Learning Based Communication Over the Air,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, February 2018.
- Mianzhi Wang, for the paper coauthored with Arye Nehorai, “Coarrays, MUSIC, and the Cramer–Rao Bound,” *IEEE Transactions on Signal Processing*, vol. 65, no. 4, February 2017.

### 2019 Chapter of the Year Award

The IEEE SPS Hyderabad Chapter has been selected as the recipient of the 2019 Chapter of the Year Award, which will be presented at the ICASSP 2020 awards ceremony in Barcelona, Spain. The award is presented annually to a Chapter that has provided its membership with the highest quality of programs, activities, and services. The SPS Hyderabad Chapter will receive a plaque and a check in the amount of US\$1,000 to support local Chapter activities. The Chapter will publish an article in a future issue of *IEEE Inside Signal Processing eNewsletter*.

### SPS members receive 2020 IEEE awards

The IEEE has announced the recipients of IEEE Technical Field Awards for contributions or leadership in specific fields of interest of the IEEE. The following SPS members have received IEEE Technical Field Awards this year.

The IEEE James L. Flanagan Speech and Audio Processing Technical Field Award will be presented to Hynek Hermansky “for contributions to speech processing and feature extraction for robust speech recognition.”

The IEEE Fourier Award for Signal Processing will be presented to Alfred O. Hero III “for contributions to the foundations of statistical signal processing with applications to distributed sensing and performance benchmarking.”

The IEEE Electronics Packaging Technology Award will be presented to Mitsumasa Koyanagi and Peter Ramm “for pioneering contributions leading to the commercialization of 3D wafer and die-level stacking packaging.”

In addition, the IEEE has announced the recipients of the 2020 IEEE medals. IEEE medals are the highest honor of awards presented by the IEEE. The medals will be presented at the 2020 IEEE Honors Ceremony. Three SPS members have been awarded with the IEEE medals for 2020.

The IEEE Jack S. Kilby Signal Processing Medal awarded for outstanding achievements in signal processing, will be presented to Ramalingam Chellappa (University of Maryland) “for contributions to image and video processing, especially applications to face recognition.”

The IEEE James H. Mulligan Jr. Education Medal, awarded for outstanding contributions to education, will be presented to Leah Jamieson (Purdue University) “for contributions to the promotion, innovation, and inclusivity of engineering education.”

The IEEE Dennis J. Picard Medal for Radar Technologies and Applications, awarded for outstanding accomplishments in advancing the fields of radar technologies and their applications, will be presented to Joseph R. Guerci (Information Systems Laboratories, Inc.) “for contributions to advanced, fully adaptive radar systems and real-time knowledge-aided, and cognitive radar processing architectures.”



Paavo  
Alku



Adel  
Belouchrani



Sen-Ching  
Cheung



Chong-Yung  
Chi



Wan  
Choi




Zhiguo  
Ding



Octavia  
Dobre



Touradj  
Ebrahimi



Sina  
Farsiu



Feifei  
Gao



Guofei  
Gu



Reinhold  
Haeb-Umbach



Mark  
Hasegawa-Johnson



Daesik  
Hong



Eduard  
Jorswieck



Markku  
Juntti



Jeff  
Krolik



Jia  
Li



Richard  
Lippmann



Meinard  
Mueller




Robert  
Murphy

## Congratulations!

*SPS Members and Affiliates Elected IEEE Fellow Class 2020*



Patrick  
Naylor



Moshe  
Nazarathy



Ana  
Perez-Neira



Markus  
Pueschel



Christ  
Richmond



Tanja  
Schultz



Gokhan  
Tur



Shahrokh  
Valaee



Dimitri  
Van De Ville



Ba-Ngu  
Vo



Aaron  
Wagner



Yunhong  
Wang



Yonggang  
Wen



Pengfei  
Xia



Jong  
Chul Ye



Bowen  
Zhou



Joel  
Tropp



Shaohua  
Zhou



# Robotics Rolls Into High Gear With Signal Processing

*A robotics revolution promises to transform global industries and services, and signal processing is at the forefront*

A new generation of exponentially more intelligent and capable robots is on the way, helped along by talented and imaginative engineers and heavy doses of signal processing. In fields spanning almost every aspect of human professional and personal life, robots are ready to perform tasks faster, better, and more efficiently than their human counterparts. Even major sports organizations are now looking into the possibility of replacing human referees and umpires with robot arbiters.

Although robots have handled rote production-line tasks for decades, new technologies and approaches are taking robotics to the next level in fields ranging from agriculture to manufacturing to emergency services. “As robots get cheaper, more flexible, and more autonomous, they will power growth around the world,” observes Rob Atkinson, author of “Robotics and the Future of Production and Work,” a report issued by the Information Technology and Innovation Foundation [1]. “The global economy needs to increase productivity, and robotics is poised to do just that,” he adds.

## In the field

A compact agricultural robot developed by researchers at the University of Illinois at Urbana-Champaign promises to transform data collection and field scouting for agronomists, seed provid-

ers, and farmers. The TerraSentia crop phenotyping robot autonomously moves between crop rows, measuring the traits of individual plants with a camera and an array of sensors, transmitting data in real time to its operator’s mobile phone or laptop computer. A custom app and tablet computer, provided with the robot, allow the operator to steer the unit with the assistance of GPS and virtual reality technology.

TerraSentia can be taught to detect and identify common diseases as well as measure a variety of traits, such as plant height, leaf area, and biomass, says Girish Chowdhary, the University of Illinois agricultural and biological engineering professor who is leading a team of students, engineers, and postdoctoral researchers in developing the robot as well as supporting machine-learning teaching algorithms (Figure 1). The compact and portable robot is light enough to roll over tender young plants without damaging them and is designed to allow an agronomist to toss it onto a truck seat or into a car trunk for convenient transport to a farm or field.

TerraSentia’s automated data collection and analytical capabilities promise to improve plant breeding by helping researchers understand why plant varieties respond to environmental conditions in very different ways. As a result, data gathered by the crop-scouting robot could help plant breeders identify genetic lineages that are most likely to produce the best quality and highest yields

in specific locations. “Our research focuses on creating autonomous robots that solve complex problems in challenging environments,” Chowdhary explains. “[Facilitating] reliable perception of the environment and estimation of the robot’s own state in the context of the environment for enabling autonomous decision making are some of the key challenges my group works on.” The researchers are also focusing on adaptation and learning in the area of robot perception and control, topics that are of interest to the signal processing community, he notes.

Robots, such as TerraSentia, promise to help growers measure every plant in a field rapidly and efficiently without having to hire an army of people. Robotic technology also eliminates the possibility of human error, allowing plants to be measured accurately and nonsubjectively.

Signal processing is integral to TerraSentia’s autonomous operation and to ensuring fast and reliable onsite data processing. “Onboard, the robot signal processing algorithms are working hard to filter noise from inertial sensors and fuse different sensor data to create a reliable estimate of robot pose using Kalman filters,” Chowdhary states. Signal processing also plays a key role in determining distance and angle relative to the rows of plants and ensuring that the data streaming in from sensors is of sufficient quality to support decision making. After the data have been uploaded into the cloud, machine-learning and

signal processing algorithms are used to find patterns that can be used to estimate traits, including plant height, plant stem width, plant leaf area index, and object count.

Signal processing also helped the researchers address signal quality issues that often affect robots operating in field environments. “For example, GPS behaves differently when the robot is under the plant canopy, in a ditch, or in open areas with a clear view of the sky,” Chowdhary explains. Since agricultural data generated by cameras tends to be noisy and unstructured, the researchers turned to machine-learning- and deep-learning-based tools to provide detailed imaging capabilities. “We found that it was difficult to come up with filters based on traditional machine vision techniques,” he says.

The team is now looking forward to building rugged, high-quality robots while keeping costs low. This means tightly integrating systems and software. “Simply making a rugged robot that can withstand all weather, from the wet monsoons of the tropics to the dry heat of southern Australia, is a great challenge,” Chowdhary observes. “Eventually, we want these robots to operate in field environments for months at a time with little supervision.”

TerraSentia is expected to become available for commercial agricultural use within a few years, with some models possibly costing lower than US\$5,000. “We’re looking to scale up manufacturing of the robots and increase the level of autonomy,” Chowdhary says. He notes that the team’s start-up company, EarthSense, produced approximately 50 robots in 2019 and is planning to scale up operations to manufacture hundreds more.

## Building tools

Engineers at the Georgia Institute of Technology’s Robot Autonomy and Interactive Learning research lab, led by Associate Prof. Sonia Chernova, have developed a training method that enables a robot to create basic tools by combining various objects. Inspired by the 1980s U.S. television series *MacGyver*, in which the title character was known for his unconventional problem-solving ability through the use of on-hand



**FIGURE 1.** Girish Chowdhary poses with the TerraSentia among crops. (Source: University of Illinois at Urbana-Champaign; used with permission.)

resources, the “MacGyvering” project recently trained an intelligent agent that allows a host robot to reason about the shape, function, and attachment of unrelated parts.

Project researcher Lakshmi Nair, a Ph.D. student in Georgia Tech’s School of Interactive Computing, says the achievement marks an important step toward creating intelligent agents capable of devising more advanced tools that could equip robots to provide support in hazardous and potentially life-threatening environments (Figure 2).

With the assistance of machine learning, a “MacGyvered” robot can learn how to match form to function, deciding which object shapes are likely to lead to a particular outcome. For instance, by learning that a concave bowl can hold liquids, a robot can understand how to build a spoon. Machine learning can also help a robot discover how to attach objects together to form tools capable of grasping

or piercing objects. So far, the researchers have been able to teach a robot how to successfully create hammers, spatulas, scoops, squeegees, and screwdrivers.

The research also aims to give robots the ability to discern how various types of material are likely to behave in real-world situations. Humans know intuitively, for example, that hammers are hard and strong, so it doesn’t make sense to create a hammer out of foam blocks. “We want to reach that level of reasoning in our work, which is something we’re working on now,” Nair says.

The project relies on two types of information: 3D point clouds and spectral data. “We chose to analyze these two types of signals because they provide a lot of information about objects that’s beneficial to tool construction,” Nair explains. “The robot needs information about shape, which comes from 3D point clouds, and materials, which comes from spectral data.”



**FIGURE 2.** Lakshmi Nair presents the “MacGyvering” project. (Source: Georgia Institute of Technology; used with permission.)

More specifically, the robot collects spectral data and 3D point clouds (using a red, green, blue depth camera) of all of the objects available to it. The collected data are then analyzed using supervised learning, a machine-learning technique that allows the robot to predict an object’s properties. Based on the predictions, the robot can decide which parts need to be selected and combined to construct a specific tool. “The input to our algorithm is a set of objects available to the robot and some action the robot needs to perform,” Nair says. “The output is a tool constructed by combining some of the available objects, for example, combining pliers and a bowl to create a makeshift scoop.”

Nair states that the research offers tremendous potential for improving robotic adaptability and resourcefulness. “Often, robots are preprogrammed to handle a variety of different situations,” she observes. “This is very tedious, and there are often many unprecedented scenarios that may arise in the real world.” She notes that the team’s research will help robots easily adapt to unpredictable situations. “Currently, our work is one of the first to demonstrate tool construction on a physical robot and is, therefore, very novel.”

The research was inspired by an incident onboard the ill-fated *Apollo 13* spacecraft, when astronauts facing a life-or-death emergency had to “MacGyver” a carbon dioxide filter using only

available parts. “Their ingenuity helped them combat increasing carbon dioxide levels in the spacecraft, enabling them to return home safely,” Nair explains. “Our applications include similarly stressful scenarios often encountered in households, space exploration, and search–rescue missions.”

Nair notes that dealing with 3D point clouds and spectral data noise was the biggest challenge the team faced. “Sensor noise is inevitable in our robot application, and, to overcome this, we had to collect tremendous amounts of data to train our models and make our system more robust.”

Improving robot manipulation skills remains an ongoing challenge. “It’s still very difficult for the robot to manipulate objects, such as picking up and joining/attaching parts,” Nair observes. “This is important for tool construction, and we had to organize the robot workspace by placing different objects far apart so that the robot can easily manipulate them.”

Nair acknowledges that more work remains to be done before “MacGyvering” techniques can be applied to commercial robots. “Our system is currently only able to construct simple tools by combining a small number of parts,” she notes. “Further, a lot of tool constructions require bimanual, two-handed manipulation, which is a challenge for robots even today.”

Meanwhile, researchers are looking forward to the day when a robot will

be able to build highly complex structures rather than just simple tools. “We would like the robot to reason about additional properties of objects, such as mass and density, when constructing tools,” Nair says.

## Swarming robots

Inspired by the navigational accuracy and efficiency of swarming insects, a research team with members drawn from the Delft University of Technology, University of Liverpool, and Radboud University of Nijmegen, The Netherlands, has developed a swarm of tiny autonomous drones capable of exploring unknown environments. The achievement marks a significant step forward in swarm robotics, notes principal investigator Guido de Croon, since the project’s drones are capable of navigating autonomously despite possessing only limited sensing and computational capabilities.

The drones’ light weight—only 33 g each—allows them to fly safely around objects and people, and their small size enables them to fly into and around narrow spaces (Figure 3). “One potential application is in search-and-rescue scenarios where the drones have to explore a building that may be damaged and about to collapse,” de Croon says.

The drones don’t need to create a map of their environment, an advantage that’s designed to conserve onboard memory and storage resources. Instead, the units travel in a preferred direction and deal with obstacles on the fly. “When exploring the environment, all of the drones fly in different, predetermined directions that make them spread out over the environment,” de Croon explains. When the drones’ batteries reach a 60% charge level, the devices begin flying back to their home base, guided by a wireless beacon. “The preferred direction during the inbound flight is the direction that will increase the received signal strength of the home beacon the most,” he says. “With this approach, the drones can navigate autonomously with almost no memory and computational requirements.”

Signal processing is essential to any autonomous robot project,



de Croon states. However, the researchers realized that, in their project, critical signal processing functions associated with complex navigation and sensing tasks would be restricted by the drones' limited onboard power, sensing, and computational capabilities. "We had to make some tough choices, and we ended up putting more priority on the individual drone's ability to avoid obstacles and navigate than on intradrone avoidance," explains Kimberly McGuire, a student researcher who participated in the project.

Most of the data generated by each drone's internal navigation devices, including an accelerometer, gyroscope, downward-pointing laser, and optical flow sensor, are processed by a software stack running an extended Kalman filter. The output is used to determine the drone's current state in terms of attitude and velocity—essential for safe, autonomous flight, especially in cramped indoor environments.

The project's main challenges, de Croon notes, was designing processing algorithms that would allow each drone to detect obstacles via a camera and monitor its proximity to other drones. Proximity tracking is accomplished by measuring the signal strength from the wireless communication chips integrated into each drone. Finding the beacon direction is also accomplished by determining signal strength. All of these capabilities rely on combining and filtering multiple signals. "For instance, for finding the best direction to the home beacon, the drone filters the received signal strength indicator to the home beacon over time," de Croon says.

Looking ahead, deCroon would like to improve the drones' navigational capabilities, eliminating the need to rely on a home beacon. He's also hoping to improve the drones' obstacle detection and avoidance capabilities, helping them to better avoid complex and thin obstacles, such as wires and tree and plant branches. "Finally, we would like to improve the robustness of all components to take the step toward real-world unstructured environments," he states.



**FIGURE 3.** One of the drones inspired by swarming insects. (Source: Delft University of Technology; used with permission.)

### Author

**John Edwards** (jedwards@johnedwardsmedia.com) is a technology writer based in the Phoenix, Arizona, area. Follow him on Twitter @TechJohnEdwards.

### Reference

[1] R. D. Atkinson, "Robotics and the future of production and work," Information Technology & Innovation Foundation, Washington, D.C., Oct. 15, 2019. [Online]. Available: <https://itif.org/publications/2019/10/15/robotics-and-future-production-and-work>

SP



### Training Data for Machine Learning

Magic Data Technology is an One-stop AI Data Service Solution provider. We are committed to providing a wide range of data services in the fields of automatic speech recognition (ASR), text to speech (TTS), computer vision recognition and Natural Language Processing (NLP).

#### Why us:

- ✓ Efficiency: Human-in-the-loop data processing, 300,000+ professional annotators around the world
- ✓ Pioneering: Task-segmentation process and strict project management; innovative data tool software
- ✓ Quality: Our value proposition is 97%-99% quality, speed and scale (50+ languages covered)
- ✓ Professional: Multilingual & Multidomain; 100,000+ hours of data acquisition and annotation;

#### Two types of service:

- ✓ Over 100,000-hour self-owned copyright training data sets for building AI models quickly
- ✓ Customized data service solutions, including design, collection, annotation and processing.

Website: <http://en.imagicdatatech.com>

Linkedin: <https://www.linkedin.com/company/magicdata>

Business contact: +86 10-85527250

[business@magicdatatech.com](mailto:business@magicdatatech.com)

## Top Downloads in IEEE Xplore

Each “Reader’s Choice” column focuses on a different publication of the IEEE Signal Processing Society. In this issue of *IEEE Signal Processing Magazine*, we highlight papers in *IEEE Transactions on Mobile Computing (T-MC)*.


*T-MC* is a bimonthly journal for publishing archival research results related to mobility of users, systems, data, and computing as well as issues in information organization and access, services, management, and applications. Research areas for these results include, but are not restricted to, nomadic computing, multimedia applications, mobile data and knowledge management, and mobile communication systems and networking.

This issue's "Reader's Choice" column lists the 15 *T-MC* papers most downloaded from January 2017 to October 2019. Your suggestions and comments are welcome and should be sent to Associate Editor H. Vicky Zhao (vzhao@tsinghua.edu.cn).

# RT-Fall: A Real-Time and Contactless Fall Detection System With Commodity Wi-Fi Devices

Wang, H.; Zhang, D.; Wang, Y.; Ma, J.; Wang, Y.; Li, S.

This paper proposes RT-Fall, a real-time, contactless, low-cost yet accurate indoor fall-detection system using Wi-Fi devices. First, the channel state information phase difference over two antennas is found to be more



sensitive than amplitude for activity recognition and enables reliable segmentation of fall and fall-like activities. Second, the sharp power profile decline pattern of the fall in the time-frequency domain is used for new feature extraction and accurate fall segmentation and detection.

2017

## WiFall: Device-Free Fall Detection by Wireless Networks

Wang, Y.; Wu, K.; Ni, L.M.

The correlation between different signal variations and activities is studied, and an obtrusive fall-detection system named *WiFall* is proposed. It employs physical layer channel state information

as the indicator of activities and detects human falls without hardware modification, extra environmental setup, or any wearable devices. WiFall is implemented on desktop computers with commodity 802.11n network interface controllers, and its performance is evaluated in three typical indoor scenarios with several layouts of transmitter–receiver links.

2017

## HEED: A Hybrid, Energy-Efficient, Distributed Clustering Approach for Ad Hoc Sensor Networks

*Younis, O.; Fahmy, S.*

The Hybrid Energy-Efficient Distributed (HEED) clustering protocol is

proposed for ad hoc sensor networks to balance load on sensor nodes and to increase network scalability and lifetime. It periodically selects cluster heads according to the node residual energy and a secondary parameter, e.g., node proximity to its neighbors and node degree. It incurs low message overhead and achieves fairly uniform cluster head distribution across the network. HEED is proved to asymptotically almost surely guarantee connectivity of clustered networks with appropriate bounds on node density and transmission ranges.

2004

### **EABS: An Event-Aware Backpressure Scheduling Scheme for Emergency Internet of Things**

*Qiu, T.; Qiao, R.; Wu, D.O.*

An event-aware backpressure scheduling (EABS) scheme for emergency Internet of Things is proposed. A backpressure queue model with emergency packets is derived based on the analysis of the arrival process of different packets. The proposed scheme combines the shortest path with backpressure scheme in the process of next-hop node selection. It forwards emergency packets in the shortest path and avoids network congestion according to the queue backlog difference. Experimental results verify the effectiveness of EABS in reducing the average end-to-end delay and increasing the average forwarding percentage.

2018

### **Cooperative Edge Caching in User-Centric Clustered Mobile Networks**

*Zhang, S.; He, P.; Suto, K.; Yang, P.; Zhao, L.; Shen, X.*

Delay-optimal cooperative edge caching in large-scale user-centric mobile networks is investigated. A greedy content placement algorithm is proposed based on the optimal bandwidth allocation. In addition, the optimal user-centric cluster size is studied, and the explicit form of the condition constraining the maximal cluster size is given, which reflects the tradeoff between caching diversity and spectrum efficiency. Extensive simula-



©PIXABAY/COMUEFERRIS

tions are conducted to validate the performance of the proposed method, and the impact of system parameters on clustering is also discussed.

2018

### **Received-Signal-Strength-Based Indoor Positioning Using Compressive Sensing**

*Feng, C.; Au, W.S.A.; Valaee, S.; Tan, Z.*

This paper proposes an accurate received-signal-strength-based indoor positioning system using compressive sensing. A coarse localizer first compares the received signal strength to a number of clusters to find the cluster to which it belongs. Then, a compressive sensing-based fine localization step is used to further refine the location. Different coarse localization schemes and access point selection approaches are investigated to improve the accuracy. The proposed positioning system is implemented on a Wi-Fi-integrated device to evaluate its performance.

2012

### **Edge Computing Assisted Adaptive Mobile Video Streaming**

*Mehrabi, A.; Siekkinen, M.; Ylä-Jääski, A.*

This paper presents an optimized solution for network-assisted bit-rate adaptation targeting mobile streaming in multiaccess edge computing environments. An integer nonlinear programming optimization model is proposed for video bit-rate selection to jointly maximize the quality of experience of

clients and proportional fairness of the bit-rate allocation. A near-optimal self-tuning greedy-based scheduling algorithm is designed to efficiently solve the client-to-edge-server mapping and the bit-rate selection problems.

2019

### **Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis**

*Sommer, C.; German, R.; Dressler, F.*

The need for bidirectional coupling of network simulation and traffic microsimulation for evaluation of intervehicle communication protocols is discussed. A hybrid simulation framework including the network simulator OMNeT++ and the road traffic simulator SUMO is developed, allowing dynamic interaction between both simulators. A proof-of-concept study is used to demonstrate its advantages based on the evaluation of two protocols for incident warning over vehicular ad hoc networks.

2011

### **Load Balancing Under Heavy Traffic in RPL Routing Protocol for Low Power and Lossy Networks**

*Kim, H.-S.; Kim, H.; Paek, J.; Bahk, S.*

The load-balancing and congestion problem of the standardized IPv6 routing protocol for low power and lossy networks (RPL) is studied. Experimental study of RPL shows that most packet losses under heavy traffic are due to congestion, and there exists a serious



load balancing problem in RPL in terms of routing parent selection. To address this issue, a simple yet effective queue utilization RPL (QU-RPL) is proposed, which improves the end-to-end packet delivery performance by balancing the traffic load within a routing tree. Experimental results show that QU-RPL significantly reduces the queue loss and improves the packet delivery ratio.

2017

### Smartphones Based Crowdsourcing for Indoor Localization

Wu, C.; Yang, Z.; Liu, Y.

Locating in Fingerprint Space (LiFS), a wireless indoor localization system based on off-the-shelf Wi-Fi infrastructure and mobile phones, is proposed. Considering user movements in a building, originally separated received signal strength fingerprints are geographically connected, and they consequently form a high-dimension fingerprint space. The fingerprint space is then automatically mapped to the floor plan. Preliminary experiment results show that LiFS achieves low human cost, rapid system deployment, and competitive location accuracy.

2015

### Combining Solar Energy Harvesting With Wireless Charging for Hybrid Wireless Sensor Networks

Wang, L.C.; Li, J.; Yang, Y.; Ye, F.

This paper proposes a hybrid framework that combines the advantages of wireless charging and solar energy harvesting technologies. The network is divided into three hierarchical levels. The first level studies how to deploy solar-powered cluster heads to minimize the overall cost. Second, the energy balance problem is examined, and a distributed head reselection algorithm is proposed to designate some wireless-powered nodes as cluster heads when

solar energy is unavailable. Third, a linear-time algorithm is proposed to optimize the joint tour consisting of both wireless charging and data-gathering sites for mobile chargers.

2018

### VeMAC: A TDMA-Based MAC Protocol for Reliable Broadcast in VANETs

Omar, H.A.; Zhuang, W.; Li, L.

A novel multichannel time-division-multiple-access media-access-control protocol, VeMAC, is introduced for vehicular ad hoc networks. It provides reliable one-hop broadcast service without the hidden terminal problem and offers efficient multihop broadcast service to disseminate information over the network. It assigns disjoint sets of time slots to vehicles moving in opposite directions and to road side units and reduces transmission collisions on the control channel caused by node mobility. Analysis and simulation results in highway and city scenarios show that VeMAC can provide significantly higher throughput on the control channel.

2013

### Ubii: Physical World Interaction Through Augmented Reality

Lin, S.; Cheng, H.F.; Li, W.; Huang, Z.; Hui, P.; Peylo, C.

The presented ubiquitous interface and interaction (Ubii) is an integrated interface system that connects smart devices and allows users to interact with physical objects, such as computers, projector screens, and printers, to complete tasks including document copying, printing, sharing, and so on. Individual augmented reality menu overlays are aligned with physical objects to present their physical affordance. Developed on Google Glass, Ubii enables gesture-based free-hand interaction with greater convenience. A new support vector machine-based detection algorithm and

computation offloading is implemented for better performance.

2017

### Mode Selection and Resource Allocation in Device-to-Device Communications: A Matching Game Approach

Kazmi, S.M.A.; Tran, N.H.; Saad, W.; Han, Z.; Ho, T.M.; Oo, T.Z.; Hong, C.S.

A distributed scalable solution for a dense device-to-device network is introduced, which jointly addresses the mode selection, resource allocation, and interference management problems. A novel learning framework based on Markov approximation is proposed, where unsupervised learning is used for mode selection, and a two-sided matching game is incorporated to address the resource allocation issue. Simulation results show that the proposed framework converges in probability, achieves interference protection, and closely approaches the optimal solution.

2017

### Energy-Efficient Dynamic Computation Offloading and Cooperative Task Scheduling in Mobile Cloud Computing

Guo, S.; Liu, J.; Yang, Y.; Xiao, B.; Li, Z.

An energy-efficient dynamic offloading and resource scheduling policy is proposed to reduce energy consumption and to shorten application completion time in mobile cloud computing. The problem is formulated as an energy-efficiency cost minimization problem with task-dependency requirement and completion time deadline constraint. A distributed algorithm is then proposed that consists of three subalgorithms: computation offloading selection, clock frequency control, and transmission power allocation.

2019

SP

# Call for Papers IEEE Signal Processing Magazine

## Special Issue: Innovation Starts with Education

Signal processing (SP) is at the very heart of our digital lives, owing to its role as the pivotal development technology across multiple disciplines. Its prominence in modern data science has created a necessity to supply industry, government labs, and academia with graduates who possess relevant crucial signal processing expertise and are well equipped to deal with the manifold challenges in current and future applications. To this end, both the ways to deliver the educational content and the core Signal Processing curriculum need to be revisited and integrated into Engineering and Computer Science degrees, to provide hands-on skills, experience, and inspiration for students.

Signal processing education in today's universities is largely influenced by three modern trends: 1) the availability of competing and complementary online and multimedia resources; 2) the fact that we live in a world in which the amount and diversity of information we generate, process and analyse is growing ever faster; 3) the explosive growth of computing power, and the rapid development of new technologies for implementing both analog and digital signal processing. These trends offer both opportunities and challenges, which we can, and must, exploit in charting dynamically adjustable courses that attract a high level of student engagement while offering a mix of essential background physics, intuition, mathematical rigor, and practical applicability of the taught material.

With such initiatives underway at many universities world-wide, this special issue aims to facilitate both keeping abreast with SP education and exploring innovative and participatory ways to present the educational materials. In effect, we cannot assume that students will be able to appreciate the scope and relevance of their courses without explicitly building a bridge between the material presented in class and cutting-edge research, societal, and practical impact of their education. This includes the convergence of educational material with other disciplines (machine learning, data science, big data, bioengineering, artificial intelligence, finance and many others).

**This special issue will revolve around three general and most pressing aspects** of modern SP education:

- **How to educate differently (better).** This includes the use of available technology, bringing research into the classroom, web resources, experiential learning, and massive open online courses (MOOC).
- **Student engagement**, such as ways to enhance student creativity and curiosity, student satisfaction issues, various forms of assessment and metrics, engagement of under-represented population, and outreach drives.
- **Promotion of the societal impact of SP**, including privacy, ethical and security concerns, wearable devices and eHealth, global interconnections through IoT, and impact on climate change, global economy and finance.

**Topics of interest** include but are not limited to:

- Mitigation of issues related to the perceived difficulty of traditional SP courses, such as strategies on how to teach SP with less maths and how to attract attendees from non-engineering departments.
- The use of emerging technologies and technologically orientated classroom, such as MOOC and web resources.
- Metrics for success of education delivery in the After Online Technology (AOT) era.
- Using the principles of signal processing to improve teaching and research in related areas, such as machine learning, bioengineering, artificial intelligence and optimization, and vice versa.
- Curricular changes to meet contemporary demands from industry, such as using practically relevant problems, exploring feasible extensions and new applications of the taught material, and curiosity driven learning.
- Preparing students for life-long learning, teaching life-long fundamentals of SP, and relevance of SP with respect to technological advance.
- Challenges and solutions in industry-run courses; design of short courses offered by academia for industry, Government Agencies and National Defence.
- Role of mentorship and initiatives to encourage and motivate students in research experiences.
- Promoting creativity in learning, especially when applying the concepts with "opportunity windows" to explore entrepreneurship and possible product developments, and cross-disciplinary aspects of our work.

**Submission Process.** Original submissions will be reviewed according to the guidelines as set out in the IEEE Signal Processing Magazine, and should be submitted online at <http://mc.manuscriptcentral.com/sps-ieee>. Prospective authors should initially submit short White Papers, as indicated below, which will be reviewed by Guest Editors. Authors of successful White Papers will be invited to submit full-length manuscripts, which will undergo the usual reviewing process in accordance with the schedule outlined below.

**White Papers (up to 4 pages) due: March 1, 2020**

Full length manuscripts due: June 1, 2020

Revised manuscripts due: September 15, 2020

Final manuscripts due: December 1, 2020

Decision on White Papers: April 1, 2020

Review results and decision notification: August 1, 2020

Acceptance notification: November 1, 2020

**Publication date: March 1, 2021**

### Guest Editors

- Mónica Bugallo, SUNY at Stony Brook, USA, [monica.bugallo@stonybrook.edu](mailto:monica.bugallo@stonybrook.edu)
- Anthony Constantinides, Imperial College London, [a.constantinides@imperial.ac.uk](mailto:a.constantinides@imperial.ac.uk)
- Danilo Mandic, Imperial College London, UK, [d.mandic@imperial.ac.uk](mailto:d.mandic@imperial.ac.uk)
- Alan Oppenheim, MIT, USA, [avo@mit.edu](mailto:avo@mit.edu)
- Roberto Togneri, University of Western Australia, [roberto.togneri@uwa.edu.au](mailto:roberto.togneri@uwa.edu.au)

Please contact Danilo Mandic for any initial questions concerning this Special Issue.

# Cooperative Wireless Mobile Caching



©ISTOCKPHOTO.COM/TONY STUDIO

## *A Signal Processing Perspective*

Soheil Mohajer, Itsik Bergel,  
and Giuseppe Caire

**C**ache-aided communications have shown potential for substantial improvement in network performance, which goes far beyond that of traditional caching. *Traditional caching* (i.e., the bringing and storing of data closer to the end users) is only efficient when a significant portion of the popular files can be locally stored. In cache-aided communications, however, information stored at one user is useful for interference mitigation even if it is requested only by another user. The core idea in cache-aided communication is to use this interference-cancellation opportunity to simultaneously serve multiple users by sending a sum of multiple packets. By creating opportunities for multicasting, the improved performance scales with the accumulated cache size at all users. This is a great advantage for modern networks, where the number of users is typically large, and a small amount of memory can easily be allocated at each user. This article presents the novel techniques of cache-aided communications while focusing on the signal processing aspects that lie in the heart of these schemes. In particular, we examine

the three well-studied signal processing problems at the core of cache-aided communications: resource allocation, beamforming design, and interference mitigation.

### **Introduction**

The increased demand for large files (e.g., media) has produced overwhelming network traffic. The nature of the data has shifted from voice and short messages to large files. It is expected that by 2020, 75% of total mobile data traffic will be attributable to video. The characteristics of this type of data can be exploited to improve the performance of data delivery networks. In particular, popular videos are typically repeatedly requested by multiple users in an asynchronous manner. Moreover, prime time is usually associated with videos, i.e., the demand peaks during certain hours of the day, rather than being uniformly distributed over time.

*Caching*, bringing the data closer to where they will be used and storing them locally, is an efficient approach used to exploit the characteristics of such large-size contents to reduce the network traffic. Currently, caching is being used for data delivery systems, e.g., Netflix and Facebook's photo

Digital Object Identifier 10.1109/MSP.2019.2962507  
Date of current version: 26 February 2020



caching. However, these systems work based on predicting content with high demand and storing popular files on local storage units close to the end users. The gain provided by such strategies is limited by the size of the local memory and the prediction accuracy. In other words, if the individual local storage units are not large enough to store a significant portion of the popular files, cached data will be almost useless (when users request other files) and the gain of caching will be negligible.

Despite the dramatic drop in the cost of memory, the storage size on individual fixed or mobile devices is still too small to store a substantial fraction of the popular data. However, mobile devices have become more popular and the number of mobile users is rapidly growing. This leads to a substantial amount of aggregate storage, which is distributed across the network. The question is: How can multiple small caches be utilized in a cooperative manner to perform as a giant cache?

Recently, Maddah-Ali and Niesen [1] introduced a novel caching technique, called *coded caching*, which provides a gain that scales with the total cache distributed over the network. Unlike classical caching, coded caching is beneficial even if packets requested by one user are stored in the cache of other users. This surprising caching benefit is due to the opportunity for *multicasting*, i.e., when multiple users can be served by a single transmission. With coded caching, each transmission is a combination of multiple requested packets, where each user can retrieve its desired packet by removing the interference using the data stored in its cache.

The new coded caching scheme was first proposed for a fixed and homogeneous single-hop network. Following this pioneering work [1], the benefits of coded caching have been studied for various scenarios and applications, including device-to-device (D2D) communication [2], dynamic settings where the placement is not necessarily controlled by the server (decentralized caching) [3], [4], multihop networks [5], cloud networks [6], and multiple-input, multiple-output (MIMO) settings [7]–[11]. Currently, the main attention of the community is on the practical challenges of coded caching under nonideal channel models, e.g., packet erasure or fading [10], [12]–[16].

Despite its name, the main feature of coded caching is not coding in the information-theoretic sense (e.g., error correction codes) but rather combining multiple packets, each requested by one receiver so that each target receiver can cancel the interfering packets using its previously stored data and retrieve its desired packet. This allows for simultaneously serving multiple users, and hence leads to an increase in the number of degrees of freedom (DoF) in the system. As a result, the cache-aided communication problem can be formulated in the signal processing language and solved using the tools and techniques of signal processing. In this article, we present the novel aspects of cache-aided communication while focusing on the signal processing problems that lie at the heart of these schemes. In particular, we

present a simple and tractable problem formulation while highlighting the relations to common signal processing problems. We show that the challenges of cache-aided communication can be largely decomposed into three main problems: resource allocation, the design and coordination of beamforming/precoding vectors, and interference mitigation. These problems have been well studied in the signal processing community in different contexts, and (close to) optimum solutions are advised. However, these solutions must be adapted to the emerging field of cache-aided communication before this technique can be utilized for practical networks.

As practical implementations of cache-aided communication systems are approaching, many practical concerns and complexity issues still need to be addressed. Thus, the signal processing community can make a significant contribution to the derivation and optimization of novel cache-aided communication schemes. Furthermore, the signal processing community's experience with solving the aforementioned problems and, in particular, with handling the imperfections and practical constraints in the network can be crucial

as cache-aided communication continues to be adopted for practical implementations.

### Relevant works

The topic of coded caching has received a considerable amount of attention in recent years and is highlighted in several survey papers. In particular, [17] addressed the scaling laws of throughput in wireless networks with caching, with a special focus on the D2D approach. The challenges of edge caching in wireless networks are studied in [18], where the differences between wired and wireless caching are outlined. Specifically, this article discusses the essential limitations of wireless caching and the possible tradeoffs between spectral efficiency, energy efficiency, and cache size. Context-aware networks using edge/cloud computing and the exploitation of big data analytics are investigated in [19]. A tutorial on coded caching is presented in [20], which provides an introduction for the seminal and pioneering papers that opened new avenues in caching as well as a brief overview of existing caching solutions from an information-theoretic perspective. Moreover, [20] surveys some of the industrial challenges of caching and identifies bottleneck issues that need to be resolved to unleash the full potential of caching in practical systems.

The main distinction of this article from other surveys is its focus on the signal processing aspects of caching. In particular, we focus on interference mitigation using caching in wireless networks. This article also emphasizes the role of spatial reuse and the open issues that must be addressed prior to a practical adoption of coded caching.

### Notation

Throughout this article, we use calligraphic symbols (e.g.,  $\mathcal{W}$ ) to denote sets. The size of a set  $\mathcal{A}$  is denoted by  $|\mathcal{A}|$ . For an integer  $N$ , we denote the set  $\{1, 2, \dots, N\}$  by  $[N]$ .

**With coded caching, each transmission is a combination of multiple requested packets, where each user can retrieve its desired packet.**

## Classical coded caching

The application of caching was conventionally limited to pre-storing popular contents in storage units close to the end users and delivering the data from the local copies to reduce the server load and bandwidth requirements [21], [22]. The main challenge there is to predict users' requests before they actually ask for the data [21], [23], [24]. More importantly, the advan-

tage provided by conventional caching schemes is limited to its capacity of local memory/cache and can be negligible if the cache size is much smaller than the total size of popular data.

In wireless networks, the available memory on individual devices is small compared to the database size; however, with the growing popularity of mobile devices, plenty of such small-size memories are available and distributed over all the

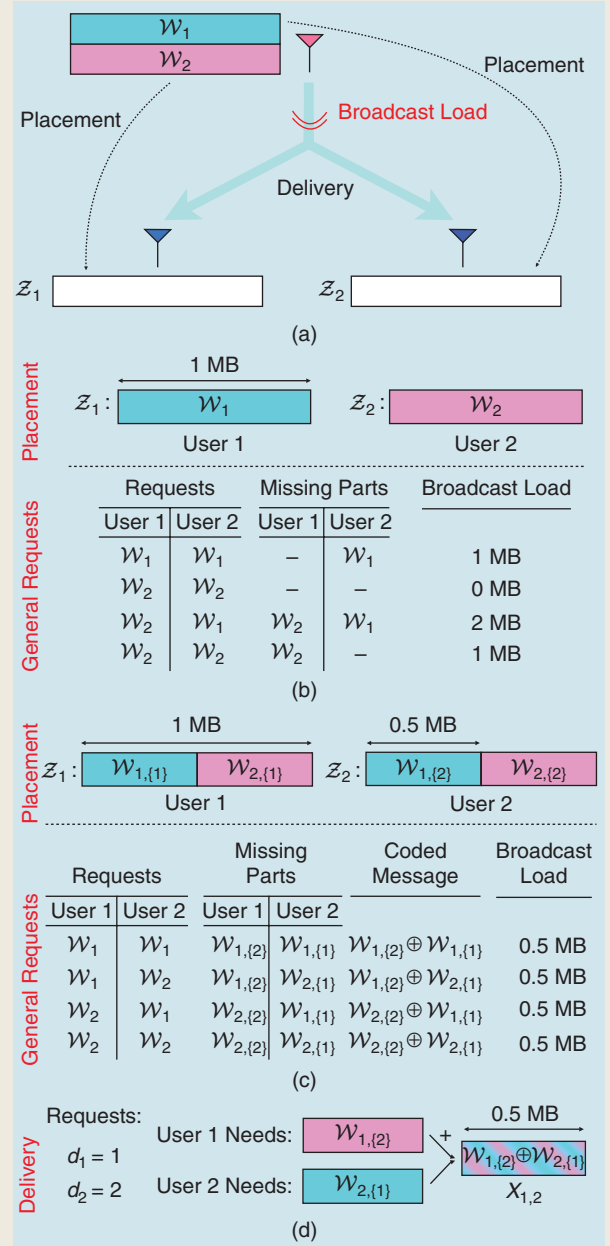
### Traditional Versus Coded Caching for Two Users

Consider the network in Figure S1(a) with  $N=2$  files, namely  $\mathcal{W}_1$  and  $\mathcal{W}_2$ , each of size 1 megabyte, and  $K=2$  users, each equipped with a local cache size of 1 megabyte. Assume a perfect broadcast link from the server to the receivers. Figure S1(b) displays a naive placement strategy, where each user caches one of the files. Once the user requests are revealed, the server must transmit all of the files that were requested and not stored at the requesting user. The tables show the files to be broadcast and the load of delivery under different demands. The load of delivery is not fixed and its average size over different demands is 1 megabyte.

Figure S1(c) depicts a smarter placement strategy wherein each file is divided into two halves, each of size 0.5 megabytes, i.e.,  $\mathcal{W}_1 = \{\mathcal{W}_{1,\{1\}}, \mathcal{W}_{1,\{2\}}\}$  and  $\mathcal{W}_2 = \{\mathcal{W}_{2,\{1\}}, \mathcal{W}_{2,\{2\}}\}$ . Each user caches one half of each file, i.e., the cache content of user 1 is  $\mathcal{Z}_1 = \{\mathcal{W}_{1,\{1\}}, \mathcal{W}_{2,\{1\}}\}$  and that of user 2 is  $\mathcal{Z}_2 = \{\mathcal{W}_{1,\{2\}}, \mathcal{W}_{2,\{2\}}\}$ . Again, the sub-files to be broadcast for each demand are listed in the table. For any demand, the server must send two subfiles, each of size 0.5 megabytes.

However, in all cases, the load of broadcast can be reduced to only 0.5 megabytes by sending a coded packet. Thus, coded caching can send twice as much information over the same period of time.

For instance, consider demands  $d_1 = 1$  and  $d_2 = 2$  shown in Figure S1(d), i.e., assume user 1 needs  $\mathcal{W}_1$  and user 2 needs  $\mathcal{W}_2$ . Because user 1 has already cached the first half of  $\mathcal{W}_1$ , it only needs the second half of  $\mathcal{W}_1$ , which is  $\mathcal{W}_{1,\{2\}}$ . Similarly, user 2 needs  $\mathcal{W}_{2,\{1\}}$  to be able to recover its desired file,  $\mathcal{W}_2$ . Instead of sending these two packets separately (uncoded caching), the server can combine them and send  $X = \mathcal{W}_{1,\{2\}} \oplus \mathcal{W}_{2,\{1\}}$ . User 1 can remove the interfering part,  $\mathcal{W}_{2,\{1\}}$ , from  $X$  using its cached data,  $\mathcal{Z}_1$ , and recover  $\mathcal{W}_{1,\{2\}}$ . Similarly, user 2 can decode its missing part,  $\mathcal{W}_{2,\{1\}}$ . Clearly,  $X$  is the summation (e.g., binary exclusive-OR) of two (binary) packets of size 0.5 megabytes, and hence we are only broadcasting 0.5 megabytes in a coded scheme, whereas separately broadcasting  $\mathcal{W}_{1,\{2\}}$  and  $\mathcal{W}_{2,\{1\}}$  requires 1 megabyte of data transmission. Similar arguments hold for all of the other demands, as presented in the figure.



**FIGURE S1.** Traditional versus coded caching. (a) System model, (b) placement of a subset of files, (c) placement of parts of files, and (d) coded delivery.

networks. A key question is: How can a file stored at one user's cache be utilized to reduce the network traffic when it is only requested by other users?

Recently, Maddah-Ali and Niesen [1] demonstrated that distributed storage units can be utilized in a cooperative manner to achieve a global gain that is proportional to the total available cache. Coded caching is a central placement mechanism that stores packets of the popular files in users' memory all over the network. This allows for an opportunistic multicasting in the delivery phase of the coded caching, in which a single (coded) packet can simultaneously serve multiple requests. The multicasting of packets to several users increases the utility of packets, and reduces the network load and congestion probability during peak traffic time (at the delivery phase). A simple coded caching scheme for a two-user scenario is depicted in "Traditional Versus Coded Caching for Two Users."

More generally, the single shared-link network studied in [1] consists of a network with  $K$  users and a server [base station (BS)] with a library of  $N$  files  $\{\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_N\}$ , each consisting of  $F$  bits, i.e.,  $|\mathcal{W}_n| = F$  for  $n \in [N]$ . For the sake of simplicity, we assume that the number of users does not exceed the size of the library, i.e.,  $N \geq K$ . Each user is equipped with a storage memory to store up to  $MF$  bits in the placement phase, which happens during the off-peak time of the network. Then, in the delivery phase, each user requests one file and the server is required to serve all the users by broadcasting a message through a perfect and shared channel.

### Placement phase

In the placement phase, prior to receiving users' requests and during the off-peak time of the network, the server selects subsets of all the bits in the library and stores them in the memory of user  $k \in [K]$ , which is denoted by  $\mathcal{Z}_k$  with  $|\mathcal{Z}_k| \leq MF$ . The placement strategy of coded caching proposed by Maddah-Ali and Niesen (referred to as the MAN scheme) is symmetric across files and users and allocates a  $1/N$  fraction of each user's cache to each file. Thus, each user will store a  $\mu = M/N$  fraction of each file in its cache.

Focusing on the case that  $\alpha = KM/N$  is an integer, the MAN scheme stores each packet in the cache of  $\alpha$  users. Thus, each file is split into  $\binom{K}{\alpha}$  equal segments, each of size  $F/\binom{K}{\alpha}$  bits. It is more convenient to index the segments by subsets of  $[K]$  of size  $\alpha$ . File  $\mathcal{W}_i$  will therefore be partitioned into  $\{\mathcal{W}_{i,S} : S \subseteq [K], |S| = \alpha\}$ , and each segment  $\mathcal{W}_{i,S}$  will be stored in the cache of every user  $k$  satisfying  $k \in S$  (see Figure 1). The content of the cache at user  $k$  will be

$$\mathcal{Z}_k^{\text{MAN}} = \{\mathcal{W}_{i,S} : S \subseteq [K], |S| = \alpha, \alpha \ni k\}. \quad (1)$$

Using this method, each user will cache  $\binom{K-1}{\alpha-1}$  segments out of a total  $\binom{K}{\alpha}$  segments, for each file. It is easy to verify that the number of bits from each file stored in each user's cache is  $KF/\alpha = \mu F$ .

### Delivery phase

Upon receiving the requests  $\{d_k : k \in [K]\}$  from the users (i.e., user  $k$  requests file  $\mathcal{W}_{d_k}$ ), the server broadcasts a message,  $X$ , which is a sequence of coded packets, to serve all user demands. The formation of this combination depends on the actual demand profile  $(d_1, \dots, d_K)$  and thus can be denoted by  $X_{(d_1, \dots, d_K)}$ . Upon receiving  $X$ , user  $k$  should be able to reconstruct its desired file  $\mathcal{W}_{d_k}$  using its cache content  $\mathcal{Z}_k$  and  $X$ , i.e.,

$$(X_{(d_1, \dots, d_K)}, \mathcal{Z}_k) \mapsto \mathcal{W}_{d_k}, \quad k \in [K].$$

In the MAN scheme, at the delivery phase, the server provides user  $k$  with missing (not cached) segments of its requested file  $\mathcal{W}_{d_k}$ , i.e., all  $\mathcal{W}_{d_k,S}$  where  $k \notin S$ . However, such a missing segment is cached in the local memory of exactly  $\alpha$  other users, indexed by elements of  $S$ . The symmetric placement guarantees a similar situation with respect to every other user in  $\mathcal{A} = S \cup \{k\}$ . Thus, at any given time, the server can simultaneously serve a subset  $\mathcal{A}$  of  $\alpha + 1$  users by multicasting a linear combination of all such packets. That is, for any subset  $\mathcal{A} \subseteq [K]$  with  $|\mathcal{A}| = \alpha + 1$ , the server sends

$$X_{\mathcal{A}} = \bigoplus_{j \in \mathcal{A}} \mathcal{W}_{d_j, \mathcal{A} \setminus \{j\}}, \quad \forall \mathcal{A} \subseteq [K], |\mathcal{A}| = \alpha + 1. \quad (2)$$




FIGURE 1. A file partitioning used for placement.



Note that we consider the file segments as bit streams of the same length, and  $\oplus$  denotes binary-exclusive-OR (XOR). It is worth mentioning that later in this article, we show that a similar operation can be performed using finite field summation on modulated sequences, which will be denoted by  $+$  instead of  $\oplus$ . Each user  $k \in \mathcal{A}$  has every term in this linear combination cached in its memory, except  $\mathcal{W}_{d_k, \mathcal{A} \setminus \{k\}}$ . So, the user can suppress the interference to recover its desired packet. Such a combined packet has utility  $(\alpha + 1)$ , because it can simultaneously serve  $(\alpha + 1)$  users. An example of the MAN scheme with  $K = 3$  users is given in “Three-User Coded Caching.”

#### Performance evaluation

The *load of delivery* is defined as the normalized size of  $X$ , which is the sum of the (normalized) size of the broadcast combinations, i.e.,

$$D = \frac{|X|}{F} = \frac{1}{F} \sum_{\substack{\mathcal{A} \subseteq [K] \\ |\mathcal{A}| = \alpha + 1}} |X_{\mathcal{A}}| = \frac{1}{F} \binom{K}{\alpha + 1} \frac{F}{\binom{K}{\alpha}} = \frac{K(1 - \mu)}{1 + \alpha}. \quad (3)$$

This leads to a delivery time [1] of

$$T_{\text{centralized-caching}} = \frac{D \cdot F}{R} = \frac{K(1 - \frac{M}{N})}{1 + \frac{KM}{N}} \frac{F}{R}, \quad (4)$$

where  $R$  is the rate supported by the common and shared links from the server to the users. This represents a significant improvement compared to that of conventional and uncoded caching, which has a load of  $D_{\text{uncoded}} = K(1 - \mu)$  and requires a delivery time of

$$T_{\text{uncoded}} = K(1 - \frac{M}{N}) \frac{F}{R}. \quad (5)$$

Another interpretation of (4) can be expressed: Each user has cached  $\mu F$  bits of its desired file and requires another  $(1 - \mu)F$  bits. Thus, a total of  $K(1 - \mu)F = K(1 - M/N)F$  bits are requested by all the users, which can be sent at a common rate of  $R$ . The factor  $1 - M/N$  is the (typically small) local caching gain due to the parts of the requested files that were prestored in the cache of the requesting users. The more

### Three-User Coded Caching

Consider  $K = 3$  users in the system, where all of the users receive information from the server through a common perfect broadcast link that supports a rate of  $R$  bits/s. Each user is equipped with a cache of size  $MF$ , where  $M = 2N/3$ , i.e., each user can prefetch two-thirds of each file. We first partition each of the files into

$$\binom{K}{MK/N} = \binom{3}{2} = 3$$

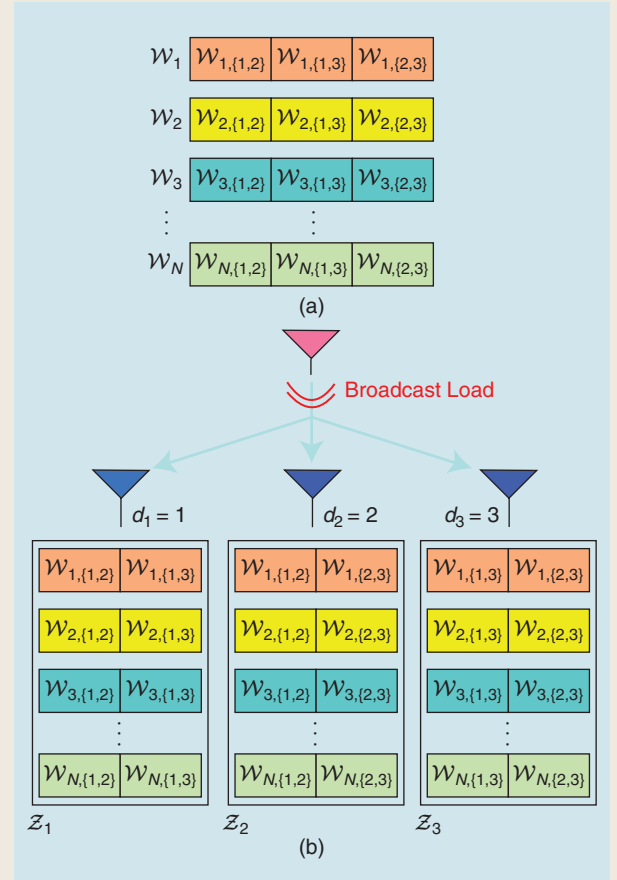
segments and label them with subsets of  $\{1, 2, 3\}$  of size  $MK/N = 2$ , as shown in Figure S2(a). Then, each user  $k$  prefetches all of the segments in  $\mathcal{W}_{i, \mathcal{S}}$  with  $k \in \mathcal{S}$ . The resulting placement is displayed in Figure S2(b). Assume that user  $k$  requested file  $\mathcal{W}_k$ , for  $k = 1, 2$ , and  $3$ . Each user has two segments of its desired file in the cache, and the remaining segment should be sent during the delivery phase.

More precisely, segments  $\mathcal{W}_{1, \{2, 3\}}$ ,  $\mathcal{W}_{2, \{1, 3\}}$ , and  $\mathcal{W}_{3, \{1, 2\}}$  must be delivered to the users. To this end, the base station (BS) transmits the combination

$$X = \mathcal{W}_{1, \{2, 3\}} \oplus \mathcal{W}_{2, \{1, 3\}} \oplus \mathcal{W}_{3, \{1, 2\}}. \quad (S1)$$

Upon receiving  $X$ , user 1 can remove  $\mathcal{W}_{2, \{1, 3\}}$  and  $\mathcal{W}_{3, \{1, 2\}}$  using its cache content and recover  $\mathcal{W}_{1, \{2, 3\}}$ . A similar argument holds for users 2 and 3.

Note that in each time slot, the BS is broadcasting a segment (combination) whose size is one-third of the size of a file. Given the common rate of  $R$  bits/s and the file length of  $F$  bits, the transmission takes only  $F/3R$  seconds. During this time, three files are delivered to the users, resulting in an overall network throughput of  $3F/(F/3R) = 9R$  bits/s.



**FIGURE S2.** A cache-aided communication with  $K = 3$  users. (a) File partitioning and (b) cache placement.

significant gain comes from the fact that each transmission can simultaneously serve  $1 + \alpha = 1 + KM/N$  users. This global caching gain is due to the joint placement and delivery scheme. In the context of wireless communication, this can be captured as a  $(1 + \alpha)$  prefactor for the rate, which is referred to as its *DoF*.

We also use network throughput to quantify the overall number of bits (including the prestored segments) delivered to the users per unit of time. Hence, the throughput of the basic MAN scheme will be  $KF/T = (1 + \alpha)/(1 - \mu)$ . Note that the network throughput grows unboundedly as  $\mu \rightarrow 1$ .

For non-integer values  $\alpha$  we can use *memory sharing*, i.e., the network can be treated as an interpolation between two separate systems, one with cache  $KM_1/N = \lfloor KM/N \rfloor$  and another with  $KM_2/N = \lceil KM/N \rceil$ . Hence, the delivery load and transmission time can be obtained as a linear combination of those of networks with two closest integer values.

With the description of the placement and delivery phases of MAN presented in this section, we provided the necessary background information to understand the signal processing aspects of cache-aided communication discussed in the “Signal Processing Problem Formulation” section. We refer interested readers to the “Coded Caching: A Broader Picture” section, where some of the most important follow-up works related to coded caching are presented.

## Signal processing problem formulation

### Linear combining versus coding

The main idea of cache-aided communication is the joint transmission of information to several users, where each user is able to decode its desired data using the data stored in its cache. The original approach used an XORing of the data required by the different users after making sure that all served users have their interfering data in their cache (as described previously). The so-called coded caching refers to the XOR operation (addition in the binary field), which is mostly used in the field of coding and not very typical as a signal processing technique. However, the practical implementation of such schemes in wireless communications will definitely be used in the field of signal processing and will typically not include XORing.

To introduce the signal processing aspects, we first note that a superposition of modulated signals instead of the XORing in (2) (see also “Traditional Versus Coded Caching for Two Users”) will do the work at a high signal-to-noise ratio (SNR) [25]. If the transmitter transmits  $X = \tilde{\mathcal{W}}_{1,\{2\}} + \tilde{\mathcal{W}}_{2,\{1\}}$  where  $\tilde{\mathcal{W}}_{i,\{k\}}$  is the modulated version of  $\mathcal{W}_{i,\{k\}}$  and  $+$  is a simple addition of the coded and modulated data, then each user can extract its desired information by simple subtraction (e.g., user 1 will estimate  $\hat{\mathcal{W}}_{1,\{2\}} = X - \tilde{\mathcal{W}}_{2,\{1\}}$ ). The only drawback of this approach is that transmitting  $\tilde{\mathcal{W}}_{1,\{2\}} + \tilde{\mathcal{W}}_{2,\{1\}}$  will take twice as much power than will transmitting  $(\mathcal{W}_{1,\{2\}} \oplus \mathcal{W}_{2,\{1\}})$ . This difference becomes negligible at a high SNR.

The problem becomes more interesting in multiantenna systems, where the information can be further differentiated in the spatial domain. In this case, the XORing alternative has a disadvantage, as it requires that all served users will be able to decode the joint message. Thus, the joint message must be sent in a “direction” and at a code rate that will allow for proper detection by all of the users. Conversely, the simple addition of the messages can be much more efficient because each user is only required to decode its own message (and interference subtraction does not require decoding).

Adding messages is therefore an important alternative. Furthermore, this allows for simple adaptations of many signal processing techniques. In the following section, we focus on this approach and give a straightforward description of a caching scheme that has no “coding,” i.e., the cache is placed in the users’ memory without coding and the BS transmits sums of modulated data for different users.

## System model and problem formulation

### Cache placement

The notations for the cache placement were introduced in the “Placement Phase” section. At the placement phase, the BS divides each file into segments indexed by the subset  $\mathcal{S}$ , and stores in the cache of user  $k$  all segments  $\mathcal{W}_{i,\mathcal{S}}$  for which  $k \in \mathcal{S}$ . Thus, the cache content of user  $k$  is  $\mathcal{Z}_k = \{\mathcal{W}_{n,\mathcal{S}} : n \in [N], \mathcal{S} \subset [K], k \in \mathcal{S}\}$ . We denote the normalized cache size of user  $k$  by  $M_k = |\mathcal{Z}_k|/F$ , where, for generality’s sake, we allow the cache sizes to be different. The overall performance is mostly dominated by the average cache size  $MF = (1/K)\sum_k M_k F$ .

### Transmission and channel model

At the beginning of the delivery phase, once the user requests,  $\{d_k\}$ , are known, each requested segment,  $\mathcal{W}_{d_k,\mathcal{S}}$ , is encoded at the rate supported by its requesting user. If multiple users request the same file, it will be encoded at the rate that will enable decoding by the weakest user (while other users may be able to decode the message after receiving only some of the coded symbols).

We denote with  $\tilde{\mathcal{W}}_{i,\mathcal{S}}$  the digitally modulated symbol set resulting from mapping the data packet  $\mathcal{W}_{i,\mathcal{S}}$  into a suitable signal space codebook. In this article, we do not discuss the details of how this can be done. Any suitable coded-modulation technique (e.g., low-density parity check code concatenated with a suitable high-order signal constellation) may be a practical implementation of the scheme.

Consider a multiple-input, single-output (MISO) communication scheme, where the BS has  $L$  antennas and each user has a single antenna. The  $n$ th symbol after match filtering and sampling at the  $k$ th user is given by

$$y_k[n] = \mathbf{h}_k \mathbf{x}[n] + z_k[n], \quad (6)$$

where  $\mathbf{x}[n]$  is the transmit vector,  $z_k[n] \sim \mathcal{CN}(0, N_0)$  is the additive noise sample,  $N_0$  is the power spectral density of the

complex white Gaussian noise at each receiver, and  $\mathbf{h}_k \in \mathbb{C}^{1 \times N}$  is the channel vector from the BS to user  $k$ . For this description, we focus on the case where the BS has perfect channel state information.

### System requirements

The system objective is to allow each user to decode each of its desired segments, while transmitting the least number of symbols. The ability to detect a message depends on the exact system definition. Using the general definition given so far, the ability to decode can be stated using information-theoretic terms: we need a sufficient amount of mutual information between the transmitted symbols and the received signal together with the cache content at the specific user.

However, this formulation does not reveal the structure of the problem and, in particular, does not tell us how to construct the transmitted symbol vectors. Thus, we turn to linear precoding to obtain more insight.

### Linear precoding

#### General linear precoding

For linear precoding, we assign symbols from several segments to be transmitted at every time interval. The transmission scheduling is described by a sequence of transmission assignments. Each transmission assignment is a set  $\mathcal{T}[n]$ , where  $(i, \mathcal{S}) \in \mathcal{T}[n]$  indicates that a symbol from  $\mathcal{W}_{i,\mathcal{S}}$  is transmitted at time  $n$ . Then, the transmitted vector is constructed by

$$\mathbf{x}[n] = \sum_{(i,\mathcal{S}) \in \mathcal{T}[n]} \sqrt{E_{i,\mathcal{S}}[n]} \cdot \mathbf{f}_{i,\mathcal{S}}[n] \tilde{u}_{i,\mathcal{S}}[n], \quad (7)$$

where  $E_{i,\mathcal{S}}[n]$  is the energy assigned for this transmission,  $\mathbf{f}_{i,\mathcal{S}}[n]$  is the precoding vector (normalized to  $\|\mathbf{f}_{i,\mathcal{S}}[n]\|^2 = 1$ ), and  $\tilde{u}_{i,\mathcal{S}}[n] \in \mathcal{W}_{i,\mathcal{S}}$  is the transmitted symbol (a different symbol from the segment at each assigned time).

Considering a specific symbol, described by  $(i, \mathcal{S}) \in \mathcal{T}[n]$ , user  $k$  will need to decode this symbol if  $d_k = i$  (i.e., if user  $k$  requested this file) and  $k \notin \mathcal{S}$  (i.e., the symbol is not stored at user  $k$ ). Attempting to decode the symbol, the desired signal gain is

$$A_{k,i,\mathcal{S}}[n] = \sqrt{E_{i,\mathcal{S}}[n]} \cdot \mathbf{h}_k \mathbf{f}_{i,\mathcal{S}}[n]. \quad (8)$$

The receiver can subtract all of the cached symbols, and the residual noise plus interference power is given by the conditional variance

$$\sigma_{k,i,\mathcal{S}}^2[n] = \text{Var}(y_k[n] - A_{k,i,\mathcal{S}}[n] \tilde{u}_{i,\mathcal{S}}[n] | \mathcal{Z}_k). \quad (9)$$

To ensure the proper decoding of all the requested files, we must verify that each segment is decodable at all requesting users. A common approximation for the decodability of a packet uses the Shannon capacity of an additive white Gaussian noise channel, with a multiplicative constant that represents that system imperfections. This constant,  $\gamma$ , is commonly referred to

as the *Shannon gap* and takes values between  $-2$  and  $-10$  dB. Thus, we say that segment  $\mathcal{W}_{d_k,\mathcal{S}}$  is decodable at user  $k$  if

$$\sum_{n:(d_k,\mathcal{S}) \in \mathcal{T}[n]} \log_2 \left( 1 + \gamma \frac{|A_{k,d_k,\mathcal{S}}[n]|^2}{\sigma_{k,d_k,\mathcal{S}}^2[n]} \right) > |\mathcal{W}_{d_k,\mathcal{S}}|. \quad (10)$$

Hence, cache-aided communication with linear precoding is solvable if there exist assignments,  $\mathcal{T}[n]$ , and matching precoding vectors, such that all users can decode all segments of their requested files. This is still a difficult assignment problem, and the optimal linear assignment is not yet known. The only approach that has been solved thus far is that of limiting the discussion to zero-forcing (ZF) precoding. We note that ZF is typically efficient at the high SNR regime.

### The ZF system

Using this approach, we make the following limiting assumptions as compared to the general linear case. These assumptions result in the most tractable problem formulation thus far:

- Each beamforming vector is selected in the unique direction that is orthogonal to the channel vectors of  $L - 1$  users.
- At any given time, exactly  $L + \alpha = L + KM/N$  users are active.
- After cache subtraction, all symbols are detected in the presence of noise only.
- The transmission rate to each user is set independently from the scheduling decisions.

Also, for simplicity's sake, we describe only the worst case scenario for the delivery phase, i.e., the case where all users request different files.

The first assumption allows for  $L - 1$  users to not be interfered with during the transmission to another user. To comply also with the second and third assumptions, we need to remove the interference at additional  $\alpha$  users by cache subtraction. Hence, every symbol must be stored at the cache of exactly  $\alpha$  users. Compared to (9), the third assumption guarantees that for all the symbols,  $\sigma_{k,i,\mathcal{S}}^2[n] = N_0$ .

The last assumption is required for simplifying the transmission scheduling. As stated previously, for each symbol, the choice of precoding vectors is completely determined by the choice of  $L - 1$  users that receive messages at the same time and cannot subtract the given message using their cache. Thus, the effective gain for each user (and its achievable rate) will be different depending on the combination of interfered users. This may significantly complicate the scheduling [25].

To resolve this, most works turned to either the high SNR or ergodic fading regime. With the high SNR regime, we consider the performance as the transmission power grows to infinity. In such an asymptotic scenario, the resulting rate is proportional to the logarithm of the transmission power. Hence, in this regime it is assumed that the actual channel gains are negligible and all rates are approximately equal.

Another approach, which does not lead to equal rates for all users, assumes that each transmission has sufficient-enough



diversity so that its rate will converge to its expected rate. For example, the ergodic fading regime can be approached in a wideband system, where the transmission bandwidth is large and spread over many frequency bins. When the bandwidth,  $B$ , is large enough compared to the channel coherence time, using the law of large numbers, the spectral efficiency converges to its expectation:

$$\frac{R_k}{B} \rightarrow \mathbb{E} \left[ \log_2 \left( 1 + \gamma \frac{E_{d_k, S}[n]}{N_0} |\mathbf{h}_k \mathbf{f}_{d_k, S}[n]|^2 \right) \right]. \quad (11)$$

This limit results in a rate per user, that is independent of the other users in the network. This independent rate can still be different for each user. Yet, it allows a tractable problem formulation.

Once the rate per user is determined and does not depend on the transmission scheduling, checking whether a segment is decodable can be done by simply counting the symbols received for this segment by the intended user. In this section, we give a mathematical statement of the conditions that guarantee decodability. Prior to that, we give a mathematical formulation of the aforementioned assumptions, i.e., we define what transmission assignments,  $\mathcal{T}[n]$ , are valid.

Let  $\mathcal{K}(\mathcal{T}[n]) = \{k : (d_k, \mathcal{S}) \in \mathcal{T}[n]\}$  be the set of active users for the specific assignment  $\mathcal{T}[n]$ . A transmission assignment,  $\mathcal{T}[n]$ , is valid if  $|\mathcal{K}(\mathcal{T}[n])| = L + \alpha$  and for each  $(d_k, \mathcal{S}) \in \mathcal{T}[n]$ :

- Each bit is stored at  $\alpha$  users, i.e.,  $|\mathcal{S}| = \alpha$ .
- All users that store the relevant segment are also actively receiving data, i.e.,  $\mathcal{S} \subset \mathcal{K}(\mathcal{T}[n])$ .
- The BS does not send a segment to a user that already stores it in its cache, i.e.,  $k \notin \mathcal{S}$ .

Note that a transmission assignment,  $\mathcal{T}[n]$ , also uniquely defines the precoding vectors. Each precoding vector,  $\mathbf{f}_{i, S}[n]$  is chosen in the unique direction that is orthogonal to the  $L - 1$  users that are active and not in the set  $\mathcal{S}$ , i.e., if  $i = d_k$ , we have  $\mathbf{h}_\ell \mathbf{f}_{i, S}[n] = 0$  for any  $\ell \in \mathcal{K}(\mathcal{T}[n]) \setminus \{k \cup \mathcal{S}\}$ .

To design and analyze the network, we do not need to choose the actual assignment of symbols for each transmission. It suffices to characterize the number of symbols sent using each assignment type. Denote by  $n_{\mathcal{T}}$  the number of symbols dedicated to a specific transmission assignment type, i.e.,  $n_{\mathcal{T}} = |\{n : \mathcal{T}[n] = \mathcal{T}\}|$ . The problem can be formulated as an optimization problem on the variables  $\{n_{\mathcal{T}}\}$  for all valid  $\mathcal{T}$ . The optimization goal is to minimize the transmission time, i.e.,  $\sum_{\mathcal{T}} n_{\mathcal{T}}$ , subject to the decodability of all the requested segments. This constraint can now be simply stated as: For each user,  $k$ , and each set  $\mathcal{S}$  with  $|\mathcal{S}| = \alpha$  and  $k \notin \mathcal{S}$ , we require that the number of received symbols will be the number of coded symbols, i.e.,

$$\sum_{\mathcal{T} : (d_k, \mathcal{S}) \in \mathcal{T}} n_{\mathcal{T}} = |\mathcal{W}_{d_k, \mathcal{S}}|. \quad (12)$$

An example of the optimal cache placement and transmission scheduling in the homogeneous case is given in “A Multiple-Input, Single-Output Cache-Aided Communication System,” while an example of nonhomogeneous rates is given in “A Multiple-Input, Single-Output Cache-Aided Communication Systems With Heterogeneous Rates.”

Posing the question as an optimization problem allows us to find the best transmission scheduling for each network. The resulting optimization problem is linear, and we may utilize a variety of algorithms for efficient solution. Note that the formulation herein focuses only on the optimization of the worst-case scenario, in which all users request different files. A more detailed formulation can lead to further improvement in cases where several users request the same file.

Keep in mind that the optimization approach rarely leads to closed-form performance expressions. For example, for the case of arbitrary user rates, the only case with such a closed-form expression is the case that  $K = L + \alpha$ , with integer  $\alpha$ . This is a limited case, as it typically requires unreasonably large users' cache. Nevertheless, it is interesting because it shows that for many user rate combinations, all  $L + \alpha$  users can be simultaneously served using  $L$  antennas, where each user receives information at its own rate. It was shown [10] that a minimal time of

$$T = \frac{KF(1 - M/N)}{\sum_{k=1}^K R_k}, \quad (13)$$

is achievable if the user rates satisfy  $R_k \leq (1/L) \sum_{\ell=1}^K R_\ell$  for every  $k$  in  $\{1, 2, \dots, K\}$ .

For the completeness of the network description, we next show that the presented scheme is indeed feasible and achieves a DoF of  $L + KM/N$ . In the next section, we present a transmission scheduling scheme for the simple homogeneous case.

### Performance in the homogeneous case

In the homogeneous case, all files are of equal size and popularity, and all users have the same rate. In this case, we show that all users can receive their desired files while  $L + \alpha$  users are served at any given time with no interference. We note that this also represents the high SNR regime, and hence proves that the system achieves the expected DoF as claimed above. Again, we focus on the worst-case scenario, where all the users request different files.

In the homogeneous case, it is reasonable to adopt the cache allocation of the centralized MAN scheme described in (1). This cache allocation divides each file into  $\binom{K}{\alpha}$  equal parts denoted by  $\mathcal{W}_{k, S}$ , where  $|\mathcal{S}| = \alpha$  and the set index,  $\mathcal{S}$ , indicates the set of users that store  $\mathcal{W}_{k, S}$  in their cache, i.e.,  $\mathcal{W}_{k, S} \in \mathcal{Z}_i$  if and only if  $i \in \mathcal{S}$ . This cache allocation is useful in many cases and, in particular, when the network is symmetric.

The delivery phase in the homogeneous case is solved by simply setting the number of symbols in each valid transmission assignment type to be identical. Inspecting the conditions for a valid transmission assignment in the previous section, we note that there are  $\binom{K}{L + \alpha}$  possible choices for the set of  $L + \alpha$  active users in the assignment. Given the active users, the transmission to each user,  $k$ , is characterized by a pair  $(d_k, \mathcal{S})$ , where  $\mathcal{S}$  is selected from the other active users; therefore, there are  $\binom{L + \alpha - 1}{\alpha}^{L + \alpha}$  possible assignments for each set of active users. Thus, in total, we have  $Q = \binom{K}{L + \alpha} \cdot \binom{L + \alpha - 1}{\alpha}^{L + \alpha}$

possible transmission assignment types, and we assign an identical number of symbols to each: i.e.,  $n\tau = c$ .

Each transmission assignment includes  $L + \alpha$  segments, and each specific segment  $\mathcal{W}_{d_k,s}$  appears in  $(L + \alpha)/$

$(K \cdot \binom{K-1}{\alpha})$  of the transmission assignments, i.e., we have a total of  $J = Q \cdot (L + \alpha) / (K \cdot \binom{K-1}{\alpha})$  assignments. To decode each segment, we need to receive all of its symbols, i.e., we need  $c = |\tilde{\mathcal{W}}_{d_k,s}|/J$ . Knowing  $c$ , and noting that the

## A Multiple-Input, Single-Output Cache-Aided Communication System

Consider a cache-aided network with  $K = 3$  users and  $L = 2$  antennas at the transmitter. The channel from the base station (BS) to user  $k$  is denoted by  $\mathbf{h}_k$ , and we assume that all the users can decode at a rate of  $R$  bits/s. Each user has a memory of size of  $M/N = 1/3$ , and thus, can prefetch only one-third of each file. We first partition each of the files into  $\binom{K}{MK/N} = \binom{3}{1} = 3$  segments, and label them as  $\mathcal{W}_{i,\{1\}}$ ,  $\mathcal{W}_{i,\{2\}}$ , and  $\mathcal{W}_{i,\{3\}}$ , respectively, for every  $i \in [N]$ , as shown in Figure S3(a). Then, each user,  $k$ , prefetches segments  $\mathcal{W}_{i,\{k\}}$  for all  $i \in [N]$ . The resulting

placement is presented in Figure S3(b). Assuming for simplicity that user  $k$  requests the file with an index of  $d_k = k$ , i.e.,  $d_1 = 1$ ,  $d_2 = 2$ , and  $d_3 = 3$ . Note that, each user has one segment of its desired file in the cache, and the two remaining segments should be sent during the delivery phase.

The transmitter broadcast message can be represented (intuitively) as

$$\begin{aligned} \mathbf{x}(1) &= \mathbf{h}_3^\perp \tilde{\mathcal{W}}_{1,\{2\}} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{2,\{3\}} + \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{3,\{1\}}, \\ \mathbf{x}(2) &= \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{1,\{3\}} + \mathbf{h}_3^\perp \tilde{\mathcal{W}}_{2,\{1\}} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{3,\{2\}}, \end{aligned} \quad (\text{S2})$$

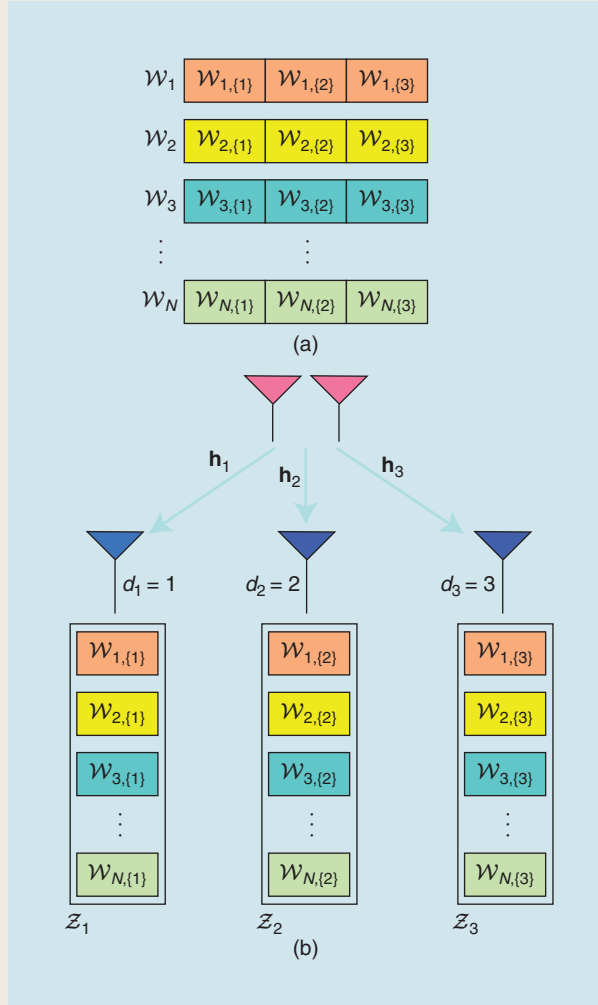
where  $\tilde{\mathcal{W}}_{i,s}$  denotes the digitally modulated sequence of symbols corresponding to the file segment  $\mathcal{W}_{i,s}$ . Also note that the precoder vector  $\mathbf{h}_k^\perp$  indicates that the beamformed signal is perpendicular to the channel of user  $k$  and will be zero forced at this user. The exact mathematical representation of the transmitted signals is given in (7).

Let us study the signal received by user 1. During the first time slot, user 1 receives

$$y_1(1) = \mathbf{h}_1 \mathbf{x}(1) + z_1(1) = \mathbf{h}_1 \mathbf{h}_3^\perp \tilde{\mathcal{W}}_{1,\{2\}} + \mathbf{h}_1 \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{3,\{1\}} + z_1(1),$$

where  $z_k(t)$  is the additive white Gaussian noise observed at user  $k$  in time slot  $t$ . Note that one interference term is zero forced, i.e.,  $\mathbf{h}_1 \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{2,\{3\}} = 0$ . Moreover,  $y_1(1)$  is a combination of the desired codeword  $\tilde{\mathcal{W}}_{1,\{2\}}$  and another interfering codeword  $\tilde{\mathcal{W}}_{3,\{1\}}$ . However, this interfering codeword can be reconstructed from the segment  $\mathcal{W}_{3,\{1\}}$  stored in  $\mathcal{Z}_1$ . Once the interference is subtracted, user 1 can decode segment  $\mathcal{W}_{1,\{2\}}$ . A similar argument holds for all the users and all the time slots, where each user can decode one desired segment at each time slot.

Note that each transmission can simultaneously serve  $L + KM/N = 2 + 1 = 3$  users. This is due to the possibility of interference cancellation using the cache content at  $KM/N = 1$  user and using zero forcing at  $L - 1 = 1$  user. In each time slot the BS is broadcasting a segment (combination) of length one-third of the size of a file, at rate  $R$  (which is decodable for all users). Thus, the duration of each transmission will only be  $F/3R$  seconds, where  $F$  is the file size. This leads to a total transmission time of  $2F/3R$  seconds. During this time, three files are delivered to the users, and therefore, we have an overall network throughput of  $3F/(2F/3R) = 4.5R$  bits/s.



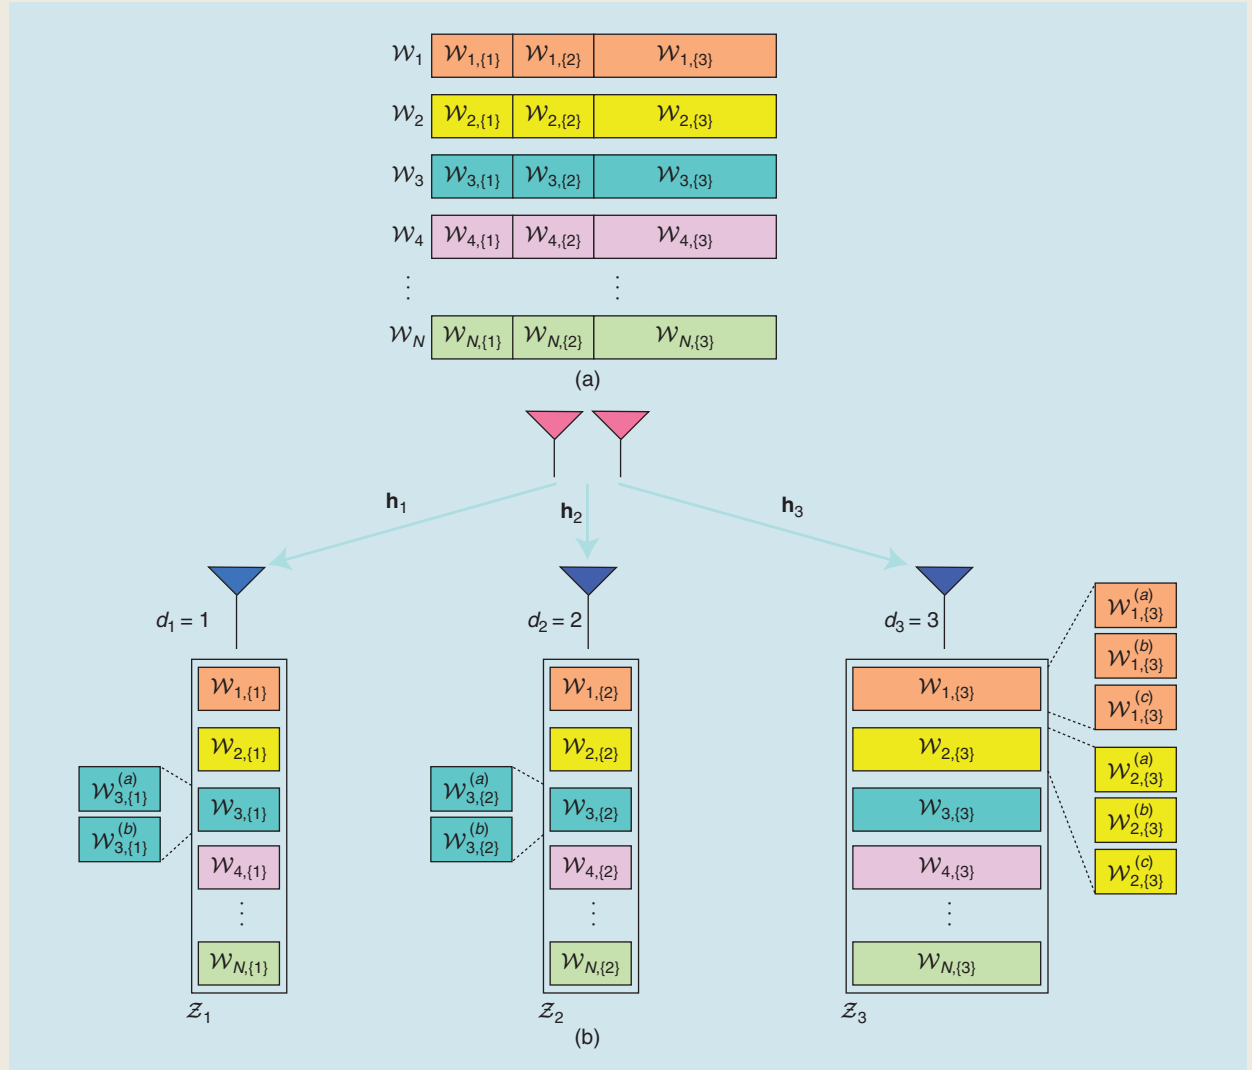
**FIGURE S3.** A multiple-input, single-output cache-aided communication system with  $L = 2$  transmit antennas,  $K = 3$  users, and normalized cache size of  $M/N = 1/3$ . The (a) file partitioning and (b) cache placement.

## A Multiple-Input, Single-Output Cache-Aided Communication Systems With Heterogeneous Rates

Consider a multiple-input, single-output (MISO) broadcast channel with  $L = 2$  transmit antennas and  $K = 3$  users with single antenna, as shown in Figure S3. We assume a total cache constraint so that only one copy of each (packet of each) file can be prefetched among all of the users, i.e.,  $\sum_{k=1}^3 |\mathcal{Z}_k| = NF$ . We denote the link capacity of user  $k$  by  $R_k$  and assume a heterogeneous topology, in which users 1 and 2 have good channels to support  $R_1 = R_2 = 2$  bits/s, while the third user is further away from the transmitter and can decode only at rate  $R_3 = 1$  bits/s. We consider a placement strategy that is symmetric across files but not necessarily symmetric across users. It is natural to expect that a larger cache will be allocated to the weak user to minimize the overall transmission time of the delivery phase.

It turns out that the optimal cache allocation places  $M_1 F/N = M_2 F/N = F/5$  bits in the cache of users 1 and 2, while user 3 stores  $M_3 F/N = 3F/5$  bits to compensate for its weakness. The cache allocation is accomplished by partitioning each file into three segments, i.e.,  $\mathcal{W}_{i,\{1\}}$ ,  $\mathcal{W}_{i,\{2\}}$ , and  $\mathcal{W}_{i,\{3\}}$ , which are stored at the cache of users 1, 2, and 3, respectively, as depicted in Figure S4(b). Note that the cache placement is performed prior to the user's request and is identical for all files. In this example, we assume that user  $k$  requested  $d_k = k$  for  $k \in \{1, 2, 3\}$ .

The delivery phase includes the broadcasting of four messages. To present the broadcast messages, we must further divide the cached messages into smaller segments as  $\mathcal{W}_{3,\{1\}} = (\mathcal{W}_{3,\{1\}}^{(a)}, \mathcal{W}_{3,\{1\}}^{(b)})$ ,  $\mathcal{W}_{3,\{2\}} = (\mathcal{W}_{3,\{2\}}^{(a)}, \mathcal{W}_{3,\{2\}}^{(b)})$ ,  $\mathcal{W}_{1,\{3\}} = (\mathcal{W}_{1,\{3\}}^{(a)}, \mathcal{W}_{1,\{3\}}^{(b)}, \mathcal{W}_{1,\{3\}}^{(c)})$ , and  $\mathcal{W}_{2,\{3\}} = (\mathcal{W}_{2,\{3\}}^{(a)}, \mathcal{W}_{2,\{3\}}^{(b)}, \mathcal{W}_{2,\{3\}}^{(c)})$ .



**FIGURE S4.** A cache-aided communication with  $L = 2$  transmit antennas. The (a) file partitioning and (b) cache placement.



## A Multiple-Input, Single-Output Cache-Aided Communication Systems With Heterogeneous Rates (Continued)

$\mathcal{W}_{2,\{3\}}^{(b)}, \mathcal{W}_{2,\{3\}}^{(c)}$ ) to keep up with the capacity of the links to the users. All three users can be served in four time slots by transmitting

$$\begin{aligned}\mathbf{x}(1) &= \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{1,\{3\}}^{(a)} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{2,\{3\}}^{(a)} + \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{3,\{1\}}^{(a)} \\ \mathbf{x}(2) &= \mathbf{h}_3^\perp \tilde{\mathcal{W}}_{1,\{2\}}^{(a)} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{2,\{3\}}^{(b)} + \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{3,\{1\}}^{(b)} \\ \mathbf{x}(3) &= \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{1,\{3\}}^{(b)} + \mathbf{h}_3^\perp \tilde{\mathcal{W}}_{2,\{1\}}^{(a)} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{3,\{2\}}^{(a)} \\ \mathbf{x}(4) &= \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{1,\{3\}}^{(c)} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{2,\{3\}}^{(c)} + \mathbf{h}_1^\perp \tilde{\mathcal{W}}_{3,\{2\}}^{(b)}\end{aligned}\quad (S3)$$

where  $\mathbf{x}(\ell)$  is the collection of all transmission vectors for the  $\ell$ th time slot and  $\tilde{\mathcal{W}}_{i,S}$  denotes the digitally modulated sequence of symbols corresponding to file segment  $\mathcal{W}_{i,S}$ . Moreover, beamforming a codeword with precoder vector  $\mathbf{h}_k^\perp$  guarantees that the received signal corresponding to that codeword will be zero forced at user  $k$ . The exact mathematical representation of the transmitted signals is given in (7).

Let us consider the decoding at user 1. For instance, during time slot 1, user 1 receives

$$y_1(1) = \mathbf{h}_1 \mathbf{x}(1) + z_1(1) = \mathbf{h}_1 \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{1,\{3\}}^{(a)} + \mathbf{h}_1 \mathbf{h}_2^\perp \tilde{\mathcal{W}}_{3,\{1\}}^{(a)} + z_1(1).$$

The user subtracts codeword  $\tilde{\mathcal{W}}_{3,\{1\}}^{(a)}$  using its cache content and then uses the remaining signal to decode  $\mathcal{W}_{1,\{3\}}^{(a)}$ . Similarly, each user can decode all of the missing segments of its requested file.

Note that each transmission time takes  $F/10$  seconds, and the total transmission time is  $0.4F$  seconds, after which, all users can decode their requested files. Because three users are served during this time, the overall network throughput is  $3F/0.4F = 7.5$  bits/s. In contrast, with equal cache placement, the minimal time to serve user 3 is  $\frac{2}{3} \cdot \frac{F}{R_3}$  seconds, which leads to a throughput of only 4.5 bits/s. This shows that an optimized coded caching in MISO brings additional gain by balancing the load of the network.

number of coded symbols in a segment is  $(F/R)/\left(\frac{K}{\alpha}\right)$ , the total number of symbols required to serve all users is

$$T = c \cdot Q = \frac{F/R}{\left(\frac{K}{\alpha}\right)} \cdot \frac{K \cdot \left(\frac{K-1}{\alpha}\right)}{L+\alpha} = \frac{K(1-M/N)}{L+\alpha} \cdot \frac{F}{R}. \quad (14)$$

Here,  $KF$  is the total number of bits in the files requested by all the users. Compared to (5), we see that the number of DoF given by

$$\text{DoF} = L + \alpha = L + KM/N, \quad (15)$$

where we see the combination of  $KM/N$  DoF that are achieved through cache-aided communication and  $L$  DoF of spatial multiplexing using  $L$  antennas.

### Related problems

#### Demand prediction

In some limited systems, the number of files in the database is relatively small, hence the notation in the previous section seems reasonable. However, in most cases, and in particular if we consider cellular systems of 5G and beyond, the database size can be huge. As can be verified from the performance measures in (14) and (15), a huge database size,  $N$ , can completely diminish the gain of caching.

Accordingly, cache-aided communications in large systems should cache only a subset of the files in the database, preferably the most popular ones [26]. Yet, the popularity of files changes over time, and a proper system operation would require a continuous update of the most popular sets [27].

Even more challenging is the prediction of the popularity of the files. This problem has already attracted much attention, although, not necessarily in the context of cache-aided communications. Because the problem is outside the scope of this article, we only refer the interested reader to [19] and [28] as two starting points where this problem in the context of caching is discussed. These highlight the importance of popularity prediction and cache updates. They also emphasize the effectiveness of machine-learning techniques, as the actual distribution of file popularity is typically unknown.

It is worth pointing out that these types of works are more related to web caching and content distribution networks or, in the wireless framework, on scenarios known as *femtocaching*. In coded caching, there is typically a substantial separation between the time scale at which users' requests are served and the time scale at which the popularity of the content evolves. As an example, the collection may include the most popular 1,000 titles from the Netflix library, and every week some new movies are included and some old ones are eliminated (while a streaming session is on the order of tens of minutes). In this context, demand prediction is virtually orthogonal and complementary with coded caching. One can use a standard scheme (e.g., as done by Netflix) to track what users want at the large time scale and use coded caching for the delivery of files from the current library.

#### Resource allocation

Resource allocation is one of the most important tasks in the optimization of any communication system. In cache-aided communication systems, traditional resource allocation problems become more challenging and new problems

need to be addressed. In this section, we discuss the power allocation as an example of a traditional resource allocation problem that becomes more complicated, and then we discuss the cache allocation as an example of a new allocation problem.

Note that any resource allocation problem is strongly related to the transmission scheduling problem. The transmission scheduling is, by itself, a resource allocation problem that determines which segments will be allocated to each transmission time. So, allocating any additional resources can only be done in coordination with transmission scheduling.

#### *Power allocation*

The power allocation problem requires finding the optimal transmission energy for each symbol, which minimizes the time required to serve all the users. The main difficulty in the optimization of the power allocation is the need for a closed-form performance expression. Such an expression, which can be written as a function of the transmission powers, currently exists only for the limited case where the number of users is exactly  $K = L + KM/N$ . As a result, this is also the only cache-aided communication scenario where the optimal power allocation is known [10], [29].

This power allocation is shown to be a variation of the water-filling algorithm with a rate saturation. As the cache-aided communication cannot take advantage of users with a very high rate, the rate of such users is saturated and the extra power is used for other users. Optimal power allocation for other cache-aided communication scenarios still remains an open problem.

#### *Cache allocation*

Cache allocation includes 1) allocating the memory (cache size) for each user and 2) allocating the cache content. The allocation of cache content means the choice of sets  $\mathcal{W}_{i,S}$  for each  $i$  and  $S$ , usually under the constraint of user cache size of  $MF$  bits.

For example, for nonuniform file popularity, it is likely that popular files will be stored more often than others. Thus, we may have  $|\mathcal{W}_{i,S}| \neq |\mathcal{W}_{j,S}|$  when files  $i$  and  $j$  have different popularity. This has mostly been studied for a single-antenna BS and formulated as linear optimization problem [30]. It was shown that the uniform placement of the centralized MAN scheme [1] is indeed optimal when the file popularity is uniform. However, in the case of nonuniform popularity, the optimal placement performs better than the uniform.

Contrarily, in the case of uniform file popularity, while the placement is still symmetric across the files, it is not necessarily symmetric across users. That is, if the BS decides that user  $k$  stores bit number  $n$ , then it will store this bit from all files in the database. So, we have  $|\mathcal{W}_{i,S}| = |\mathcal{W}_{j,S}|$  for all  $i$  and  $j$ , but not necessarily  $|\mathcal{W}_{i,S_1}| = |\mathcal{W}_{i,S_2}|$ .

For such a case, we observe that the cache placement optimization is much simpler for  $|\mathcal{S}| = KM/N = 1$ . In this case, there is no meaning to the question “Which specific bit is stored at which user?” Hence, the only cache allocation decision that has

an effect on the performance is the users’ cache sizes. Conversely, for  $KM/N > 1$ , there will be overlap between users’ caches (i.e., a bit may be stored at more than one user). The choice of which users’ caches will overlap can have an effect on performance in a nonhomogenous network.

Another improvement can be gained in networks that allow for an optimization of the cache size per user depending on the channel qualities. Recall that cache allocation occurs during the placement phase, and in most scenarios, at this stage the system has no knowledge of the future states of the channels. The typical approach employed is to use an equal cache size for all users.

In some cases, the users’ delivery rates can be predicted. Two examples of this are 1) in fixed wireless networks, where the rates are rather fixed or 2) in low-orbit satellite communications, where the rates change faster but more predictably because satellites travel in known orbits. In such cases, the cache size at each user can be adjusted to partly compensate for the different rates and increase the network throughput (an illustrative example is discussed in “A Multiple-Input, Single-Output Cache-Aided Communication Systems With Heterogeneous Rates”). In the case of multiantennas, this problem was framed as a linear optimization problem, albeit with a number of variables that grows exponentially with the number of users [10].

#### *Caching for scalable coding*

The cache-aided communication problem is usually studied in the strict fairness setup. The typical problem formulation is of equal size files, and each user requests a single file. Thus, the objective is to supply to each user the same amount of data, regardless of the channel qualities. The strict fairness constraint makes the network very sensitive to the performance of the weakest user. This is in contrast to other network-optimization works, which typically focus on the total network throughput.

An interesting scenario that challenges this fairness assumption is scalable coding. With scalable coding, the same content (e.g., a movie) can be compressed at divergent rates to produce distinct types of quality for a variety of users. Reproducing lower-quality versions by using lower data rates is desired, e.g., for diverse equipment types (e.g., different screen sizes), dissimilar channel conditions, or a range of content pricing. Because of this, the same file may be requested in the network at several quality levels. Scalable coding allows for the encoding of the file at several layers. Users that decode only the first layer will reconstruct the content at the lowest quality (highest distortion), while users that decode multiple layers will be able to refine the reconstruction and realize higher quality.

Having several layers for each file makes cache allocation much more interesting [31]. In particular, if the desired quality of each user is known in advance, both the cache allocation and placement can be optimized for the specific desired quality of each user. Obviously, such an optimization results in a better utilization of the cache and higher network throughput.

### Effect of transmitter cache

In a network with multiple single-antenna BSs, cache can be used to facilitate cooperation between the BSs. This was studied in a network where all BSs are synchronized and make joint scheduling decisions but do not store the complete database. In such a case, each bit can be transmitted only from the BSs that store them in their cache. Thus, a ZF precoder that nulls the interference caused by the transmission of a symbol can only be used in these BSs.

Denoting the cache size at each BS with  $M_T F$  and assuming that  $LM_T/N$  is an integer (where  $L$  is the number of BSs), each bit is accessible to  $LM_T/N$  antennas. Each transmission can therefore only be zero forced at  $LM_T/N - 1$  users. Combined with the gain of the receivers cache, the number of users that can be served simultaneously (hence, the DoF of the system) [32] is

$$\text{DoF} = \frac{LM_T + KM}{N}. \quad (16)$$

Note that if each BS has access to the whole database, complete cooperation can be established between all BSs, and the performance will be characterized by the single BS case that is equipped with  $L$  antennas. In this case, we have  $M_T = N$ , and (16) simplifies to (15).

It is important to emphasize that (16) gives an achievability result, and the maximal performance is not yet known. For example, in the specific case of  $L = 3$  and  $M_T/N = 2/3$  and  $M = 0$ , it was shown [32] that when using interference alignment the DoF is 18/7, which is significantly higher than the DoF of 2 that is achievable according to (16). The best-known general upper bound on the DoF of linear precoding (also termed *one-shot coding*) [32] is two times the DoF given in (16).

### Multiserver wireline network

Surprisingly, a similar channel model can be found in wireline multiserver networks. Using the concept of network coding, in some scenarios it is beneficial for each server in the network to forward a linear combination of its incoming packets. This linear combination is performed over a large-enough finite field so that the network is invertible. The operation of the whole network can then be represented by a matrix.

A central network manager controls multiple servers, which need to jointly serve multiple users. These multiple servers act as antennas in the wireless model presented above. Accordingly, the output at each user is described as a multiplication of a transmitted vector and a channel matrix, just as in (6), with the only differences being the operations over a large finite field instead of over the field of complex numbers [8] and the absence of noise. The same transmission schemes are applicable and result in the same performance gain.

### The subpacketization problem and the role of spatial multiplexing

For the single shared-link network presented in [1], the major problem of the applicability of coded caching schemes to real-

world systems is represented by the very large subpacketization order, i.e., the number of subpackets (segments) that each file  $\mathcal{W}_i$  in the library must be partitioned into. In the MAN scheme with  $K$  users,  $N$  files, and cache memory per user  $MF$  bits, assuming that  $\alpha = KM/N$  is an integer, each file must be split into  $\binom{K}{\alpha}$  subpackets. Letting  $\mu = M/N \in [0, 1]$  denote the fractional memory level, i.e., the fraction of the whole library that each user can cache, the subpacketization order of MAN is given by  $\binom{K}{K\mu} \approx \exp(K\mathcal{H}(\mu))$ , where  $\mathcal{H}(\mu) = -\mu \log \mu - (1 - \mu) \log (1 - \mu)$  is the binary entropy function. Therefore, the subpacketization order is exponential in  $K$  for a given fractional memory level  $\mu$ .

Denoting the length of each file by  $F$  bits, it is clear that  $F$  must be at least as large as the subpacketization order, i.e., the file size required by MAN is also  $O(\exp(K))$ . (In practice, the file size should be even larger as the transmission always processes a packet of many bits at a time.) For example, a system with  $K = 500$  users where each user caches a fraction  $\mu = 0.01$  of the library, would require a minimum file size in bits  $F \approx 2.5 \cdot 10^{11}$ . This means that the required file size is at least 250 gigabits. Taking into account that the typical size of a movie encoded in standard definition is of the order of 1 gigabits, we see that the file size required by such a scheme is between 1 and 2 orders of magnitude larger, and therefore, completely impractical.

The subpacketization order problem is exacerbated by the fact that, in a practical media streaming system, a streaming session consists of a sequence of demands of video “chunks,” corresponding to a few seconds of video. To cope with asynchronous streaming sessions, each video chunk of each video file should be treated as a “file”  $\mathcal{W}_i$  in the formalism of MAN, presented in this section [1]. It follows that, in practical video steaming applications, the actual size of the video chunks,  $F$ , is of an order of a few megabytes, which is four or five orders of magnitude smaller than that required file size in the example above.

Furthermore, it was proved in [33] that any decentralized caching scheme based on symmetric random caching, i.e., where each user caches a fraction  $\mu$  of the bits of each file selected at random with uniform probability, independently of the other users, must have  $F$  that grows superexponentially in  $K$  to achieve a  $\text{DoF} = O(K)$ . For example, the load expression of the decentralized MAN scheme in [3] [see (22)] is valid only in the limit of  $F \rightarrow \infty$  and fixed  $K$ . In contrast, in [33] it is shown that if  $F$  grows less than superexponentially in  $K$ , then the maximum achievable DoF does not exceed 2.

### Cache replication

Several methods have been proposed to cope with the subpacketization order problem (e.g., see [34]–[39]). These methods can be (roughly) classified into optimization-based (e.g., [34]–[36]) and graph-theoretic/combinatorial methods (e.g., [37] and [38]). Unfortunately, they are typically quite complicated and not flexible in terms of system parameters.

A much simpler approach consists of cache replication. For the centralized case, we create a MAN packetization



for a nominal number of users  $G$ , such that the subpacketization order  $\approx \exp(G\mathcal{H}(\mu))$  is kept to a reasonable value. For example, for  $G = 100$  and  $\mu = 0.01$  we have  $F = 100$ . Then, we divide the user population in  $G$  groups of  $K/G$  users each (assume for simplicity's sake that  $K/G$  is an integer). Users in the same group  $g = 1, \dots, G$  cache exactly the same packets, i.e., the  $g$ th cache configuration is replicated across all  $K/G$  users in each group  $g$ . In the delivery phase, a delivery array of dimension  $G \times K/G$  is formed by arranging the requests of the users in the same group by rows. Hence, the system delivers the requests by serving sequentially the columns of the delivery array. Note that each column forms a MAN scheme with  $G$  users because, by construction, the users in each column belong to different groups. It follows that all requests can be delivered with a load equal to the load of a  $G$ -user MAN scheme, given by  $G(1 - \mu)/(1 + \mu G)$  times the number of columns  $K/G$ . This yields

$$D = \frac{K}{G} \cdot \frac{G(1 - \mu)}{1 + \mu G}. \quad (17)$$

By letting  $G$  be a function of  $K$ , this scheme achieves a very desirable tradeoff between subpacketization order and DoF.

A decentralized version of the cache replication approach is proposed and analyzed in [39] (see also [40]). In this case, users simply choose one of the  $g$  groups (and the corresponding cache content) randomly and independently. We refer to the number of users in group  $g$  as the  $g$ th *occupancy number*, denoted as  $\ell_g$ . The vector of occupancy numbers  $(\ell_1, \dots, \ell_G)$  is random and follows a multinomial distribution over all possible  $G$ -way integer partitions  $(\eta_1, \dots, \eta_G)$  of  $K$ . It follows that the delivery for this problem, for given occupancy numbers, is identical to the shared caches network studied in [41], for which an optimal delivery under symmetric caching (as enforced by the cache replication construction) consists of sorting the occupancy numbers  $\ell_{(1)} \geq \dots \geq \ell_{(G)}$  in a nonincreasing order, such that  $\ell_{(g)}$  denotes the size of the  $g$ th most populated group, and forms a delivery array of size  $G \times \ell_{(1)}$ . Such an array has empty elements because, in general, groups have fewer than  $\ell_{(1)} = \max\{\ell_g\}$  users. Then, each column of the array is served by an improved MAN scheme that avoids sending XORs when they are not useful for at least one user (in fact, each column is served by the improved delivery scheme of [42]). The resulting load for  $\{\ell_{(g)}\}$  is given by

$$D(\ell_1, \dots, \ell_G) = \sum_{g=1}^{G(1-\mu)} \ell_{(g)} \frac{\binom{G-g}{\mu G}}{\binom{G}{\mu G}}, \quad (18)$$

Here, for simplicity's sake, we assume that  $\mu G$  is an integer. If not, the usual memory-sharing argument holds and the convex envelope of the load/memory points for  $\mu G \in \{0, 1, 2, \dots, G\}$  is achievable.

It turns out that this load is information-theoretically optimal for any given configuration of the occupancy numbers [41], in the case of distinct demands (which requires  $N \geq K$ ) and uncoded placement.

As a rough upper bound for the average load, where the averaging is with respect to the occupancy numbers, we have (trivially)

$$\mathbb{E}[D(\ell_1, \dots, \ell_G)] \leq \mathbb{E}[\ell_{(1)}] \cdot \frac{G(1 - \mu)}{1 + \mu G}. \quad (19)$$

By comparing (17) and (19), we note that the difference is generally small. In fact,  $\mathbb{E}[\ell_{(1)}]$  is the expectation of the maximum of  $G$  multinomial variables. Because each occupancy number  $\ell_g$  is marginally binomially distributed with parameters  $(K, 1/G)$ , its expected value is  $\mathbb{E}[\ell_g] = K/G$ . Now, although the ordered statistics of a jointly multinomial random vector are generally difficult to characterize,  $\mathbb{E}[\ell_{(1)}]$  behaves as  $K/G$  up to the logarithmic factors in  $G$ . Therefore, the average load  $\mathbb{E}[D(\ell_1, \dots, \ell_G)]$  of the decentralized scheme has essentially the same behavior as that of the centralized scheme in (17). A more refined analysis is given in [40].

### Exploiting spatial multiplexing

In the previous section, we explained how cache replication provides a viable approach for achieving a competitive tradeoff between DoF and subpacketization order. However, the fact that the delivery array columns are served one by one in sequence prevents such schemes from reaching low subpacketization order with the same (ideal) DoF of the MAN scheme for the single shared-link network. Intuitively, if we can “parallelize” the different columns of the delivery array using spatial multiplexing, we should be able to achieve both low subpacketization order and optimal DoF at the same time.

Following this idea, we revisit the MISO broadcast channel with caches at the receivers, previously discussed in the “Signal Processing Problem Formulation” section, and present a different scheme achieving the low subpacketization proposed in [43] and the same optimal DoF of the ZF precoding scheme in the “Signal Processing Problem Formulation” section. Consider a MISO broadcast channel with  $K$  users, a BS with  $L$  antennas, a library of  $N$  files of size  $F$  where each user has cache memory  $MF$  bits, yielding fractional memory level  $\mu = M/N$ , such that  $\mu K$  is an integer. For the sake of simplicity, assume that  $L$  divides both  $K$  and  $\mu K$ . We partition the users into  $G = K/L$  groups of  $L$  users each and use the cache replication approach. The cache placement consists of cache replication as explained previously: create a MAN subpacketization with parameters  $(N, M, G, \text{ and } F)$ , and let all users in group  $g$  to cache the same content. Note that this subpacketization consists of  $\binom{G}{\mu G}$  subpackets. Because  $G = K/L$ , the subpacketization order of this scheme is  $\approx \exp(K/L\mathcal{H}(\mu)) = \sqrt[L]{\exp(K\mathcal{H}(\mu))}$ , the  $L$ th root of the subpacketization order of the “classical” scheme presented in the “System Model and Problem Formulation” section. Referring to a previous example, for  $K = 500$ ,


$\mu = 0.01$ , and  $L = 8$  antennas, we find  $F \approx 27$  in contrast to the case of  $L = 1$  for which  $F \approx 2.5 \cdot 10^{11}$ .

The delivery consists of simultaneously serving the  $L$  sections of the delivery array by combining a  $K/L$ -user MAN scheme with the  $L$ -fold spatial multiplexing obtained by ZF MIMO precoding. Figure 2 qualitatively shows the delivery array of dimension  $K/L \times L$ , where each column is formed by users belonging to distinct caching groups. Therefore, each column can be served using a  $K/L$ -user MAN scheme; however, unlike for the single shared-link network, here, the inherent spatial multiplexing of the MIMO channel can be used to serve all the  $L$  slices simultaneously.

Due to space limitations, we omit the details of the combined coded caching and MIMO precoding scheme, which can be found in [43]. The important point to notice here is that, although in the basic cache replication scheme for the single shared-link network we have to serve the sections of the delivery array in sequence [see (18)], in the MIMO case, by exploiting spatial precoding, we can serve up to  $L$  sections simultaneously. In general, each user receives interference from other users in the same section and from users in different sections. If the number of interfering sections is not larger than the number of antennas  $L$ , then this second type of interference (intergroup interference) can be zero forced by MIMO precoding, while the first type of interference (intragroup interference) is handled in the usual coded caching way: it can be canceled at the receiver because each user has all of the interfering packets in its cache, except the one it needs to decode.

Because the  $L$  sections of the delivery array are served simultaneously (by spatial multiplexing), the load is the same as that of a single MAN scheme ( $N$ ,  $M$ ,  $G$ , and  $F$ ). Therefore,

$$D = \frac{G(1-\mu)}{1+\mu G} = \frac{K/L(1-\mu)}{1+\mu K/L} = \frac{K(1-\mu)}{L+\mu K}. \quad (20)$$



**FIGURE 2.** A qualitative representation of the delivery array of the MISO broadcast channel scheme in [43], where coded caching operates “on the columns” and spatial multiplexing operates “on the rows.”

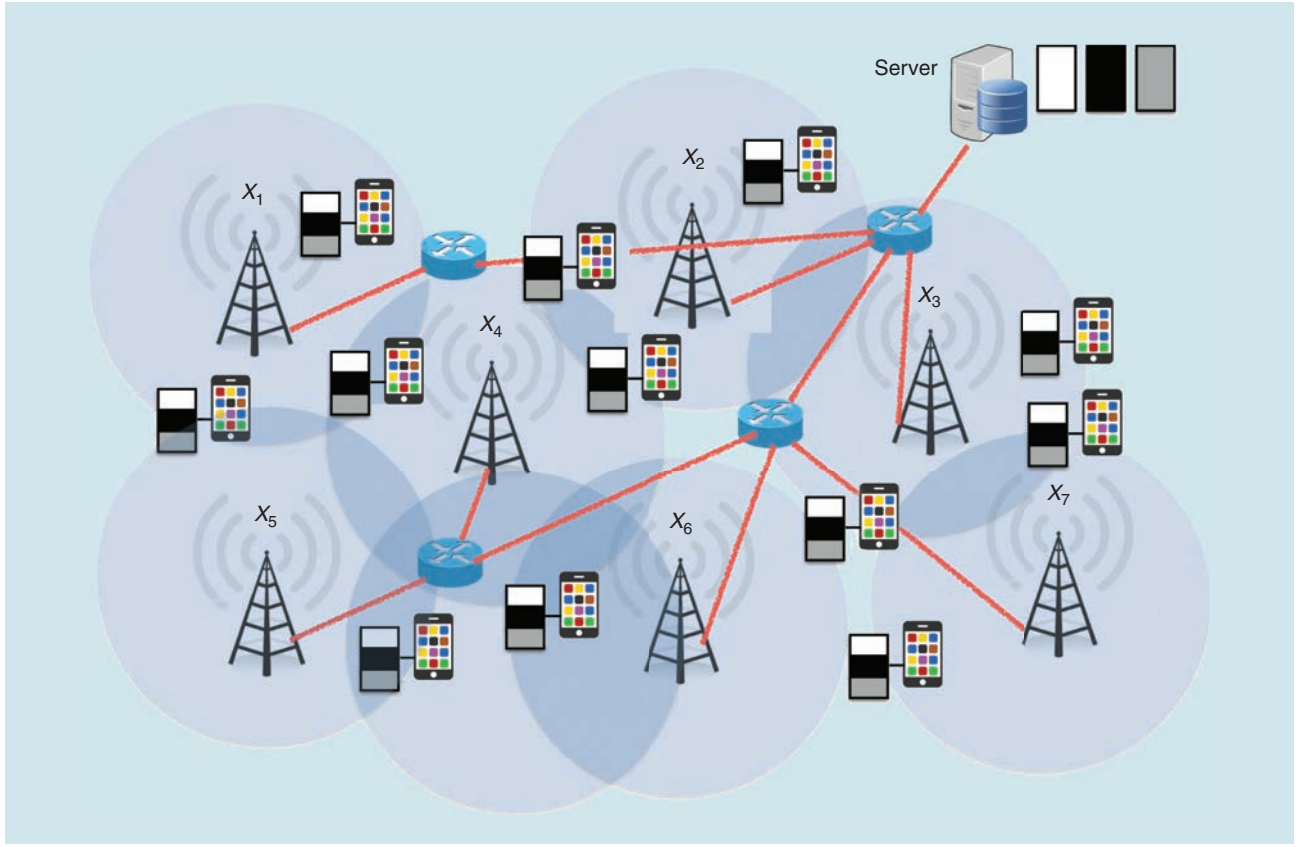
Note that the achieved DoF of this scheme are given by  $L + \mu K = L + KM/N$ , as previously obtained in the “Performance in the Homogeneous Case” section, which are known to be optimal under uncoded placement and distinct user demands.

### Exploiting spatial reuse in extended (cellular) networks

In this section, we present a simple construction able to achieve the same DoF of the single BS MISO broadcast channel, in the case of an extended cellular network with spatial reuse. Consider a multicell system covering a certain geographic area of size  $A$  with  $B$  single-antenna BSs. A population of users is initially distributed as a Poisson point process (PPP) of intensity  $\lambda$ . The number of users  $K$  in the area  $A$  is a Poisson-distributed random variable with a mean of  $\lambda A$ . The users move around the coverage area according to some random walk with independent increments. It is well known that, in this case, the marginal spatial distribution of the users at each point in time is also a PPP with intensity  $\lambda$ . We wish to design a scheme robust to mobility, i.e., the cache placement is done a priori and independent of the user positions, with the system capable of delivering the user requests for any realization of the PPP.

As discussed previously, we partition the users into  $G$  groups and use a MAN placement with parameters ( $N$ ,  $M$ ,  $G$ , and  $F$ ) by replicating the cache content for all users in the same group  $g = 1, \dots, G$ . Users select groups at random, such that each user group forms an independent thinning of the original PPP so that each user group is distributed according to an independent PPP with intensity  $\lambda/G$ . Assuming a symmetrical layout where each cell has the same area,  $A/B$ , it follows that the number of users of a given group  $g$  in each cell is an independent Poisson random variable with a mean of  $\lambda A/(BG)$ . Let us focus on a given reference cell and denote by  $\ell_1, \dots, \ell_G$  the number of users in groups  $g = 1, \dots, G$  inside the cell. As mentioned previously, these occupancy numbers are independent and Poisson distributed with the same mean:  $\lambda A/(BG)$ . The delivery process at each cell works in the same way, as explained in the “Cache Replication” section. The BS forms a delivery array with occupancy numbers  $\ell_1, \dots, \ell_G$  and serves each column of the delivery array in sequence, incurring a load given by (18). The average cell load can be trivially upper bounded by (19), where now  $\mathbb{E}[\ell_{(1)}]$  is the expected value of the maximum of  $G$ -independent Poisson random variables. Accordingly, it appears that  $\mathbb{E}[\ell_{(1)}]$  is similar to  $\lambda A/(BG)$  (the mean of an individual occupancy number), up to the logarithmic terms in  $G$ .

For large area  $A$ , the number of users  $K$  is very close to its mean, i.e.,  $K \approx \lambda A$ . Using  $\lambda = K/A$ , we find that the mean of the occupancy numbers is  $K/(BG)$ . Thus, a sensible choice for  $G$  is  $G = K/B$ , yielding that each cell contains, on average, one user per group. Using the simple upper bound (19) with this choice of  $G$ , we find the average load of each cell, that is,



**FIGURE 3.** A qualitative representation of a cache-aided multicell system, handling mobility and the small subpacketization order.

$$\begin{aligned}
 \mathbb{E}[D(\ell_1, \dots, \ell_G)] &\leq \frac{G(1-\mu)}{1+\mu G} \cdot \mathbb{E}[\ell_{(1)}] \\
 &= \frac{K/B(1-\mu)}{1+\mu K/B} \cdot \mathbb{E}[\ell_{(1)}] \\
 &= \frac{K(1-\mu)}{B+\mu K} \cdot \mathbb{E}[\ell_{(1)}]. \quad (21)
 \end{aligned}$$

Comparing (21) with (20), we see that the multicell system yields DoF equal to  $(B + \mu K)/\mathbb{E}[\ell_{(1)}]$ , where  $\mathbb{E}[\ell_{(1)}]$  is of the order of  $\log G$ . Neglecting this logarithmic term, the DoF of the multicell system take on the same form as the DoF of a single-cell, cache-aided MISO broadcast channel with  $L = B$  antennas. Because the subpacketization is formed for  $G = K/B$  groups, as long as the number of users per cell  $K/B$  is a constant that does not grow with  $K$ , this system achieves a fixed subpacketization order  $\approx \exp(G\mathcal{H}(\mu))$ , while both the number of users  $K$  and the number of cells  $B$  grows arbitrarily large as long as they grow with fixed-ratio  $G$ . Note that present cellular systems are designed to balance the number of users per cell such that each cell handles a constant number of users to avoid congestion; therefore, the operating conditions of  $K/B = O(1)$  are realistic for a well-designed cellular system.

Finally, we observe that the aforementioned simple analysis is done with the assumption that all cells can operate simultaneously on the same frequency band (frequency reuse 1). If the intercell interference resulting from reuse 1 is too large,

then a standard frequency reuse scheme with some reuse factor  $m$  can be employed. In this case, the load is increased by  $m$ , which is typically a small integer (e.g., for the classical hexagonal layout typical values of  $m$  are 3, 4, or 7). Figure 3 shows the architecture of a multicell network based on these ideas. A more refined analysis of the achievable average load as well as a mixed-integer linear programming optimization problem for the case of reuse 1 and intercell interference modeled by the so-called protocol model, which serves as a simple conceptual model for a collision-based all-or-nothing interference channel (see the definition in [2]), is given in [44].

### Coded caching: A broader picture

The basic MAN scheme was first introduced for the single shared-link network in [1]. Since then, several important follow-up works have appeared in the literature and present different aspects of coded caching. For the sake of completeness, in this section we briefly review some of the main challenges of coded caching addressed in these works.

#### Centralized versus decentralized caching

In many applications, the centralized placement of the MAN scheme is not practically feasible. For example, the set of users present in the network may change from the placement to the delivery phase because users join or leave the network in a dynamic manner. A decentralized caching scheme is proposed in [3] where the placement phase for each user is

performed individually and independently from other users. More precisely, each user  $k$  stores  $MF/N$  packets from each file, chosen uniformly at random, and independently across the files and users. Similar to a centralized scheme, the server tries to maximize the utility of a multicasting message by combining requested packets; however, in the absence of a centralized placement, a packet intended for user  $k$  may be cached at a random number of other users, rather than at exactly  $\alpha = KM/N$  users. Therefore, a wide range of utilities for multicasting packets is available, which results in an expected delivery time of

$$T_{\text{decentralized-caching}} = \frac{D \cdot F}{R} = K \left(1 - \frac{M}{N}\right) \frac{N}{KM} \left(1 - \left(1 - \frac{M}{N}\right)^K\right) \frac{F}{R}. \quad (22)$$

This leads to a loss compared to the delivery time of the centralized case [3].

### Fundamental limits and optimality

The scheme proposed in [1] offers a significant global gain over an uncoded caching scheme that serves the users individually. However, a priori it is not clear whether it can be further improved using a more sophisticated placement and delivery scheme. Characterizing the optimum gain and the exact tradeoff between the cache size and the delivery load are not yet fully addressed. A cut-set-type argument was provided in [1], which proves that the basic MAN scheme is within a constant multiplicative gap from the optimum scheme. Several tighter outer bounds on the optimum tradeoff have been developed [45]–[50].

In particular, Wan et al. in [51] and [52] proved that the basic MAN scheme is optimal if all of the following conditions are fulfilled:

- 1) the cache contents are limited to being collections of segments of the files without any precoding (uncoded placement)
  - 2) the cache contents are jointly and centrally optimized (centralized caching)
  - 3) the user requests are all distinct (the worst demand profile).
- This result was generalized in [42], where assumptions 2) and 3) are relaxed. More precisely, the optimum exact tradeoff between the cache size and the delivery load under the assumption of uncoded placement is characterized in [42]. It is shown that, when there is no overlap between the users' requests, the schemes of [1] and [3] are optimum for centralized and decentralized caching, respectively. Moreover, a novel caching strategy is introduced in [42] to exploit commonality among user demands and improve upon the gain of the basic MAN scheme, which shows that the new scheme is information-theoretically optimum. This fully characterizes the optimum tradeoff for the uncoded placement and the single shared-link network.

### Coded versus uncoded placement

Even in the original work [1], it was observed that a coded placement [i.e., when the data placed in the caches are functions (e.g., XORs) of the original files] can improve the overall system performance and further reduce the load of delivery. Characterizing the optimum tradeoff and developing cache de-

sign for coded placement under centralized setting is studied in [53]–[56]. The common feature in the proposed cache designs is interfile coding, which allows for the combining of packets from different files and caching the coded packets during the placement phase. Note that such coded prefetched packets will be useless if the interference cannot be canceled during the delivery phase. In contrast, intrafile coded placement is introduced in [4] and [57], where it is demonstrated that individually encoding the files in the database using an erasure code can reduce the delivery load in the decentralized scheme. This is further improved upon using subspace coding in [58]. Finally, [59] develops an information-theoretic converse bound (infeasibility) that applies to any placement (coded or uncoded) and proves the optimality of the scheme in [60] within a factor of 2, which means that any more-complicated coded placements can gain at most a factor of 2 in the load with respect to the conceptually simpler uncoded placement of [60].

### Network topologies

Beyond the single shared-link network considered in [1], several other coded caching network topologies have been studied in the literature. Here we summarize the most popular ones, for which often exact optimality or order optimality (i.e., the minimum worst-case load optimality up to multiplicative factors) have been determined.

#### Tree networks

In [1], the single shared-link network is generalized to a tree network where the server is at the root and the users are at the leaves. Intermediate nodes simply route the XOR-ed packets from one tree layer to the next. It is shown that a combination of routing and the original MAN scheme is order optimal for the tree network. The routing algorithm is very simple: consider an intermediate node in the tree at layer  $\ell$ . Such a node receives XOR-ed packets,  $X_s$ , from its parent at layer  $\ell - 1$ , and routes them to its  $i$ th child at layer  $\ell + 1$  whenever at least one user  $k \in \mathcal{S}$  is present in the subtree rooted at the  $i$ th child. That is, an XOR-ed packet is passed “down” to a node if it is useful for at least one (grand)child of that node.

#### Hierarchical two-level network

In [5], a network formed by the server, a layer of relays, and a layer of end users is considered. The server communicates with the relays via a single shared-link network and each relay also communicates with a subset of users via a “local” single shared-link network. Each user receives from only one relay. Caching memory is present at the relays and at the users. This network is called *hierarchical coded caching* because it is composed of a two-level hierarchy of single shared-link networks.

#### Shared caches network with arbitrary occupancy numbers

In [41], a variant of the single shared-link network is considered, where a server communicates with a layer of intermediate nodes, each of which has cache  $M$  via a single shared-link network. Each intermediate node serves a different number of possible users via ideal infinite capacity links. The number of users connected to a



given intermediate node is denoted as the *occupancy number* of that node. The model may be motivated by a network of small-cell BSs with caches, which receive information from a controlling server or macro BS via a broadcast link (single shared-link between server and intermediate nodes), and serve their own users independently via a much faster local “access network,” in each small cell. This model is formally identical to the case of shared user caches where the group of users connected to the same intermediate node actually share the same cache. It is also isomorphic to the case of users with identical copies of the cache, as originated by cache replication discussed more extensively in the “Exploiting Spatial Reuse in Extended (Cellular) Networks” section. In fact, all of the users connected to the same intermediate node behave as if they had their own individual cache, but the placement scheme replicates the same cache content in each one of them. Furthermore, the model is also related to the case of multiple requests; in fact, we can identify each individual node as a user, but each user makes multiple requests (one for each of the actual users connected to the intermediate node).

#### Multiserver linear network

This topology, presented and studied in [8], considers  $L \geq 1$  servers, each of which has access to the full file library, serving  $K$  users, each of which has a cache size of  $M$ . The relation between the  $L$  inputs and the  $K$  outputs is given by  $\mathbf{y} = \mathbf{H}\mathbf{x}$ , where  $\mathbf{y} \in \mathbb{F}_q^K$ ,  $\mathbf{x} \in \mathbb{F}_q^L$  and  $\mathbf{H} \in \mathbb{F}_q^{K \times L}$  are defined over a (typically large) finite field,  $\mathbb{F}_q$ . The rationale for this model is that the  $L$  servers communicate to the  $K$  users via some network for which end-to-end linear network coding is used instead of Internet Protocol routing. In this way, each user receives a linear combination of the information packets sent by the servers. The finite field size is chosen such that matrix  $\mathbf{H}$  has rank  $\min\{L, K\}$  with high probability when the network coding combination coefficients are chosen randomly.

The results for the multiserver linear network apply immediately to the case of a physical MISO downlink channel where the server is colocated with the BS and has  $L$  antennas, and the  $K$  users have a single antenna each. The Gaussian multiuser MISO version of the problem was studied in [25] and [43] and is examined in greater detail in the “Signal Processing Problem Formulation” section.

#### Gaussian interference channel

A generalization of the MISO downlink coded caching problem is a scenario in which the transmit antennas are separated transmitters, each of which has an individual cache of size  $M_T$  not necessarily equal to  $M$  (cache at the receivers). In this case, a necessary condition for the successful delivery of any user request is that the whole library can be stored in the network, i.e.,  $M_T L + MK \geq N$ . A one-shot precoding solution for this network was provided in [32], while more elaborate schemes based on interference alignment with dimension expansion or signal-level expansion was presented in [61]. A recent extension of the one-shot precoding scheme to the case of nonfully connected interference channels arising from a cellular topology is presented in [62].

#### D2D coded caching

In [2], a D2D version of the coded caching problem is proposed and an order-optimal scheme is provided. The D2D network consists of a shared ideal channel where all the nodes can broadcast to all the other nodes, but only one node can talk at a time. This network may be motivated by a carrier sense multiple access D2D scheme (e.g., Wi-Fi Direct) or a token-ring medium access control protocol, where a collision avoidance mechanism permits only one node to be active at a time. However, when a node is active, all the other nodes can listen and decode its transmission. A necessary condition for the feasibility of the D2D coded caching network is that the whole library can be stored in the network, i.e.,  $KM \geq N$ .

#### Combination network

The combination network consists of a server, a layer of  $L$  relays, and a layer of users. For a certain degree of connectivity,  $r$ , there are exactly  $K = \binom{L}{r}$  users, one for each distinct combination of  $r$  relays. All the links connecting the server to relays and relays to users are orthogonal, i.e., there are no broadcast or interference constraints. In particular, a user connected to  $r$  relays can simultaneously receive the  $r$ -transmitted signals from these relays without interference. Coded caching for combination networks is studied in [63], where it is shown that a speed-up factor of  $1/r$  in the delivery time with respect to the single shared-link network is possible for this network. Building on the combination network, in [64], a scheme for a multicell system with macrodiversity order  $r$  is proposed and analyzed in the case of MISO fading channels and distance-dependent path loss. Variants of the combination network, including the case of caches at the relays, have been analyzed in [65], while the improved strategies and information-theoretic optimality results are given in [66]–[69].

### Challenges and open issues

The field of cache-aided communication is still in its infancy, and its challenges outnumber its achievements. Although research has shown significant throughput gains in various scenarios, questions remain unanswered and many issues must be resolved to allow for practical implementation. In this section, we describe some of these challenges and open issues.

#### Physical layer

Even though it is easy to obtain an ideal DoF scheme, the best cache-aided communication scheme for the general case is not yet known. An optimality proof exists only for the single-antenna homogenous case. For the multiple antenna case, even the linear optimal scheme is not yet known (all the results presented for this case are based on ZF). Research advances in this direction are needed to allow for a better understanding of the capabilities and to enable the efficient implementation of cache-aided communication schemes also at a low SNR.

More effort is also required to deal with physical-layer practicalities. For example, the effect of an imperfect channel state at the BS has been studied in only limited scenarios. In MIMO systems, such uncertainty will affect the ZF accuracy and may

be addressed by a variety of known signal processing techniques. More importantly, this will create a rate uncertainty at the BS, which can have significant effects on transmission scheduling.

A scheme that can handle channel variations during the transmission stage has a significant advantage because this scheme will enable cache-aided communication in a wide area of cellular networks where the channel rates vary relatively fast. A simpler variation can apply cache-aided communication in networks where the rate changes over time but its average is known in advance.

### *Beyond linear precoding*

For multiantenna BSs, we have discussed the use of linear precoding only. Linear precoding for cache-aided communications still requires further study; thus far, only ZF precoding has been studied extensively. Other approaches beyond linear precoding require even more attention.

Nonlinear precoding was thus far considered through two opposite approaches: on one hand, transmitting a single message at a time, where this message is an XOR of multiple file segments that can be used by multiple users (where each user can use its cache to cancel out the unintended file segments). This message should be beamformed to a direction that is favorable to all users served and can obtain some amount of array gain [9]. On the other hand, to obtain a multiplexing gain, the same work suggests serving multiple groups of users simultaneously, where the transmission to different groups is separated using ZF precoding, and each user in a group can extract its own message by XORing with its cache content. Neither of these schemes is optimal and combining both approaches in a single system to optimize the balance for specific network conditions can possibly offer a better performance. Nevertheless, the best approach for combining XORing and linear precoding remains unknown.

Taking it one step further, there is still much room for improvement using sophisticated nonlinear precoding schemes. For example, dirty paper coding and vector perturbation are more energy efficient than is linear ZF. Although not previously considered, the combinations of such schemes with multiantenna cache-aided communications may lead to significant gains, particularly in the low-to-medium SNR regime, where the ZF may be inefficient.

### *Scheduling and resource allocation*

As shown in the previous section, cache-aided communication is much simpler in a homogenous network (where all the user rates are identical and all files are of the same popularity). Even in this case, existing scheduling methods require a very thin subpacketization, which, in most cases becomes unpractical for a large number of users. Thus, the research for scheduling approaches that will require a smaller number of packets is ongoing.

In a nonhomogenous case, the situation is even more complicated, and the only known solutions require solving a large optimization problem. Hence, a scheduling algorithm with an acceptable complexity that can handle network inhomogeneities is needed. In particular, if the user rates are not identical, each file segment is encoded to a different number of symbols

that depend on the rate of the requesting user. Accordingly, the transmission scheduling must allocate more transmissions to users with a low rate, while still trying to serve a maximal number of users simultaneously at any given time.

Furthermore, in nonhomogenous networks, resource allocation is also a major challenge for cache-aided communications. Existing approaches are either highly inefficient or very complicated to implement. As a result, there is an acute need for low-complexity resource-allocation schemes (optimal or suboptimal) that enable the benefits of cache-aided communication in practical systems.

Additionally, further performance analysis and closed-form performance expressions of cache-aided communications in nonhomogeneous networks are needed. Such expressions are required to better predict the performance in various networks and for network planning. Performance expressions are also needed for network optimization, e.g., for power allocation, parameter selection, and so on.

### *Higher layers*

Another major difficulty in cache-aided communications is the necessity of dividing the data into many subpackets. As discussed previously, this problem can be made simpler when using multiple antennas at the BS. Yet, this approach requires additional research, particularly for the nonhomogeneous case.

Another high-layer issue that is crucial for practical implementation is the handling of network and content dynamics. Users' disconnection or movement from one BS (cell) to another causes changes in the network connectivity. Similarly, content dynamics can occur, for example, as a result of variations in the popularity of files. A practical network will likely meet all types of dynamics and must be robust to such changes. These aspects have hardly been addressed thus far and still require much research, e.g., How can the cache content be updated at minimal overhead? Are there schemes that allow for such a cache adaptation at low complexity? What are the performance costs of such a scheme?

As for network dynamics, schemes are needed for the adaptation of the transmission scheduling following a network change. Such schemes can consider intermediate planning (i.e., adapting to changes that happen between the cache placement and transmission stages) and online planning (i.e., adapting to changes that occur during the transmission phase). The development of such schemes is likely to be the final catalyst for the practical implementation of cache-aided communications.

## **Conclusions**

Cache-aided communications have shown significant potential for throughput increase in wideband communication networks. The possibility of using the data stored at one user, even if they are only requested by another user, allows for combining the small size memories employed at different users and using them as an effectively large cache. Because the network performance depends on the total memory size of all the users, the network throughput scales linearly with

the number of users. Thus, cache-aided communication is expected to take a significant role in large networks.

Yet, many challenges must be overcome prior to practical implementation. These challenges are mostly in the field of signal processing, and include low-complexity optimization, practical system design, and the handling of network imperfections. This article aimed at presenting this promising technology in a tractable manner that reflects its potential and open challenges.

## Authors

**Soheil Mohajer** (soheil@umn.edu) received his B.Sc. degree in electrical engineering from Sharif University of Technology, Iran, in 2004 and his M.Sc. and Ph.D. degrees in communication systems, both from École Polytechnique Fédérale de Lausanne, Switzerland, in 2005 and 2010, respectively. He is currently an assistant professor in the Department of Electrical and Computer Engineering at the University of Minnesota, Twin Cities. Previously, he was a postdoctoral researcher at Princeton University, New Jersey (2010–2011), and the University of California, Berkeley (2011–2013). He received the National Science Foundation CAREER Award in 2018. He currently serves as an editor of *IEEE Transactions on Communications*. His research interests include information theory and its applications in distributed storage systems, wireless networks, bioinformatics, and statistical machine learning. He is a Member of the IEEE.

**Itsik Bergel** (itsik.bergel@biu.ac.il) received his B.Sc. degree in electrical engineering and his B.Sc. degree in physics from Ben Gurion University, Beersheba, Israel, in 1993 and 1994, respectively, and his M.Sc. and Ph.D. degrees in electrical engineering from Tel Aviv University, Israel, in 2000 and 2005, respectively. Currently, he is with the Faculty of Engineering, Bar-Ilan University, Israel. From 2001 to 2003, he was a senior researcher with the Intel Communications Research Laboratory. In 2005, he was a postdoctoral researcher at Politecnico di Torino, Italy. He has been an associate editor of *IEEE Transactions on Signal Processing* and currently serves as an editor of *IEEE Transactions on Wireless Communications*. His main research interests include interference mitigation in wireline and wireless communications, cooperative transmission, and cross-layer optimization of random ad hoc networks. He is a Senior Member of the IEEE.

**Giuseppe Caire** (caire@tu-berlin.de) received his B.Sc. degree in electrical engineering from Politecnico di Torino, Turin, Italy, in 1990, his M.Sc. degree in electrical engineering from Princeton University, New Jersey, in 1992, and his Ph.D. degree from Politecnico di Torino in 1994. He is currently an Alexander von Humboldt Professor with the Faculty of Electrical Engineering and Computer Science at the Technical University of Berlin, Germany. He received the Jack Neubauer Best System Paper Award from the IEEE Vehicular Technology Society in 2003, IEEE Communications Society & Information Theory Society Joint Paper Award in 2004 and in 2011, IEEE Communications Society Leonard G. Abraham Prize for best *IEEE Journal on Selected Areas in Communications* paper in 2019, Okawa Research Award in 2006, Alexander von

Humboldt Professorship in 2014, Vodafone Innovation Prize in 2015, and European Research Council Advanced Grant in 2018. He served as an officer on the Board of Governors of the IEEE Information Theory Society and, in 2011, was elected president of the IEEE Information Theory Society. He is a Fellow of the IEEE.


## References

- [1] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, 2014. doi: 10.1109/TIT.2014.2306938.
- [2] M. Ji, G. Caire, and A. F. Molisch, "Fundamental limits of caching in wireless D2D networks," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 849–869, 2016. doi: 10.1109/TIT.2015.2504556.
- [3] M. A. Maddah-Ali and U. Niesen, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *IEEE/ACM Trans. Netw.*, vol. 23, no. 4, pp. 1029–1040, 2015. doi: 10.1109/TNET.2014.2317316.
- [4] H. Reiszadeh, M. A. Maddah-Ali, and S. Mohajer, "Erasure coding for decentralized coded caching," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2018, pp. 1715–1719. doi: 10.1109/ISIT.2018.8437574.
- [5] N. Karamchandani, U. Niesen, M. A. Maddah-Ali, and S. N. Diggavi, "Hierarchical coded caching," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3212–3229, 2016. doi: 10.1109/TIT.2016.2557804.
- [6] R. Tandon and O. Simeone, "Cloud-aided wireless networks with edge caching: Fundamental latency trade-offs in fog radio access networks," in *Proc. 2016 IEEE Int. Symp. Information Theory (ISIT)*, pp. 2029–2033. doi: 10.1109/ISIT.2016.7541655.
- [7] N. Mital, D. Gündüz, and C. Ling, "Coded caching in a multi-server system with random topology," in *Proc. 2018 IEEE Wireless Communications and Networking Conf. (WCNC)*, pp. 1–6. doi: 10.1109/WCNC.2018.8377365.
- [8] S. P. Shariatpanahi, S. A. Motahari, and B. H. Khalaj, "Multi-server coded caching," *IEEE Trans. Inf. Theory*, vol. 62, no. 12, pp. 7253–7271, 2016. doi: 10.1109/TIT.2016.2614722.
- [9] S. P. Shariatpanahi, G. Caire, and B. H. Khalaj, "Multi-antenna coded caching," in *Proc. 2017 IEEE Int. Symp. Information Theory (ISIT)*, pp. 2113–2117. doi: 10.1109/ISIT.2017.8006902.
- [10] I. Bergel and S. Mohajer, "Cache-aided communications with multiple antennas at finite SNR," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 8, pp. 1682–1691, 2018. doi: 10.1109/JSAC.2018.2844618.
- [11] K.-H. Ngo, S. Yang, and M. Kobayashi, "Cache-aided content delivery in MIMO channels," in *Proc. 2016 54th IEEE Annu. Allerton Conf. Communication, Control, and Computing (Allerton)*, pp. 93–100. doi: 10.1109/ALLERTON.2016.7852215.
- [12] S. S. Bidokhti, M. Wigger, and R. Timo, "Noisy broadcast networks with receiver caching," *IEEE Trans. Inf. Theory*, vol. 64, no. 11, pp. 6996–7016, 2018. doi: 10.1109/TIT.2018.2835507.
- [13] S. S. Bidokhti, M. Wigger, and A. Yener, "Benefits of cache assignment on degraded broadcast channels," in *Proc. 2017 IEEE Int. Symp. Information Theory (ISIT)*, pp. 1222–1226. doi: 10.1109/ISIT.2017.8006723.
- [14] K.-H. Ngo, S. Yang, and M. Kobayashi, "Scalable content delivery with coded caching in multi-antenna fading channels," *IEEE Trans. Wireless Commun.*, vol. 17, no. 1, pp. 548–562, 2018. doi: 10.1109/TWC.2017.2768361.
- [15] S. Yang, K.-H. Ngo, and M. Kobayashi, "Content delivery with coded caching and massive MIMO in 5G," in *Proc. 2016 9th IEEE Int. Symp. Turbo Codes and Iterative Information Processing (ISTC)*, pp. 370–374. doi: 10.1109/ISTC.2016.7593139.
- [16] M. M. Amiri and D. Gündüz, "Cache-aided content delivery over erasure broadcast channels," *IEEE Trans. Commun.*, vol. 66, no. 1, pp. 370–381, 2018. doi: 10.1109/TCOMM.2017.2751608.
- [17] M. Ji, G. Caire, and A. F. Molisch, "Wireless device-to-device caching networks: Basic principles and system performance," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 1, pp. 176–189, 2016. doi: 10.1109/JSAC.2015.2452672.
- [18] D. Liu, B. Chen, C. Yang, and A. F. Molisch, "Caching at the wireless edge: Design aspects, challenges, and future directions," *IEEE Commun. Mag.*, vol. 54, no. 9, pp. 22–28, 2016. doi: 10.1109/MCOM.2016.7565183.
- [19] E. Zeydan, E. Bastug, M. Bennis, M. A. Kader, I. A. Karatepe, A. S. Er, and M. Debbah, "Big data caching for networking: Moving from cloud to edge," *IEEE Commun. Mag.*, vol. 54, no. 9, pp. 36–42, 2016. doi: 10.1109/MCOM.2016.7565185.
- [20] G. S. Paschos, G. Iosifidis, M. Tao, D. Towsley, and G. Caire, "The role of caching in future communication systems and networks. 2018. [Online]. Available: arXiv:1805.11721
- [21] A. Meyerson, K. Munagala, and S. Plotkin, "Web caching using access statistics," in *Proc. 12th Annu. ACM-SIAM Symp. Discrete Algorithms*. Philadelphia: Society for Industrial and Applied Mathematics, 2001, pp. 354–363.



- [22] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. 2010 IEEE INFOCOM*, pp. 1–9. doi: 10.1109/INFCOM.2010.5461964.
- [23] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," in *Proc. IEEE INFOCOM'99, 18th Annu. Joint Conf. Computer and Communications Societies*, 1999, vol. 1, pp. 126–134. doi: 10.1109/INFCOM.1999.749260.
- [24] J. Wang, "A survey of web caching schemes for the internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 5, pp. 36–46, 1999. doi: 10.1145/505696.505701.
- [25] S. P. Shariatpanahi, G. Caire, and B. H. Khalaj, "Physical-layer schemes for wireless coded caching," *IEEE Trans. Inf. Theory*, vol. 65, no. 5, pp. 2792–2807, 2018. doi: 10.1109/TIT.2018.2888615.
- [26] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Order-optimal rate of caching and coded multicasting with random demands," *IEEE Trans. Inf. Theory*, vol. 63, no. 6, pp. 3923–3949, 2017. doi: 10.1109/TIT.2017.2695611.
- [27] R. Pedarsani, M. A. Maddah-Ali, and U. Niesen, "Online coded caching," *IEEE/ACM Trans. Netw.*, vol. 24, no. 2, pp. 836–845, 2016. doi: 10.1109/TNET.2015.2394482.
- [28] E. Bastug, M. Bennis, and M. Debbah, "Living on the edge: The role of proactive caching in 5G wireless networks," *IEEE Commun. Mag.*, vol. 52, no. 8, pp. 82–89, 2014. doi: 10.1109/MCOM.2014.6871674.
- [29] S. Mohajer and I. Bergel, "Optimal power allocation in MISO cache-aided communication," in *Proc. 2018 IEEE 19th Int. Workshop Signal Processing Advances Wireless Communications (SPAWC)*, pp. 1–5. doi: 10.1109/SPAWC.2018.8445770.
- [30] S. Jin, Y. Cui, H. Liu, and G. Caire, "Structural properties of uncoded placement optimization for coded delivery." 2017. [Online]. Available: arXiv:1707.07146
- [31] Q. Yang and D. Gündüz, "Coded caching and content delivery with heterogeneous distortion requirements," *IEEE Trans. Inf. Theory*, vol. 64, no. 6, pp. 4347–4364, 2018. doi: 10.1109/TIT.2018.2805331.
- [32] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, "Fundamental limits of cache-aided interference management," *IEEE Trans. Inf. Theory*, vol. 63, no. 5, pp. 3092–3107, 2017. doi: 10.1109/TIT.2017.2669942.
- [33] K. Shanmugam, M. Ji, A. M. Tulino, J. Llorca, and A. G. Dimakis, "Finite-length analysis of caching-aided coded multicasting," *IEEE Trans. Inf. Theory*, vol. 62, no. 10, pp. 5524–5537, 2016. doi: 10.1109/TIT.2016.2599110.
- [34] Q. Yan, M. Cheng, X. Tang, and Q. Chen, "On the placement delivery array design for centralized coded caching scheme," *IEEE Trans. Inf. Theory*, vol. 63, no. 9, pp. 5821–5833, Sept. 2017. doi: 10.1109/TIT.2017.2725272.
- [35] Q. Yan, X. Tang, and Q. Chen, "On the placement and delivery schemes for decentralized coded caching system." 2017. [Online]. Available: arXiv:1710.04884
- [36] S. Jin, Y. Cui, H. Liu, and G. Caire, "Uncoded placement optimization for coded delivery," in *Proc. 2018 IEEE 16th Int. Symp. Modeling and Optimization Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, pp. 1–8. doi: 10.23919/WIOPT.2018.8362816.
- [37] K. Shanmugam, A. M. Tulino, and A. G. Dimakis, "Coded caching with linear subpacketization is possible using Ruzsa-Szemerédi graphs," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2017, pp. 1237–1241. doi: 10.1109/ISIT.2017.8006726.
- [38] L. Tang and A. Ramamoorthy, "Low subpacketization schemes for coded caching," in *Proc. 2017 IEEE Int. Symp. Information Theory (ISIT)*, pp. 2790–2794. doi: 10.1109/ISIT.2017.8007038.
- [39] S. Jin, Y. Cui, H. Liu, and G. Caire, "Order-optimal decentralized coded caching schemes with good performance in finite file size regime," in *Proc. IEEE Global Communications Conf. (GLOBECOM)*, Dec. 2016, pp. 1–7. doi: 10.1109/GLOCOM.2016.7842115.
- [40] S. Jin, Y. Cui, H. Liu, and G. Caire, "New order-optimal decentralized coded caching schemes with good performance in the finite file size regime." 2016. [Online]. Available: arXiv:1604.07648
- [41] E. Parrinello, A. Unsul, and P. Elia, "Fundamental limits of caching in heterogeneous networks with uncoded prefetching." 2018. [Online]. Available: arXiv:1811.06247
- [42] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "The exact rate-memory tradeoff for caching with uncoded prefetching," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 1281–1296, 2018. doi: 10.1109/TIT.2017.2785237.
- [43] E. Lampsiris and P. Elia, "Adding transmitters dramatically boosts coded-caching gains for finite file sizes," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1176–1188, 2018. doi: 10.1109/JSAC.2018.2844960.
- [44] A. Asadzadeh and G. Caire, "Coded caching with small subpacketization via spatial reuse and content base replication," in *Proc. 2019 IEEE Int. Symp. Information Theory (ISIT)*, pp. 2982–2986. doi: 10.1109/ISIT.2019.8849460.
- [45] C. Tian, "Symmetry, outer bounds, and code constructions: A computer-aided investigation on the fundamental limits of caching," *Entropy*, vol. 20, no. 8, p. 603, 2018. doi: 10.3390/e20080603.
- [46] C.-Y. Wang, S. H. Lim, and M. Gastpar, "A new converse bound for coded caching," in *Proc. 2016 IEEE Information Theory and Applications Workshop (ITA)*, pp. 1–6. doi: 10.1109/ITA.2016.7888186.
- [47] H. Ghasemi and A. Ramamoorthy, "Improved lower bounds for coded caching," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4388–4413, 2017. doi: 10.1109/TIT.2017.2705166.
- [48] A. Sengupta and R. Tandon, "Improved approximation of storage-rate tradeoff for caching with multiple demands," *IEEE Trans. Commun.*, vol. 65, no. 5, pp. 1940–1955, 2017. doi: 10.1109/TCOMM.2017.2664815.
- [49] S. H. Lim, C.-Y. Wang, and M. Gastpar, "Information-theoretic caching: The multi-user case," *IEEE Trans. Inf. Theory*, vol. 63, no. 11, pp. 7018–7037, 2017. doi: 10.1109/TIT.2017.2733527.
- [50] N. Ajaykrishnan, N. S. Prem, V. M. Prabhakaran, and R. Vaze, "Critical database size for effective caching," in *Proc. IEEE 2015 21st Nat. Conf. Communications (NCC)*, pp. 1–6. doi: 10.1109/NCC.2015.7084871.
- [51] K. Wan, D. Tuninetti, and P. Piantanida, "On caching with more users than files," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2016, pp. 135–139. doi: 10.1109/ISIT.2016.7541276.
- [52] K. Wan, D. Tuninetti, and P. Piantanida, "On the optimality of uncoded cache placement," in *Proc. 2016 IEEE Information Theory Workshop (ITW)*, pp. 161–165. doi: 10.1109/ITW.2016.7606816.
- [53] Z. Chen, P. Fan, and K. B. Letaief, "Fundamental limits of caching: Improved bounds for small buffer users." 2014. [Online]. Available: arXiv:1407.1935
- [54] C. Tian and J. Chen, "Caching and delivery via interference elimination," *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1548–1560, 2018. doi: 10.1109/TIT.2018.2794543.
- [55] M. M. Amiri and D. Gündüz, "Fundamental limits of coded caching: Improved delivery rate-cache capacity tradeoff," *IEEE Trans. Commun.*, vol. 65, no. 2, pp. 806–815, 2017. doi: 10.1109/TCOMM.2016.2638841.
- [56] S. Sahraei and M. Gastpar, "K users caching two files: An improved achievable rate," in *Proc. IEEE 2016 Annu. Conf. Information Science and Systems (CISS)*, pp. 620–624. doi: 10.1109/CISS.2016.7460574.
- [57] Y.-P. Wei and S. Ulukus, "Novel decentralized coded caching through coded prefetching," in *Proc. 2017 IEEE Information Theory Workshop (ITW)*, pp. 1–5. doi: 10.1109/ITW.2017.8278044.
- [58] H. Reisizadeh, M. A. Maddah-Ali, and S. Mohajer, "Subspace coding for coded caching: Decentralized and centralized placements meet for three users," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2019, pp. 677–681. doi: 10.1109/ISIT.2019.8849613.
- [59] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Characterizing the rate-memory tradeoff in cache networks within a factor of 2," *IEEE Trans. Inf. Theory*, vol. 65, no. 1, pp. 647–663, 2019. doi: 10.1109/TIT.2018.2870566.
- [60] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "The exact rate-memory tradeoff for caching with uncoded prefetching." 2016. [Online]. Available: arXiv:1609.07817
- [61] J. Hachem, U. Niesen, and S. Diggavi, "Degrees of freedom of cache-aided wireless interference networks," *IEEE Trans. Inf. Theory*, vol. 64, no. 7, pp. 5359–5380, 2018. doi: 10.1109/TIT.2018.2825321.
- [62] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, "Cache-aided interference management in wireless cellular networks," *IEEE Trans. Commun.*, vol. 67, no. 5, pp. 3376–3387, 2019. doi: 10.1109/TCOMM.2019.2893669.
- [63] M. Ji, M. F. Wong, A. M. Tulino, J. Llorca, G. Caire, M. Effros, and M. Langberg, "On the fundamental limits of caching in combination networks," in *Proc. 2015 IEEE 16th Int. Workshop Signal Processing Advances Wireless Communications (SPAWC)*, pp. 695–699. doi: 10.1109/SPAWC.2015.7227127.
- [64] M. Bayat, R. K. Mungara, and G. Caire, "Achieving spatial scalability for coded caching via coded multipoint multicasting," *IEEE Trans. Wireless Commun.*, vol. 18, no. 1, pp. 227–240, 2019. doi: 10.1109/TWC.2018.2878845.
- [65] A. A. Zewail and A. Yener, "Coded caching for combination networks with cache-aided relays," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, June 2017, pp. 2433–2437. doi: 10.1109/ISIT.2017.8006966.
- [66] K. Wan, D. Tuninetti, M. Ji, and P. Piantanida, "State-of-the-art in cache-aided combination networks," in *Proc. 2017 IEEE 51st Asilomar Conf. Signals, Systems, and Computers*, pp. 641–645. doi: 10.1109/ACSSC.2017.8335420.
- [67] K. Wan, M. Ji, P. Piantanida, and D. Tuninetti, "Novel outer bounds and inner bounds with uncoded cache placement for combination networks with end-user-caches." 2017. [Online]. Available: arXiv:1701.06884
- [68] K. Wan, M. Ji, P. Piantanida, and D. Tuninetti, "On the benefits of asymmetric coded cache placement in combination networks with end-user caches," in *Proc. 2018 IEEE Int. Symp. Information Theory (ISIT)*, pp. 1550–1554. doi: 10.1109/ISIT.2018.8437462.
- [69] K. Wan, M. Ji, P. Piantanida, and D. Tuninetti, "Caching in combination networks: Novel multicast message generation and delivery by leveraging the network topology," in *Proc. 2018 IEEE Int. Conf. Communications (ICC)*, pp. 1–6. doi: 10.1109/ICC.2018.8422197.





©ISTOCKPHOTO.COM/MATEJMO

Yuejie Chi and Maxime Ferreira Da Costa

# Harnessing Sparsity Over the Continuum

*Atomic norm minimization for superresolution*

**A**t the core of many sensing and imaging applications, the signal of interest can be modeled as a linear superposition of translated or modulated versions of some template [e.g., a point spread function (PSF) or a Green's function] and the fundamental problem is to estimate the translation or modulation parameters (e.g., delays, locations, or Dopplers) from noisy measurements. This problem is centrally important to not only

target localization in radar and sonar, channel estimation in wireless communications, and direction-of-arrival estimation in array signal processing, but also modern imaging modalities such as superresolution single-molecule fluorescence microscopy, nuclear magnetic resonance imaging, and spike localization in neural recordings, among others.

Typically, the temporal or spatial resolution of the acquired signal is limited by the sensing or imaging devices, due to factors such as the numerical aperture of a microscope, the wavelength of the impinging electromagnetic or optical waves, or the sampling

Digital Object Identifier 10.1109/MSP.2019.2962209  
Date of current version: 26 February 2020

rate of an analog-to-digital converter. This resolution limit is well known and often referred to as the *Rayleigh limit (RL)* (see “What Is the Rayleigh Limit?”). The performance of matched filtering, or periodogram, which creates a correlation map of the acquired signal against the range of parameters, is limited by the RL, regardless of the noise level.

On the other hand, the desired resolution of parameter estimation can be much higher, a challenge known as *super-resolution*. There is a long history of pursuing superresolution algorithms in the signal processing community [1]–[3]. The oldest one probably dates back to de Prony’s root-finding method in as early as 1795 [4], and variants of this method that are better suited for noisy data have also been proposed over time (see, e.g., [5]). Subspace methods based on the computation of eigenvector or singular vector decompositions, such as MUSIC [6], ESPRIT [7], and matrix pencil [8], have been another class of popular approaches since their inception in the 1980s. Different forms of maximum likelihood estimators have also been studied extensively [9], [10]. Collectively, these algorithms have superresolution capabilities, that is, they can resolve the parameters of interest at a resolution below the RL when the noise level is sufficiently small.

While a plethora of traditional methods already exists, convex optimization has recently emerged as a compelling

framework for performing superresolution, garnering significant attention from multiple communities spanning signal processing, applied mathematics, and optimization. Due to (the relative) tractability of convex analysis and convex optimization, the new framework offers several benefits. First, strong theoretical guarantees are rigorously established to back its performance up, even in the presence of noise and corruptions. Second, it is versatile enough to include prior knowledge into the convex program to handle a wide range of measurement models that are out of the reach of traditional methods. Third, leveraging the rapid progress in large-scale convex optimization opens up the possibility of applying efficient solvers tailored to real-world applications.

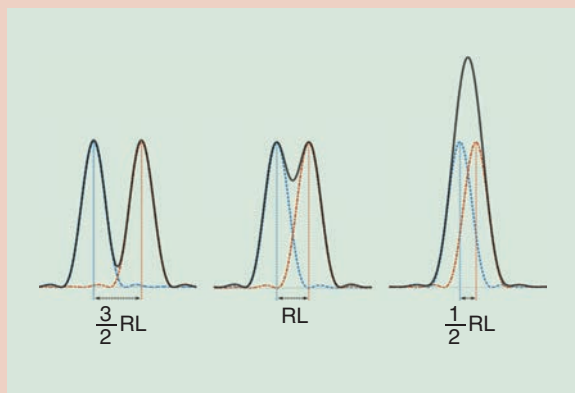
The goal of this article is to offer a friendly exposition to atomic norm minimization (ANM) [11] as a canonical convex approach for superresolution. The atomic norm is first proposed in [12] as a general framework for designing tight convex relaxations to promote simple signal decompositions, where one seeks to use a minimal number of atoms to represent a given signal from an atomic set composed of an ensemble of signal atoms. Celebrated convex relaxations such as the  $\ell_1$  norm approach for cardinality minimization [13] and the nuclear norm approach for rank minimization [14], can be viewed as particular instances of atomic norms for appropriately defined atomic sets. Specializing the atomic set to a dictionary containing all translations of the template signal over the continuous-valued parameter space, estimating the underlying translation parameters is then equivalent to identifying a sparse decomposition in an infinite-dimensional dictionary.

This key observation allows one to recast superresolution as solving an infinite-dimensional convex program [15], a special form of ANM considered in this article. We first highlight its mathematical formulation through a pedagogical yet useful model of superresolution that amounts to line spectrum estimation, where this infinite-dimensional convex program can be equivalently reformulated as a semidefinite program (SDP). We then demonstrate its versatility by discussing how it can be adapted to address measurement models that traditional methods may not apply easily. Finally, we illustrate its utility in superresolution image reconstruction for single-molecule fluorescence microscopy [16], where the infinite-dimensional convex program can be solved efficiently via tailored solvers.

Throughout this article, we use boldface letters to represent matrices and vectors, e.g.,  $\mathbf{a}$  and  $\mathbf{A}$ . We use  $\mathbf{A}^\top$ ,  $\mathbf{A}^H$ ,  $\text{Tr}(\mathbf{A})$  to represent the transpose, Hermitian transpose, and trace of  $\mathbf{A}$ , respectively. The conjugate of a complex scalar  $a$  is denoted as  $a^*$ . We use  $\mathbf{A} \succeq 0$  to represent  $\mathbf{A}$  is positive semidefinite. The matrix  $\text{toep}(\mathbf{u})$  denotes the Hermitian Toeplitz matrix whose first column is equal to  $\mathbf{u}$ , and  $\text{diag}(\mathbf{g})$  denotes the diagonal matrix with diagonal entries given as  $\mathbf{g}$ . The inner product between two matrices  $\mathbf{X}$  and  $\mathbf{P}$  is defined as  $\langle \mathbf{X}, \mathbf{P} \rangle = \text{Tr}(\mathbf{X}^H \mathbf{P})$ . Additionally, the notation  $f(n) = O(g(n))$  means that there exists a constant  $c > 0$  such that  $|f(n)| \leq c|g(n)|$ .

## What Is the Rayleigh Limit?

The Rayleigh limit (RL) is an empirical criterion characterizing the resolution of an optical system due to diffraction. In a conventional fluorescence microscope, for example, the observed diffraction patterns of two fluorescent point sources become visually more difficult to distinguish as the point sources get closer to each other, as illustrated in Figure S1. They are no longer resolvable when their separation is below the RL.



**FIGURE S1.** The combined response for two translated point spread functions under different separations of the point sources. The RL is an indication of the separability of the two sources.



## What is the atomic norm?

An everlasting idea in signal processing is decomposing a signal into a linear combination of judiciously chosen basis vectors and seeking compact and interpretable signal representations that are useful for downstream processing. For example, decomposing time series into sinusoids, speeches and images into wavelets, total system responses into impulse responses, and so on.

To fix ideas, consider the task of representing a signal  $\mathbf{x}$  in a vector space using atoms from a collection of vectors in  $\mathcal{A} = \{\mathbf{a}_i\}$  called an *atomic set*. The set  $\mathcal{A}$  can contain either a finite or infinite number of atoms. We wish to expand  $\mathbf{x}$  using the atoms in a form of

$$\mathbf{x} = \sum_i c_i \mathbf{a}_i, \quad \mathbf{a}_i \in \mathcal{A}, \quad (1)$$


where  $c_i > 0$  specifies the coefficients of the decomposition. In many applications, the size of  $\mathcal{A}$  can be much larger than the dimension of the signal, leading to an overcomplete representation, and there are an infinite number of possibilities to decompose  $\mathbf{x}$ . Which representation, then, shall we pick? Among the many plausible criteria, one meaningful approach is to pursue the Occam's razor principle and seek a parsimonious decomposition of the signal  $\mathbf{x}$  involving the smallest possible number of atoms in  $\mathcal{A}$ , i.e., the sparsest solution to (1). The corresponding representation is known as a *sparse representation* [17]. Many real-world signals admit sparse representations for appropriately chosen atomic sets. As a simple example, natural images are approximately sparse by selecting  $\mathcal{A}$  as a wavelet frame. Low-rank matrices, another class of signals that have enjoyed wide success in signal processing [18], are sparse with respect to an atomic set  $\mathcal{A}$  that is the collection of all unit-norm rank-one matrices.

Given a signal  $\mathbf{x}$ , how do we find its sparse representation in the atomic set  $\mathcal{A}$ ? In general, this problem is nonconvex and can be NP-hard due to the combinatorial aspect of cardinality minimization. The key motivation behind ANM, proposed by Chandrasekaran et al. [12], is to relax the nonconvex sparsity cost by its tight convex surrogate and instead solve the resulting convex relaxation, which is more tractable. This idea is a generalization of the popular  $\ell_1$  minimization for sparse vector recovery [19], [20] when  $\mathcal{A}$  is a finite set. Therein, one seeks to solve a linear program that minimizes the sum instead of the cardinality of the nonzero coefficients.

To extend the same idea to the case where  $\mathcal{A}$  is an arbitrary and possibly infinite-dimensional set, we first take the convex hull of  $\mathcal{A}$ , denoted as  $\text{conv}(\mathcal{A})$ , and then define its associated Minkowski functional (or gauge function) as [12]

$$\|\mathbf{x}\|_{\mathcal{A}} \triangleq \inf\{t \geq 0 : \mathbf{x} \in t \cdot \text{conv}(\mathcal{A})\}, \quad (2)$$

which is the solution to a convex program. When  $\mathcal{A}$  is centrally symmetric about the origin, definition (2) leads to a valid norm and is called the *atomic norm* of  $\mathbf{x}$ . Figure 1 illustrates



**FIGURE 1.** An atomic set  $\mathcal{A}$  (in red) and its convex hull  $\text{conv}(\mathcal{A})$  (in orange). The atomic norm of a vector  $\mathbf{x}$  can be interpreted as the smallest dilation factor  $t \geq 0$  such that  $\mathbf{x}$  belongs to  $t\text{conv}(\mathcal{A})$  (in blue).

this concept, where the atomic norm is the smallest nonnegative scaling of  $\text{conv}(\mathcal{A})$  until it intersects  $\mathbf{x}$ . Following definition (2), a fundamental geometric property is that the atomic norm ball, i.e.,  $\{\mathbf{x} : \|\mathbf{x}\|_{\mathcal{A}} \leq 1\}$ , is exactly  $\text{conv}(\mathcal{A})$ .

More interestingly, consider the case when  $\mathbf{x}$  lies in an  $n$ -dimensional vector space. Carathéodory's theorem [21] guarantees that any point in  $\text{conv}(\mathcal{A})$  can be decomposed as a convex combination of at most  $n + 1$  points in  $\mathcal{A}$ , where  $\mathcal{A}$  is not necessarily convex. Therefore, one may rewrite (2) as


$$\|\mathbf{x}\|_{\mathcal{A}} = \inf \left\{ \sum_i c_i : \mathbf{x} = \sum_i c_i \mathbf{a}_i, \quad c_i > 0, \quad \mathbf{a}_i \in \mathcal{A} \right\}, \quad (3)$$

as long as the centroid of  $\text{conv}(\mathcal{A})$  is the origin. The decomposition  $\sum_i c_i \mathbf{a}_i$  that obtains the infimum is referred to as the *atomic decomposition* of  $\mathbf{x}$  onto  $\mathcal{A}$ . It is not hard to see that the atomic norm indeed subsumes the  $\ell_1$  norm as a special case but accommodates the more general case where  $\mathcal{A}$  can be an infinite-dimensional set.

Several central questions regard how to properly select the atomic set, compute the atomic norm, and find the atomic decomposition, and when does the atomic decomposition coincide with the sparse representation, i.e., the convex relaxation is tight. Clearly, the answers highly depend on the atomic set as well as the signal itself. These questions have been addressed extensively in the study of  $\ell_1$ -norm minimization for sparse vector recovery [19], [20], [22]. In the context of superresolution, we first address these questions under a simple model that amounts to the classical problem of line spectrum estimation, which has deep connections to systems and control theory.

## A mathematical model of superresolution, equivalent to the line spectrum estimation

We first focus on a simple yet widely applicable model of superresolution that describes the convolution of a sequence of point sources with a PSF that is resolution-limited and is illustrated in Figure 2. Let  $x(t)$  be a spike signal given as



**FIGURE 2.** The mathematical model of superresolution. The spike signal  $x(t)$  is convolved with a point spread function  $g(t)$ , leading to the degradation of its resolution, which is further exacerbated by an additive noise  $\epsilon(t)$ , producing an output signal  $y(t)$ .

$$x(t) = \sum_{k=1}^r c_k \delta(t - \tau_k). \quad (4)$$

Here,  $r$  is the number of spikes,  $c_k \in \mathbb{C}$  and  $\tau_k \in [0, 1)$  are the complex amplitude and delay of the  $k$ th spike. Without loss of generality, the maximal delay is normalized to 1. Such a spike signal can model many physical phenomena, such as firing times of neurons, locations of fluorescence molecules, and so on. Let  $g(t)$  be the PSF whose bandwidth is limited due to the RL. Its Fourier transform  $G(f)$  satisfies  $G(f) = 0$  whenever  $|f| > B/2$  for some bandwidth  $B > 0$ . Its convolution with  $x(t)$ , contaminated by an additive noise  $\epsilon(t)$ , can be written as

$$y(t) = x(t) * g(t) + \epsilon(t) = \sum_{k=1}^r c_k g(t - \tau_k) + \epsilon(t),$$

where  $*$  denotes the convolution operator. Sampling the Fourier transform of the above equation at the frequencies  $\ell = -\lfloor B/2 \rfloor, \dots, 0, \dots, \lfloor B/2 \rfloor$ , we obtain the measurements

$$Y_\ell = G_\ell \cdot X_\ell + E_\ell = G_\ell \cdot \left( \sum_{k=1}^r c_k e^{-j2\pi\ell\tau_k} \right) + E_\ell, \quad (5)$$

where  $G_\ell$ ,  $X_\ell$ ,  $E_\ell$ , and  $Y_\ell$  are the Fourier transforms of  $g(t)$ ,  $x(t)$ ,  $\epsilon(t)$ , and  $y(t)$  evaluated at frequency  $\ell$ , respectively. The total number of samples is  $n = 2\lfloor B/2 \rfloor + 1 \approx B$ . We write (5) in a vector form as

$$\mathbf{y} = \text{diag}(\mathbf{g})\mathbf{x} + \boldsymbol{\epsilon}, \quad (6)$$

where  $\mathbf{y} = [Y_\ell]$ ,  $\mathbf{g} = [G_\ell]$ ,  $\mathbf{x} = [X_\ell]$ , and  $\boldsymbol{\epsilon} = [E_\ell]$ . The problem of superresolution is then to estimate  $\{c_k, \tau_k\}_{1 \leq k \leq r}$  accurately from  $\mathbf{y}$ , without knowing the model order  $r$  a priori. Here, the RL is inversely proportional to the bandwidth  $B$  and, roughly speaking, is about  $1/n$ .

When the PSF  $g(t)$  is known, one can equalize (5) by multiplying  $G_\ell^{-1}$  to both sides, provided that  $G_\ell$ 's are nonzero. The observation  $\mathbf{z} = [G_\ell^{-1} Y_\ell]$  relates to  $\mathbf{x}$  as

$$\mathbf{z} = \mathbf{x} + \tilde{\boldsymbol{\epsilon}}, \quad (7)$$

where  $\tilde{\boldsymbol{\epsilon}}$  is the additive noise. With a slight abuse of notation, we map the index of  $\ell$  from  $\lfloor B/2 \rfloor, \dots, -\lfloor B/2 \rfloor$  to  $0, \dots, n-1$  for convenience, and write  $\mathbf{x}$  as a superposition of complex sinusoids as

$$\mathbf{x} = \sum_{k=1}^r c_k \mathbf{a}(\tau_k), \quad (8)$$

where  $\mathbf{a}(\tau) \in \mathbb{C}^n$  is a vector defined as

$$\mathbf{a}(\tau) = [1, e^{j2\pi\tau}, \dots, e^{j2\pi(n-1)\tau}]^T, \quad \tau \in [0, 1). \quad (9)$$

Notably, the above simplified model (7) also amounts to the classical problem of line spectrum estimation, which consists of estimating a mixture of sinusoids (with frequencies  $\tau_k \in [0, 1)$ ) from equispaced time samples (sampled at integers  $\{0, \dots, n-1\}$ ) of the time-domain signal  $x_{\text{ls}}(t) = \sum_{k=1}^r c_k e^{j2\pi\tau_k t}$ . This finds applications in speech processing, power system monitoring, systems identification, and so on. The same model also describes direction-of-arrivals estimation using a uniform linear array, which is studied extensively in the literature of spectrum analysis [1].

### Line spectrum superresolution via ANM


In the absence of noise, one could think of superresolution as estimating the continuous-time spike signal  $x(t)$  in (4) from its discrete-time moment measurements  $\mathbf{x}$  in (8), which are related through

$$\mathbf{x} = \int_0^1 \mathbf{a}(\tau) d\mathbf{x}(\tau). \quad (10)$$

One can also think of  $x(t)$  as the representation of  $\mathbf{x}$  over a continuous dictionary

$$\mathcal{A}_0 = \{\mathbf{a}(\tau) : \tau \in [0, 1)\}, \quad (11)$$





**FIGURE 3.** A visualization of the continuous-valued atomic set for line spectrum superresolution. (a) The moment curve  $\mathcal{A}_0$  restricted to three real moments ( $\{\cos(2\pi\tau), \cos(4\pi\tau), \cos(6\pi\tau)\}^T: \tau \in [0, 1]\}$ ) and (b) its convex hull. (c) The phased version of the moment curve  $\mathcal{A}_{1D}$  restricted to three real moments ( $\{\cos(2\pi\tau + \phi), \cos(4\pi\tau + \phi), \cos(6\pi\tau + \phi)\}^T: \tau \in [0, 1], \phi \in [0, 2\pi]\}$ ) and (d) its convex hull.

which forms a one-dimensional variety of  $\mathbb{C}^n$  called the *moment curve*, illustrated in Figure 3(a). It is well known that the convex hull of  $\mathcal{A}_0$ , illustrated in Figure 3(b), is a body of  $\mathbb{C}^n$  that can be parameterized by a set of linear matrix inequalities [23] and has close relationships with the positivity of Hermitian Toeplitz matrices. This fundamental property of the moment curve has many implications in control and signal processing [24], [25] and is key for developing a superresolution theory based on ANM.

It is clearly possible to obtain the same  $\mathbf{x}$  from different  $x(t)$ . However, if we impose some sparsity assumption, namely constraining how many spikes are allowed in  $x(t)$ , this representation can be ensured to be unique. In particular, the representation (8) is unique as long as  $r \leq \lfloor n/2 \rfloor$  and the support set  $\mathcal{T} = \{\tau_k\}_{1 \leq k \leq r}$  contains distinct elements.

### Atomic norm for line spectrum superresolution

To apply the framework of ANM for superresolution, one must first define the atomic set properly. Since the complex amplitudes  $c_k$ 's can take arbitrary phases, we introduce the following augmented atomic set taking this into account:

$$\mathcal{A}_{1D} = \{e^{j\phi} \mathbf{a}(\tau) : \tau \in [0, 1], \phi \in [0, 2\pi]\}. \quad (12)$$

See an illustration of  $\mathcal{A}_{1D}$  and its convex hull in Figure 3(c) and (d). Writing  $c_k = |c_k| e^{j\phi_k}$ ,  $\mathbf{x}$  can be represented as a positive combination of the atoms in  $\mathcal{A}_{1D}$  as  $\mathbf{x} = \sum_{k=1}^r |c_k| e^{j\phi_k} \mathbf{a}(\tau_k)$ . It is easy to verify that  $\mathcal{A}_{1D}$  is centrally symmetric around the origin and consequently it induces an atomic norm over  $\mathbb{C}^n$ , as defined in (2) and (3). It is worth noting that minimizing the atomic norm of  $\mathbf{x}$  is equivalent to minimizing the total variation of  $x(t)$ , i.e.,

$$\min \|x(t)\|_{TV} \quad \text{s.t.} \quad \mathbf{x} = \int_0^1 \mathbf{a}(t) dx(t), \quad (13)$$

and both viewpoints are used frequently in the literature.

Remarkably, this atomic norm admits an equivalent SDP characterization, thanks to the following Carathéodory–Fejér–Pisarenko decomposition [26]:


$$\|\mathbf{x}\|_{\mathcal{A}} = \inf_{\substack{\mathbf{u} \in \mathbb{C}^n \\ t > 0}} \left\{ \frac{1}{2n} \text{Tr}(\text{toep}(\mathbf{u})) + \frac{1}{2} t : \begin{bmatrix} \text{toep}(\mathbf{u}) & \mathbf{x} \\ \mathbf{x}^H & t \end{bmatrix} \succeq 0 \right\}. \quad (14)$$

## From Bounded Polynomials to Linear Matrix Inequalities

Many applications encountered in signal processing, systems and control theory involve comparing the magnitudes of two real trigonometric polynomials,  $R(\tau) = \text{Re} \langle \mathbf{a}(\tau), \mathbf{r} \rangle$  and  $S(\tau) = \text{Re} \langle \mathbf{a}(\tau), \mathbf{s} \rangle$ , e.g., bounding the frequency response of a finite impulse response filter by a desired shape. Although such inequalities, in appearance, require verification over a continuous set of parameters, they can easily be translated into linear matrix inequalities of finite dimension, which are amenable to optimization.

Central to the equivalence is a Gram parametrization of real trigonometric polynomials [24], [S1], [S2], by noticing that every real trigonometric polynomial  $R(\tau) = \text{Re} \langle \mathbf{a}(\tau), \mathbf{r} \rangle$  can be equivalently represented as a quadratic form  $R(\tau) = \mathbf{a}(\tau)^H \mathbf{G} \mathbf{a}(\tau)$  for a family of Hermitian matrices  $\mathbf{G} \in \mathcal{G}(R)$ , where  $\mathbf{G}$  is related to  $\mathbf{r}$  through the following Gram mapping (Figure S2):

$$\mathbf{G} \in \mathcal{G}(R) \iff \text{Tr}(\mathbf{G}) = \text{Re}(r_0), \sum_{i=1}^{n-k} G_{i,i+k} = \frac{r_k}{2}, k = 1, \dots, n-1. \quad (\text{S1})$$



**FIGURE S2.** The Gram mapping for two polynomials satisfying  $R(\tau) \leq S(\tau)$  over  $[0, 1]$ , and the corresponding  $\mathbf{G}, \mathbf{H}$  in the respective Gram sets such that  $\mathbf{G} \leq \mathbf{H}$ .

Contrary to its abstract definition in (2), the reformulation (14) provides a tractable approach to compute the quantity  $\|\mathbf{x}\|_{\mathcal{A}}$ , which can be accomplished using generic off-the-shelf convex solvers [27]. The Vandermonde decomposition of  $\text{toep}(\mathbf{u})$ , i.e.,  $\text{toep}(\mathbf{u}) = \sum_{k=1}^{r'} |c'_k| \mathbf{a}(\tau'_k) \mathbf{a}(\tau'_k)^H$  can then be used to identify the support  $\hat{\mathcal{T}} = \{\tau'_k\}$  of the atomic representation of  $\mathbf{x}$ , as well as the atomic norm  $\|\mathbf{x}\|_{\mathcal{A}} = \sum_{k=1}^{r'} |c'_k|$ .

### Duality and atomic decomposition

The Lagrangian duality theory marks an important aspect in understanding the atomic norm. The Lagrange dual problem associated with the ANM (2) reads [11]

$$\max_{\mathbf{p}} \text{Re} \langle \mathbf{x}, \mathbf{p} \rangle \quad \text{subject to} \quad \|\mathbf{p}\|_{\mathcal{A}}^* \leq 1, \quad (15)$$

where the dual atomic norm  $\|\mathbf{p}\|_{\mathcal{A}}^*$  of a vector  $\mathbf{p} \in \mathbb{C}^n$  is defined with respect to the atomic set  $\mathcal{A}_{1D}$  as

$$\|\mathbf{p}\|_{\mathcal{A}}^* \triangleq \sup_{\mathbf{a} \in \mathcal{A}_{1D}} \text{Re} \langle \mathbf{a}, \mathbf{p} \rangle = \sup_{\tau \in [0, 1]} \underbrace{|\langle \mathbf{a}(\tau), \mathbf{p} \rangle|}_{P(\tau)}. \quad (16)$$

The last equality of (16) suggests that the dual atomic norm can be interpreted as the supremum of the modulus of a complex trigonometric polynomial  $P(\tau) = \langle \mathbf{a}(\tau), \mathbf{p} \rangle = \sum_{\ell=0}^{n-1} p_{\ell} e^{-j2\pi\ell\tau}$  with coefficients given by the vector  $\mathbf{p}$ . Constraints of this type

A remarkable property (e.g., [25, Lemma 4.23]) is that the Gram mapping preserves the partial order between the polynomials and the Hermitian matrices. Let  $\mathbf{G} \in \mathcal{G}(R)$ , then  $R(\tau) \leq S(\tau)$  holds for every  $\tau \in [0, 1]$  if and only if there exists  $\mathbf{H} \in \mathcal{G}(S)$  such that  $\mathbf{G} \leq \mathbf{H}$ .

As an example, consider the dual norm constraint  $\|\mathbf{p}\|_{\mathcal{A}}^* \leq 1$  in (16), which amounts to upper bounding  $R(\tau) = |\langle \mathbf{a}(\tau), \mathbf{p} \rangle|^2$  by  $S(\tau) = 1$ . Since  $R(\tau) = \mathbf{a}(\tau)^H \mathbf{p} \mathbf{p}^H \mathbf{a}(\tau)$ , it is clear that  $\mathbf{p} \mathbf{p}^H \in \mathcal{G}(R)$ . The constraint holds if and only if there exists a matrix  $\mathbf{H} \in \mathcal{G}(S)$  satisfying  $\mathbf{H} \geq \mathbf{p} \mathbf{p}^H$ . Rewriting this condition using the Schur's complement, as well as expanding the Gram mapping of  $S(\tau) = 1$ , we can obtain the semidefinite constraint in (17), a consequence also known as the *bounded real lemma*.

### References

- [S1] O. Teke and P. P. Vaidyanathan, "On the role of the bounded lemma in the SDP formulation of atomic norm problems," *IEEE Signal Process. Lett.*, vol. 24, no. 7, pp. 972–976, 2017. doi: 10.1109/LSP.2017.2700442.
- [S2] M.-D. Choi, T. Y. Lam, and B. Reznick, "Sums of squares of real polynomials," *Proc. Symp. Pure Math.*, vol. 58, no. 2, pp. 103–126, 1995. doi: 10.1090/pspum/058.2/1327293.

are known to be equivalent to linear matrix inequalities involving the positivity of some Hermitian matrices (see “From Bounded Polynomials to Linear Matrix Inequalities”). The dual program (15) can be reformulated into an SDP as

$$\begin{aligned} & \max_{\mathbf{p} \in \mathbb{C}^n, \mathbf{H} \in \mathbb{C}^{n \times n}} \quad \text{Re} \langle \mathbf{x}, \mathbf{p} \rangle \\ & \text{subject to} \quad \begin{bmatrix} \mathbf{H} & \mathbf{p} \\ \mathbf{p}^H & 1 \end{bmatrix} \succeq 0 \\ & \quad \sum_{i=1}^{n-k} H_{i,i+k} = \delta_k, k = 0, \dots, n-1, \end{aligned} \quad (17)$$

where  $H_{i,j}$  is the  $(i,j)$ th entry of the matrix  $\mathbf{H}$  and the indicator function  $\delta_k$  equals 1 if  $k = 0$  and 0 otherwise.

Another merit of the dual formulation is that the support set of the atomic decomposition can be inferred from the optimal solution  $\hat{\mathbf{p}}$  to the dual problem (15), by examining the dual polynomial  $\hat{P}(\tau) = \langle \mathbf{a}(\tau), \hat{\mathbf{p}} \rangle$ . We identify the spikes as the locations of the extreme values of the modulus of  $\hat{P}(\tau)$ :

$$\hat{\mathcal{T}} = \{\tau : |\hat{P}(\tau)| = 1\}. \quad (18)$$

This is possible, because, under strong duality, both the primal and the dual problems must share the same optimal objective value, i.e.,  $\|\mathbf{x}\|_{\mathcal{A}} = \sum_{k=1}^{r'} |c'_k|$ , where  $\mathbf{x} = \sum_{k=1}^{r'} c'_k \mathbf{a}(\tau_k)$  is the atomic decomposition of  $\mathbf{x}$ . Consequently, the optimal value of the dual program becomes

$$\text{Re} \langle \mathbf{x}, \hat{\mathbf{p}} \rangle = \text{Re} \left\langle \sum_{k=1}^{r'} c'_k \mathbf{a}(\tau_k), \hat{\mathbf{p}} \right\rangle = \text{Re} \sum_{k=1}^{r'} c'_k \hat{P}(\tau_k),$$


indicating  $\hat{P}(\tau_k) = \text{sgn}(c'_k) = c'_k / |c'_k|$  whenever the atomic decomposition is nonvanishing at  $\tau_k$ . This approach is illustrated in Figure 4 for a length-33 signal with six spikes, where the peaks of  $\hat{P}(\tau)$  match the locations of the true spikes, indicating the atomic decomposition perfectly recovers the true sparse representation.

The atomic norm offers an approach for line spectrum superresolution that is drastically different from traditional methods, which rely heavily on the correctness of model order estimation. The dual polynomial approach, in contrast, does not require any prior knowledge of the model order and can estimate the spikes with an infinitesimal precision.

### Exact recovery guarantees

So far, we have explained the algorithmic approach of ANM for line spectrum superresolution. A central question regards understanding whether this convex relaxation is tight or not. More precisely, one would like to identify the conditions under which the estimated support  $\hat{\mathcal{T}}$  coincides with the true support  $\mathcal{T}$  of the signal  $\mathbf{x}$  and, correspondingly, the atomic decomposition  $\mathbf{x} = \sum_{k=1}^{r'} c'_k \mathbf{a}(\tau_k)$  coincides with the sparsest representation  $\mathbf{x} = \sum_{k=1}^r c_k \mathbf{a}(\tau_k)$  over the atomic set  $\mathcal{A}_{1D}$ .

Such questions were extensively addressed in the context of  $\ell_1$  norm minimization, where the atomic set  $\mathcal{A}$  has a



**FIGURE 4.** A spike localization via pinpointing the peaks of the dual polynomial  $|\hat{P}(\tau)|$  (in blue) associated with the optimal solution  $\hat{\mathbf{p}}$  of the dual program (15) for a signal  $\mathbf{x}$  of length  $n = 33$  with six true spikes (in black).

finite number of elements. The performance guarantees often depend on specific structural properties of  $\mathcal{A}$ , formalized in the notion of restricted isometry property (RIP) [28] or certain incoherence properties [29]. Unfortunately, these properties do not hold when considering a continuous dictionary such as  $\mathcal{A}_{1D}$ , since two atoms  $\mathbf{a}(\tau)$  and  $\mathbf{a}(\tau + \delta)$  can be more and more correlated with each other as their separation  $\delta$  tends to zero, leading to arbitrarily small RIP or coherence constants.

Nonetheless, one could question for which class of signals is the relaxation tight. Leveraging duality theory, the atomic norm approach is tight for a fixed signal  $\mathbf{x}$ , i.e.,  $\hat{\mathcal{T}} = \mathcal{T}$ , as long as there exists a dual certificate  $\mathbf{p}_*$ , such that  $P_*(\tau) = \langle \mathbf{a}(\tau), \mathbf{p}_* \rangle$  satisfies [11]

$$P_*(\tau_k) = \text{sgn}(c_k), \quad \forall \tau_k \in \mathcal{T}, \quad (19a)$$

$$|P_*(\tau)| < 1, \quad \forall \tau \notin \mathcal{T}. \quad (19b)$$

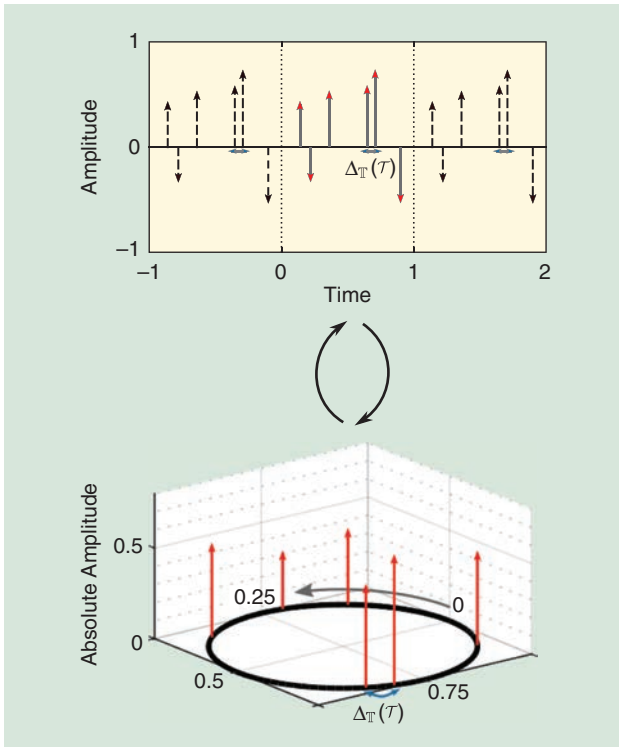
In other words, it amounts to finding an  $(n-1)$ -order trigonometric polynomial that interpolates sign patterns of the spike signal at the spike locations, as well as is bounded in magnitude by 1.

Intuitively, the difficulty of interpolation depends on the separations between the spikes in  $\mathcal{T}$  and, more precisely on the minimal separation, or the minimal wrap-around distance between any pair of distinct spikes in  $\mathcal{T}$ , defined formally as

$$\Delta_{\mathcal{T}}(\mathcal{T}) \triangleq \inf_{\tau, \tau' \in \mathcal{T}, \tau \neq \tau'} \min_{q \in \mathbb{Z}} |\tau - \tau' + q|. \quad (20)$$

This metric is illustrated in Figure 5 and reflects the periodic behavior of the atom  $\mathbf{a}(\tau + q) = \mathbf{a}(\tau)$  for any integer  $q \in \mathbb{Z}$ . For instance, if  $\mathcal{T} = \{0.1, 0.9\}$ , then  $\Delta_{\mathcal{T}}(\mathcal{T}) = 0.2$ .

A remarkable result, established by Candès and Fernandez-Granda in [11], is that for sufficiently large  $n$ , a valid certificate can be constructed in a deterministic fashion, as long as the separation condition  $\Delta_{\mathcal{T}}(\mathcal{T}) \geq 4/(n-1)$  holds, regardless of the complex amplitudes of the spikes. Furthermore, this result does not make any randomness assumptions on the signal.



**FIGURE 5.** A representation of the minimal wrap-around distance  $\Delta_T(\mathcal{T})$  for a given set of spikes  $\mathcal{T}$ . The distance corresponds to the minimal gap between any element  $\tau \in \mathcal{T}$  and any distinct elements in the aliased set  $\mathcal{T} + \mathbb{Z}$ .

Later, this separation condition was further improved by Fernandez-Granda [30] to

$$\Delta_T(\mathcal{T}) > \frac{2.52}{n-1}.$$

Conversely, there exist some spike signals with  $\Delta_T(\mathcal{T}) < 2/(n-1)$  such that ANM fails to resolve [31].

### Atomic norm denoising

In practice, the observations are corrupted by noise and no estimator can exactly recover the spike signal  $x(t)$ . This raises a natural question regarding the robustness of the estimate produced by ANM methods. When the noise is additive and the observation  $\mathbf{z}$  obeys the noisy model (7), it has been proposed to estimate  $\mathbf{x}$  by searching around the observation  $\mathbf{z}$  for signals with small atomic norms [32]:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 + \lambda \|\mathbf{x}\|_{\mathcal{A}}, \quad (21)$$

where  $\lambda > 0$  is a regularization parameter that draws a tradeoff between the fidelity to the observation and the size of the atomic norm. This method, known as *atomic norm denoising*, can be interpreted as a generalization of the celebrated LASSO estimator [33].

When the noise vector  $\tilde{\epsilon}$  is composed of independent identically distributed (i.i.d.) complex Gaussian entries with zero

mean and variance  $\sigma^2$ , the mean square error (MSE) of the estimate  $\hat{\mathbf{x}}$  returned by (21) can be bounded as [32]

$$\frac{1}{n} \|\hat{\mathbf{x}} - \mathbf{x}\|_2^2 = O\left(\sigma \sqrt{\frac{\log n}{n}} \sum_{k=1}^r |c_k|\right)$$

with a high probability by setting  $\lambda = \eta \sigma \sqrt{n \log n}$  for some constant  $\eta \in (1, \infty)$ , e.g.,  $\eta = 1.2$  in practice. This error rate can be significantly improved when the spikes satisfy the separation condition  $\Delta_T(\mathcal{T}) \geq 4/(n-1)$  where with high probability one has [34]

$$\frac{1}{n} \|\hat{\mathbf{x}} - \mathbf{x}\|_2^2 = O\left(\sigma^2 \frac{r \log n}{n}\right).$$

This last error rate is near-optimal up to some logarithmic factor, since no estimator can achieve an MSE below the rate  $O(\sigma^2 r \log(n/r)/n)$  [34].

A more important performance criteria in superresolution concerns the stability of the support estimate  $\hat{\mathcal{T}}$ , which has been studied in [34]–[38]. When the spikes satisfy the separation condition  $\Delta_T(\mathcal{T}) > 5.0018/(n-1)$ , and the complex amplitudes of the coefficients  $\{c_k\}_{1 \leq k \leq r}$  have approximately the same modulus, then it is established in [36] that the atomic decomposition of the output  $\hat{\mathbf{x}}$  of (21) is composed of the same number of spikes, i.e.,  $|\hat{\mathcal{T}}| = |\mathcal{T}| = r$  and that the estimated parameters satisfy

$$|c_k| |\hat{\tau}_k - \tau_k| = O\left(\sigma \sqrt{\frac{\log n}{n^{3/2}}}\right), |c_k - \hat{c}_k| = O\left(\sigma \sqrt{\frac{\log n}{n}}\right)$$

with high probability. Altogether, it can be seen that atomic norm denoising achieves near-optimal performance guarantees as long as the spikes are separated by a few times the RL (see “Is the Separation Condition Necessary?”).

It is natural to wonder how atomic norm denoising fares compared with classical approaches such as Prony and MUSIC for line spectrum estimation. We examine their ability to resolve close-located spikes with opposite signs, where the reconstruction performance is measured in terms of the MSE of the estimated spike locations  $\hat{\mathcal{T}}$ , and for different values of separation  $\Delta_T(\mathcal{T}) = \alpha/n$ . Figure 6 shows the MSE of atomic norm denoising, Prony’s method with Cadzow denoising [39], and root-MUSIC [40] with respect to the SNR defined as  $\|\mathbf{x}\|_2^2/(n\sigma^2)$ , benchmarked against the Cramér–Rao bound (CRB), when the separation parameter  $\alpha = 2, 1.75, 1.5$ , respectively. It is clear that atomic norm denoising outperforms classical approaches and approaches the CRB at a much lower SNR.

### A faster algorithm via ADMM

While the SDP formulation is tractable, its computational complexity is prohibitive when solving large-dimensional problems. Fortunately, it is possible to develop tailored algorithms that are significantly faster. For conciseness, we will discuss one approach based on the alternating direction method of multipliers (ADMM) [32]. The general




## Is the Separation Condition Necessary?

One might wonder if requiring a separation condition makes atomic norm minimization inferior, since many methods do not require such a separation in the noise-free case. However, some form of separation is unavoidable for stable recovery in noisy superresolution, no matter which method is used [S3]. In particular, [S3] shows that when  $\Delta_{\mathcal{T}}(\mathcal{T}) < 2/n$ , there exists a pair of spike signals  $x(t)$  and  $x'(t)$  with the same minimal separation, such that no estimator can distinguish them. Figure S3(a) exhibits such a pair of positive spike signals (see [S3] for its construction) with a minimal separation

$1.7/n$ , where their observations are very close. Figure S3(b) further demonstrates the distance between their observations as the signal dimension increases, for different separation parameter  $\alpha$ , where  $\Delta_{\mathcal{T}}(\mathcal{T}) = \alpha/n$ . It is clear that their observations are increasingly indistinguishable as the signal dimension tends to infinity when  $\Delta_{\mathcal{T}}(\mathcal{T}) < 2/n$ .

### Reference

[S3] A. Moitra, "Super-resolution, extremal functions and the condition number of Vandermonde matrices," in *Proc. 47th Annual ACM Symp. Theory of Computing*, 2015, pp. 821–830. doi: 10.1145/2746539.2746561.



**FIGURE S3.** (a) A pair of spike signal  $x(t)$  and  $x'(t)$  with a minimal separation  $\Delta_{\mathcal{T}}(\mathcal{T}) = 1.7/n$  (top), yet their observations produced according to (10) are almost indistinguishable (bottom). (b) The Euclidean distance between the observations of  $x(t)$ ,  $x'(t)$  with a minimal distance  $\Delta_{\mathcal{T}}(\mathcal{T}) = \alpha/n$  as a function of the signal length  $n$ , for different values of the separation parameter  $\alpha$ .

principle of ADMM is to split the quadratically augmented Lagrangian function of an optimization problem into a sum of separable subfunctions [41]. Each iteration of the algorithm consists of performing independent local minimization on each of those quantities, while ensuring that the feasibility constraints are always satisfied. The iterations run until both primal and dual residuals satisfy a pre-defined tolerance level.

We take atomic norm denoising (21) as an example, which, in light of (14), can be equivalently rewritten as

$$\begin{aligned} \min_{x,u,t} \quad & \frac{1}{2} \|x - z\|_2^2 + \frac{\lambda}{2} \left( \frac{1}{n} \text{Tr}(\text{toep}(u)) + t \right) \\ \text{subject to} \quad & S = \begin{bmatrix} \text{toep}(u) & x \\ x^H & t \end{bmatrix}, S \succeq 0. \end{aligned}$$

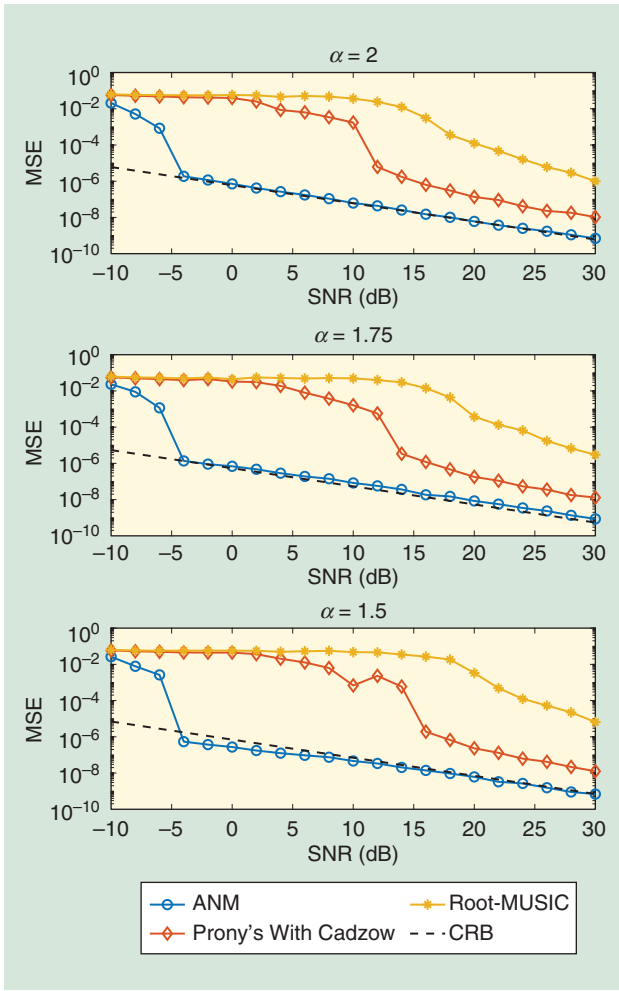
The above program has been “augmented” by introducing an intermediate variable  $S$  for the purpose of decoupling the positive semidefinite constraint on the matrix  $S$  from the linear constraints on its structure. The augmented Lagrangian  $\mathcal{L}$  is given as

$$\begin{aligned} \mathcal{L}(x, u, t, \Sigma, S) = & \frac{1}{2} \|x - z\|_2^2 + \frac{\lambda}{2} \left( \frac{1}{n} \text{Tr}(\text{toep}(u)) + t \right) \\ & + \left\langle \Sigma, S - \begin{bmatrix} \text{toep}(u) & x \\ x^H & t \end{bmatrix} \right\rangle \\ & + \frac{\rho}{2} \left\| S - \begin{bmatrix} \text{toep}(u) & x \\ x^H & t \end{bmatrix} \right\|_F^2, \end{aligned}$$

where  $S$  and  $\Sigma$  are  $(n+1)$ -dimensional Hermitian matrices and  $\rho > 0$  is a regularization parameter. The successive update steps to minimize the augmented Lagrangian are given in Algorithm 1. Closed-form solutions can be found for the first update step, yielding a very efficient implementation. The second update is the most costly part, as a projection over the cone of positive semidefinite Hermitian matrices has to be computed. This computation is typically achieved using power methods [42], with a computational complexity of  $O(n^3)$  per iteration.

### Can we discretize?

It may be worthwhile to pause and compare ANM to other approaches based on convex optimization for superresolution, particularly  $\ell_1$  minimization that is widely popular



**FIGURE 6.** The MSE of different methods for estimating two spikes with opposite signs separated by  $\Delta_T(\mathcal{T}) = \alpha/n$ , averaged more than 200 Monte Carlo trials and benchmarked against the CRB for different separation parameter  $\alpha = 2, 1.75, 1.5$ . Here, the signal length is  $n = 101$ . SNR: signal-to-noise ratio.

#### Algorithm 1. ADMM for atomic norm denoising [32].

**Input:** Observation  $\mathbf{z}$ ; parameters  $\lambda, \rho > 0$ ;  
Initialize  $j = 0$ , and  $\Sigma_0, \mathbf{S}_0$  to zero matrices  
**Repeat** until stopping criteria  
 $(\mathbf{x}_{j+1}, \mathbf{u}_{j+1}, t_{j+1}) \leftarrow \argmin_{\mathbf{x}, \mathbf{u}, t} \mathcal{L}(\mathbf{x}, \mathbf{u}, t, \Sigma_j, \mathbf{S}_j);$   
 $\mathbf{S}_{j+1} \leftarrow \argmin_{\mathbf{S} \geq 0} \mathcal{L}(\mathbf{x}_{j+1}, \mathbf{u}_{j+1}, t_{j+1}, \Sigma_j, \mathbf{S});$   
 $\Sigma_{j+1} \leftarrow \Sigma_j + \rho \left( \mathbf{S}_{j+1} - \begin{bmatrix} \text{toep}(\mathbf{u}_{j+1}) & \mathbf{x}_{j+1} \\ \mathbf{x}_{j+1}^H & t_{j+1} \end{bmatrix} \right);$   
 $j \leftarrow j + 1;$   
**Output:**  $\mathbf{x}_j$

for high-resolution imaging and localization in the recent literature due to compressed sensing (CS) [43], [44].

The  $\ell_1$  norm can be seen as a discrete approximation of the atomic norm. Indeed, taking the atomic set  $\mathcal{A}_{1D}$ , one can pick a desired resolution  $Q$  and discretize it as

$$\mathcal{A}_{1D, \text{discrete}} = \left\{ e^{j\phi} \mathbf{a}\left(\frac{q}{Q}\right) : q = 0, \dots, Q-1, \phi \in [0, 2\pi) \right\},$$

and then perform  $\ell_1$  minimization over  $\mathcal{A}_{1D, \text{discrete}}$ . The convex hull of  $\mathcal{A}_{1D, \text{discrete}}$  approaches that of  $\mathcal{A}_{1D}$  as the discretization gets finer, which suggests the performance of  $\ell_1$  minimization over the discretized dictionary approaches that of ANM asymptotically [45]. If the spike signal meets a so-called nondegenerate source condition [46, Definition 2], this approach will return a sparse solution supported on the elements of the discretized  $\mathcal{A}_{1D, \text{discrete}}$  surrounding the ground-truth spikes, when the noise is small enough [46], [47]. However, it remains unclear which class of spike signals satisfies the nondegenerate source condition in practice.

However, this discretization may come with several undesired consequences when the grid size  $Q$  is finite in practice. The theory of  $\ell_1$  minimization only provides exact recovery guarantees when the spikes of  $\mathbf{x}(t)$  lie on the grid, which is unrealistic. In fact, there is always an inevitable basis mismatch [48] between the spikes represented in the discretized dictionary  $\mathcal{A}_{1D, \text{discrete}}$  and the true spikes, no matter how fine the grid is. Perfect recovery is not possible in this situation, even in the absence of noise due to this mismatch. Furthermore, one can find signals whose representations in  $\mathcal{A}_{1D, \text{discrete}}$  are not compressible due to spectral leakage and therefore are poorly recovered using  $\ell_1$  minimization, e.g., the recovery may contain many spurious spikes. Therefore, cautions are needed to account for such consequences when applying discretization, and efforts to mitigate the basis mismatch have been proposed extensively, e.g., [49] and [50].

### Generalizations of atomic sets

The tool of atomic norms can be extended easily to handle a wide range of scenarios in a unified manner, by properly adjusting the atomic set for signal decompositions such as incorporating prior information and dealing with multidimensional settings and multiple measurement vectors, to illustrate a few.

#### Atomic set for positive spikes

In some applications, there exist additional information about the spikes, such as the coefficients of the spikes in (8) are positive, i.e.,  $c_k > 0$ . Examples include neural spike sorting, fluorescence microscopy imaging, or covariance-based spectrum estimation for noncoherent sources [1].

In this case, the atomic set reduces to the moment curve  $\mathcal{A}_0$  in (11). The induced  $\|\mathbf{x}\|_{\mathcal{A}}$  is no longer a norm, since  $\mathcal{A}_0$  is not centrally symmetric but, nonetheless, similar SDP characterization still holds. To be specific, the dual program now becomes

$$\max_{\mathbf{p} \in \mathbb{C}^n} \text{Re}\langle \mathbf{x}, \mathbf{p} \rangle \quad \text{subject to} \quad \sup_{\tau \in [0, 1)} \text{Re}\langle \mathbf{a}(\tau), \mathbf{p} \rangle \leq 1,$$

where the constraint bounds the real part of the trigonometric polynomial  $P(\tau) = \langle \mathbf{a}(\tau), \mathbf{p} \rangle$ . Using the Fejér-Riesz theorem (see, e.g., [25, Th. 1.1] and “From Bounded Polynomials to Linear Matrix Inequalities”), this can be equivalently represented as

$$\begin{aligned}
& \max_{p \in \mathbb{C}^n, H \geq 0} \operatorname{Re}\langle \mathbf{x}, \mathbf{p} \rangle \\
& \text{subject to } \operatorname{Tr}(\mathbf{H}) + \operatorname{Re}(p_0) = 1, \\
& \sum_{i=1}^{n-k} H_{i,i+k} + p_k/2 = 0, \quad k = 1, \dots, n-1.
\end{aligned}$$

It is long established [51], [52] that the spikes can be perfectly localized as long as  $r \leq \lfloor (n-1)/2 \rfloor$ , without requiring any separation between the spikes, as long as they are positive. The stability of this approach as well as the implications of nonnegative constraints for other atomic sets are further studied in [53]–[56]. As a comparison, Figure 7 shows the MSE of atomic norm denoising with and without positive constraints, Prony’s method with Cadzow denoising [39], and root-MUSIC [40] with respect to the SNR defined as  $\|\mathbf{x}\|_2^2/(n\sigma^2)$  for resolving two spikes with positive signs, separated by  $\Delta_{\mathbb{T}}(\mathcal{T}) = \alpha/n$  for  $\alpha = 1, 0.75, 0.5$ , respectively. It can be seen that atomic norm denoising still outperforms classical approaches and, in particular, incorporating the positive constraint leads to further improvements.

### Atomic set for multidimensional spikes

When the spikes reside in a multidimensional space, one can extend the 1D model in a straightforward manner. Here, we illustrate the setup for the 2D case, where each entry of the signal  $\mathbf{X}_{2D} \in \mathbb{C}^{n_1 \times n_2}$  can be expressed as a superposition of  $r$  complex sinusoids propagating in two directions:

$$\mathbf{X}_{2D} = \sum_{k=1}^r c_k \mathbf{a}_1(\tau_{1,k}) \mathbf{a}_2(\tau_{2,k})^\top, \quad (22)$$

where  $c_k$  and  $\boldsymbol{\tau}_k = [\tau_{1,k}, \tau_{2,k}]^\top \in [0, 1]^2$  are the complex amplitude and location of the  $k$ th spike, and  $\mathbf{a}_i(\tau)$  is given by (9) with the dimension parameter replaced by  $n_i$ ,  $i = 1, 2$ . It is natural to define the corresponding atomic set as [11], [57]

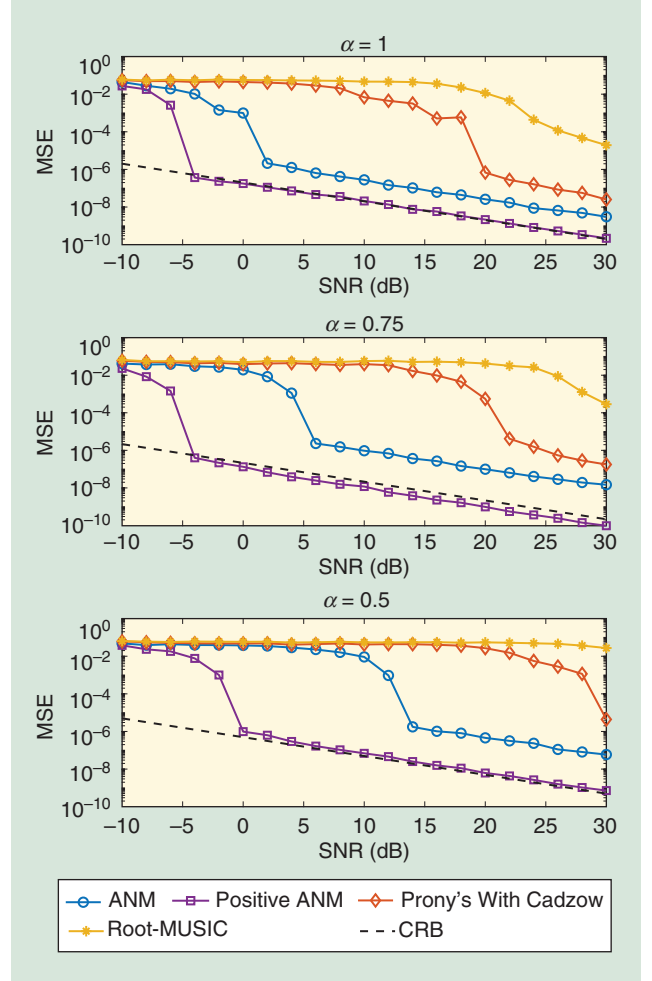
$$\mathcal{A}_{2D} = \{e^{j\phi} \mathbf{a}_1(\tau_1) \mathbf{a}_2(\tau_2)^\top : \boldsymbol{\tau} \in [0, 1]^2, \phi \in [0, 2\pi)\},$$

and the atomic norm according to (2). To localize the spikes, one could similarly study the associated dual problem


$$\max_{\mathbf{P} \in \mathbb{C}^{n_1 \times n_2}} \operatorname{Re}\langle \mathbf{X}, \mathbf{P} \rangle \quad \text{subject to } \|\mathbf{P}\|_{\mathcal{A}}^* \leq 1,$$

where the dual atomic norm can be reinterpreted as the supremum of a bivariate complex trigonometric polynomial  $P(\tau_1, \tau_2) = \langle \mathbf{a}_1(\tau_1) \mathbf{a}_2(\tau_2)^\top, \mathbf{P} \rangle$  with the matrix  $\mathbf{P}$  as its coefficients. Again, one can localize the spikes by examining the extremal points of the dual polynomial, which is illustrated in Figure 8. Cautions need to be taken when attempting to solve the dual program in two or higher dimensions, since the bounded real lemma [24], [25] does not hold anymore. Instead, a precise characterization requires solving a hierarchy of sum-of-squares relaxations and, fortunately in practice, the first level usually suffices [24], [57], [58].

The tightness of the ANM approach is closely related to a separation condition analogous to the 1D case [30]. Namely,



**FIGURE 7.** The MSE of different methods for estimating two spikes with positive signs separated by  $\Delta_{\mathbb{T}}(\mathcal{T}) = \alpha/n$ , averaged more than 200 Monte Carlo trials and benchmarked against the CRB for different separation parameter  $\alpha = 1, 0.75, 0.5$ . Here, the signal length is  $n = 101$ .



**FIGURE 8.** The spike localization using the dual polynomial approach for 2D spikes via ANM. Here, we set  $n_1 = 12$ ,  $n_2 = 10$ , and  $r = 7$ .

the atomic decomposition is unique and exact, as soon as there exists a universal constant  $C > 0$  such that the set of spikes  $\mathcal{T} = \{\boldsymbol{\tau}_k\}_{1 \leq k \leq r}$  satisfies

$$\Delta_{\mathbb{T}^2}(\mathcal{T}) \triangleq \inf_{\substack{\boldsymbol{\tau}, \boldsymbol{\tau}' \in \mathcal{T} \\ \boldsymbol{\tau} \neq \boldsymbol{\tau}'}} \min_{\mathbf{q} \in \mathbb{Z}^2} \|\boldsymbol{\tau} - \boldsymbol{\tau}' + \mathbf{q}\|_{\infty} > \frac{C}{\min(n_1, n_2) - 1}.$$



Moreover, if the signal  $X_{2D}$  is real valued,  $C = 4.76$  suffices to guarantee exact recovery of the spikes.

### Atomic set for multiple measurement vectors

One can collect multiple snapshots of observations, where they share the same spike locations with varying coefficients. Consider  $T$  snapshots, stacked in a matrix,  $X_{MMV} = [x_1, \dots, x_T]$ , which is expressed similarly to (8) as

$$X_{MMV} = \sum_{k=1}^r a(\tau_k) c_k^\top, \quad (23)$$

where  $c_k = [c_{1,k}, \dots, c_{L,k}] \in \mathbb{C}^T$  is the coefficient of the  $k$ th spike across the snapshots. Following the recipe of atomic norms, we define the atoms as

$$A(\tau, b) = a(\tau) b^\top,$$


where  $\tau \in [0, 1)$ ,  $b \in \mathbb{C}^T$  with  $\|b\|_2 = 1$ . The atomic set is defined as

$$\mathcal{A}_{MMV} = \{A(\tau, b) : \tau \in [0, 1), \|b\|_2 = 1\}.$$

The atomic norm can then be defined following (2), which turns out sharing similar nice SDP characterizations for primal and dual formulations as for the single snapshot model [59]. The atomic norm  $\|X_{MMV}\|_{\mathcal{A}}$  can be written equivalently as

$$\|X_{MMV}\|_{\mathcal{A}} = \inf_{\substack{u \in \mathbb{C}^n \\ W \in \mathbb{C}^{T \times T}}} \left\{ \frac{1}{2n} \text{Tr}(\text{toep}(u)) + \frac{1}{2} \text{Tr}(W) : \begin{bmatrix} \text{toep}(u) & X_{MMV} \\ X_{MMV}^H & W \end{bmatrix} \succeq 0 \right\}. \quad (24)$$

A curious comparison can be drawn to the nuclear norm by noticing that one recovers the nuclear norm of  $X_{MMV}$  by replacing the principal block  $\text{toep}(u)$  in (24) with an arbitrary positive semidefinite matrix. The fact that  $\text{toep}(u)$  has significantly fewer degrees of freedom ( $n$  versus  $n^2$ ) is in parallel to that  $a(\tau)$  has significantly fewer degrees of freedom than an arbitrary vector (1 versus  $n$ ).



**FIGURE 9.** The dual polynomial using multiple snapshots (in red) successfully localize all of the spikes while the one using a single snapshot (in blue) fails for the same spike signal. Here,  $r = 6$ ,  $n = 21$ , and  $T = 6$ .

Again, one can determine the atomic decomposition and localize the spikes by resorting to the dual program in a similar fashion. Figure 9 illustrates how multiple snapshots improve the performance of localization over the single snapshot case when the coefficients across snapshots exhibit some kind of diversity, e.g., generated with i.i.d. complex Gaussian entries.

### Generalizations of measurement models

So far, we have seen that ANM provides a means for super-resolution via convex relaxation in additive Gaussian noise. The framework of convex optimization is quite versatile and can be extended to handle models when the measurements are partially observed, corrupted, contain interfering sources, or even come from unknown modulations. This is an important advantage over classical methods such as MUSIC or ESPRIT, as most of them cannot be extended easily to these variants of models.

### Compressed spectral sensing

CS [43], [44] has suggested that it is possible to recover a signal using a number of measurements that is proportional to its degrees of freedom, rather than its ambient dimension. Consider the problem where only a subset of entries of  $x$  is observed,

$$y_{CS} = A_{CS} x,$$

where  $A_{CS} \in \mathbb{C}^{m \times n}$ , and  $m \ll n$  representing compressive acquisition of the signal  $x$ . The goal is to recover  $x$  and its spectral content from  $y_{CS} \in \mathbb{C}^m$ , the compressive measurements. This has applications in wide-band spectrum sensing and cognitive radio [60], for example.

One can easily extend the framework of ANM and recover  $x$  by solving the program

$$\min_x \|x\|_{\mathcal{A}} \quad \text{subject to} \quad y_{CS} = A_{CS} x.$$

When  $A_{CS}$  is a partial observation matrix, namely, a subset of  $m$  entries of  $x$  is observed uniformly at random, then  $x$  can be perfectly recovered with high probability using  $m = O(\log^2 n + r \log r \log n)$  measurements as long as  $x$  satisfies the separation condition, and with random signs of the coefficients  $c_k$ 's [15, Th. II.3]. More generally, a broader class of measurement matrices  $A_{CS}$  can be allowed where its rows are drawn independently from some isotropic and incoherent distribution [61], [62], and exact recovery is possible under the same separation condition using a number of measurements on the order of  $r$  up to some logarithmic factors. In addition, quantized measurements are further dealt in [63] with theoretical guarantees.

### Demixing sinusoids and spikes

Due to sensor failures or malicious environments, the measurements are susceptible to corruptions that can take arbitrary magnitudes. To this end, consider the problem when the observations are contaminated by sparse outliers, where

$$y_{\text{corrupt}} = x + s.$$

Here,  $\mathbf{s}$  is a sparse vector, where its nonzero entries correspond to corruptions of the observations. The goal is to decompose  $\mathbf{x}$  and  $\mathbf{s}$  from  $\mathbf{y}_{\text{corrupt}}$ , a problem intimately related to the uncertainty principle of signal decomposition in [17] and [64] and sparse error correction in CS [65].

Leveraging low-dimensional structures in both  $\mathbf{x}$  and  $\mathbf{s}$ , we seek  $\mathbf{x}$  with a small atomic norm and  $\mathbf{s}$  with a small  $\ell_1$  norm that satisfies the observation constraint [66]

$$\min_{\mathbf{x}, \mathbf{s}} \|\mathbf{x}\|_{\mathcal{A}} + \lambda \|\mathbf{s}\|_1 \quad \text{subject to} \quad \mathbf{y}_{\text{corrupt}} = \mathbf{x} + \mathbf{s},$$

where  $\lambda > 0$  is some regularization parameter. As long as the sample size is sufficiently large [66, Th. 2.2] and the spikes satisfy the separation condition, then the above algorithm perfectly localizes the spikes with high probability, even when the corruption amounts to near a constant fraction of the measurements.

### Demixing interfering sources

A scenario of increasing interest is when the observation is composed of a mixture of responses from multiple exciting or transmitting sources, and the goal is to simultaneously separate and localize the sources at a high resolution. For example, an electrode probing the activities in a brain records firing patterns of a few neighboring neurons, each with a distinct PSF. For pedagogical reasons, let us consider a generalization of the model (6) with two interfering sources, where the observation is given as

$$\mathbf{y}_{\text{mix}} = \text{diag}(\mathbf{g}_1)\mathbf{x}_1 + \text{diag}(\mathbf{g}_2)\mathbf{x}_2,$$

where  $\mathbf{g}_1$  and  $\mathbf{g}_2$  correspond to the frequency-domain response of the PSFs, and  $\mathbf{x}_i = \sum_{k=1}^n c_{i,k} \mathbf{a}(\tau_{i,k})$ ,  $i = 1, 2$ . The goal is to separate and recover the spikes in both  $\mathbf{x}_1$  and  $\mathbf{x}_2$  from  $\mathbf{y}_{\text{mix}}$ , where  $\mathbf{g}_1$  and  $\mathbf{g}_2$  are assumed to be known.

Using ANM, one seeks to recover both  $\mathbf{x}_1$  and  $\mathbf{x}_2$  simultaneously by minimizing the weighted sum of their atomic norms [67]:

$$\min_{\mathbf{x}_1, \mathbf{x}_2} \|\mathbf{x}_1\|_{\mathcal{A}} + \lambda \|\mathbf{x}_2\|_{\mathcal{A}} \\ \text{subject to} \quad \mathbf{y}_{\text{mix}} = \text{diag}(\mathbf{g}_1)\mathbf{x}_1 + \text{diag}(\mathbf{g}_2)\mathbf{x}_2,$$

where  $\lambda > 0$  is some regularization parameter. Unlike the single-source case, the success of demixing critically depends on how easy it is to tell two PSFs apart; the more similar  $\mathbf{g}_1$  and  $\mathbf{g}_2$  are, the harder it is to separate them. A random model can be used to generate dissimilar PSFs, namely it is assumed the entries of  $\mathbf{g}_i \mathbf{s}$  are i.i.d. generated from a uniform distribution over the complex circle [67]. The algorithm then succeeds with high probability as long as the sample size is sufficiently large and the spikes within the same signal satisfy the separation condition [67, Th. 2.1], without requiring any separation for spikes coming from different sources.

### Blind superresolution

So far, all algorithms have assumed the PSF as known, which is a reasonable assumption for problems where one can design

and calibrate the PSF a priori. In general, one might need to estimate the PSF at the same time, possibly due to the fact that the PSF may drift and needs to be calibrated on the fly during deployment. In this case, we need to revisit (6) and estimate  $\mathbf{g}$  and  $\mathbf{x}$  simultaneously from their bilinear measurements,

$$\mathbf{y}_{\text{BR}} = \text{diag}(\mathbf{g})\mathbf{x}.$$

This problem is terribly ill-posed, as the number of unknowns far exceeds the number of observations. One remedy is to exploit additional structures of  $\mathbf{g}$ . For example, if  $\mathbf{g}$  lies in a known low-dimensional subspace  $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n]^T \in \mathbb{C}^{n \times d}$  with  $d \ll n$ , then the degrees of freedom of  $\mathbf{g}$  is greatly dropped, since one only needs to estimate the coefficient  $\mathbf{h} \in \mathbb{C}^d$  of  $\mathbf{g} = \mathbf{B}\mathbf{h}$  in that subspace, which has a much smaller dimension. Even such, the measurement  $\mathbf{y}_{\text{BR}}$  is still bilinear in  $\mathbf{h}$  and  $\mathbf{x}$ , and one cannot directly apply ANM to  $\mathbf{x}$  as it does not lead to a convex program.

Interestingly, a lifting trick can be applied [68], which rewrites  $\mathbf{y}_{\text{BR}} = \mathcal{X}(\mathbf{Z})$  as linear measurements of a higher-dimensional object  $\mathbf{Z} = \mathbf{x}\mathbf{h}^T \in \mathbb{C}^{n \times d}$  similar to (23):

$$\mathbf{y}_{\text{BR},i} = \mathbf{b}_i^T \mathbf{h} \mathbf{e}_i^T \mathbf{x} = \langle \mathbf{x}\mathbf{h}^T, \mathbf{e}_i \mathbf{b}_i^H \rangle, \quad i = 1, \dots, n,$$

where  $\mathbf{e}_i$  is the  $i$ th standard basis vector. Consequently, one can apply ANM to  $\mathbf{Z}$  with respect to (24), leading to the algorithm [68]

$$\min_{\mathbf{Z}} \|\mathbf{Z}\|_{\mathcal{A}} \quad \text{subject to} \quad \mathbf{y}_{\text{BR}} = \mathcal{X}(\mathbf{Z}).$$

This approach succeeds with high probability; as soon as the sample size is sufficiently large, the spikes are well separated and the PSF satisfies certain incoherence properties [68, Th. 1]. Moreover, it can be further extended to demixing a mixture of sources with unknown PSFs, where each PSF lies in a distinct subspace [69].

## Beyond line spectrum estimation: Superresolution imaging for single-molecule fluorescence microscopy

When the atomic set is composed of a family of complex sinusoidal signals, an exact implementation of the atomic norm in SDP is available. In the most general setting, ANM is an infinite-dimensional convex program whose computation needs to be addressed carefully. Encouragingly, tailored solvers have been proposed and applied successfully to practical applications such as superresolution imaging for single-molecule fluorescence microscopy, which we now present as a case study to show its promise.

### Imaging principle

The development of superresolution fluorescence microscopy, which was awarded the 2014 Nobel Prize in Chemistry, is considered to fundamentally impact biological science and medicine. To date, a partial list of superresolution fluorescence microscopy technologies includes PALM [70], STORM [71], and fPALM [72], which share a similar imaging principle. A very

nice introduction can be found in [73]. While optical microscopy is desirable for imaging complex biological processes in live cells due to its noninvasive nature, due to diffraction limit, which is about hundreds of nanometers, it cannot image detailed internal structures of cells, which are often below 100 nm.

To deal with this challenge, biologists have come up with a clever idea of divide and conquer. To begin, imagine that every point within a cell is equipped with a photoswitchable fluorescent molecule, which means, once excited, the molecule will emit light stochastically over a duration of time to identify its location. This allows one to divide the imaging process into many frames, where in each frame a random and sparse subset of fluorescent molecules (point sources) are activated and localized at a resolution below the diffraction limit using imaging algorithms. The final image is thus obtained by superimposing the localization outcomes of all of the frames. Therefore, the high spatial resolution is achieved by sacrificing the temporal resolution. To speed up the imaging process and improve the temporal resolution, it is desirable to develop localization algorithms that are capable of identifying more fluorescent molecules per frame, which is known as the *emitter density*.

Very interestingly, this imaging principle can be used to reconstruct a 3D biological structure from 2D image frames [74]. One way is to introduce a cylindrical lens to modulate the ellipticity of the PSF based on the depth of the fluorescent object, which can be modeled as a Gaussian pulse with varying ellipticity along the  $x$  and  $y$  directions,

$$g(x, y | z) = \frac{1}{2\pi\sigma_x(z)\sigma_y(z)} e^{-\left[\frac{x^2}{2\sigma_x(z)^2} + \frac{y^2}{2\sigma_y(z)^2}\right]},$$


where  $\sigma_x(z)$  and  $\sigma_y(z)$  are functions of the depth in the  $z$  direction, and can be calibrated in advance. For a 3D scene of point sources,

$$\zeta(x, y, z) = \sum_{i=1}^r c_i \delta(x - x_i, y - y_i, z - z_i),$$

its convolution with the PSF  $g(x, y | z)$  is given as a 2D image in the form of

$$(\zeta * g)(x, y) = \sum_{i=1}^r \frac{c_i}{2\pi\sigma_x(z_i)\sigma_y(z_i)} e^{-\left[\frac{(x-x_i)^2}{2\sigma_x(z_i)^2} + \frac{(y-y_i)^2}{2\sigma_y(z_i)^2}\right]}.$$

Therefore, to perform superresolution, one needs to decode simultaneously the ellipticity as well as the location of the PSF, which is much more challenging. In practice, the situation is even more complex, since the continuous spatial function  $(\zeta * g)(x, y)$



**FIGURE 10.** The mathematical model of 3D imaging in superresolution fluorescence microscopy.

needs to be further discretized due to pixelization of the detector, leading to a discretized 2D image,  $\omega \in \mathbb{R}^{n_1 \times n_2}$ , where each entry of  $\omega$  corresponds to the integration of  $(\zeta * g)(x, y)$  over the area of a pixel. The final image,  $z$ , counting the number of photons hitting the detector at each pixel, is modeled as a Poisson distribution with rate  $\omega$ . The whole process is summarized in Figure 10.

### Applying ANM

Luckily, due to linearity,  $\omega$  can be viewed a sparse superposition of atoms that are parameterized by the 3D point sources,

$$\omega = \sum_{i=1}^r c_i \mathbf{a}(x_i, y_i, z_i) = \int_{x,y,z} \mathbf{a}(x, y, z) d\zeta(x, y, z),$$

where each atom  $\mathbf{a}(x, y, z)$  corresponds to the image of a point source at  $(x, y, z)$  after convolution and pixelization. The atomic set is then given as  $\mathcal{A}_{3D} = \{\mathbf{a}(x, y, z) : x, y, z \in \text{imaging range}\}$ . The goal is thus to recover  $\zeta(x, y, z)$ , or the atomic decomposition of  $\omega$ , from the observation  $z \sim \text{Poisson}(\omega)$  as accurately as possible.

A natural approach is to seek the sparsest  $\omega$  such that the likelihood function of the observation  $z$  is maximized. To that end, we consider a constrained maximum likelihood estimation, where we seek to solve  $\omega$  via

$$\min_{\omega} -\log p(z | \omega) \quad \text{subject to} \quad \|\omega\|_{\mathcal{A}} \leq \eta, \quad (25)$$

where  $p(z | \omega)$  is the Poisson likelihood function,  $\|\omega\|_{\mathcal{A}}$  is the induced atomic norm with respect to  $\mathcal{A}_{3D}$ , and  $\eta$  is some regularization parameter that may be tuned in practice.

Early efforts such as CSSTORM [75], which are based on  $\ell_1$  minimization by directly discretizing the parameter space, suffer from high computational complexity, due to the need of storing and manipulating a large dictionary of atoms, as fine discretization is required along all three spatial dimensions. On the other end, recent algorithmic developments such as alternating descent conditional gradient (ADCG) [76] and CoGenT [77] solve sparse inverse problems over continuous dictionaries with general convex loss functions at much reduced memory and computation requirements. In a nutshell, the ADCG method is an acceleration of conditional gradient, also known as *Frank–Wolfe*, to solve (25) with a general atomic set  $\mathcal{A} = \{\mathbf{a}(\theta) : \theta \in \Theta\}$ , where  $\theta$  is a short-hand notation for the parameter space. In particular, it directly estimates the atomic decomposition of  $\omega = \int \mathbf{a}(\theta) d\zeta(\theta)$ .

The standard Frank–Wolfe adds one new atom at every iteration to reduce the negative log-likelihood function; however, it will introduce many spurious atoms and lose sparsity as the iteration increases. To deal with this, ADCG introduces pruning and local refinements with a hope to maintain a sparse representation at all iterations. The detailed procedure of ADCG is given in Algorithm 2. In the  $j$ th iteration, denote the support and coefficient of the current estimate of  $\zeta(\theta)$  as  $\mathcal{T}_j$  and  $\mathbf{c}_j$ , and the current estimate of  $\omega$  as  $\mathcal{A}(\mathcal{T}_j)\mathbf{c}_j$ . Like Frank–Wolfe, ADCG starts by adding a spike to the estimated support  $\mathcal{T}_j$  that maximally correlates with the derivative of  $\log p(z | \omega)$

with respect to  $\omega$  at the current estimate,  $\mathcal{A}(\mathcal{T}_j)\mathbf{c}_j$ . Because the spike location will be refined next, in practice, this step can be solved approximately by searching over a coarse grid of  $\Theta$  to save computation.

ADCG then deviates from the standard Frank–Wolfe, and tries to improve the updated estimate by performing alternating descent over the coefficient and the support. It iterates between coefficients update via  $\ell_1$  minimization, support pruning, and local refinement of the support by holding the coefficients fixed. The last step leverages the fact that  $\mathbf{a}(\theta)$  is differentiable with respect to  $\theta$ , and a simple local search via gradient descent allows one to adjust the support to further reduce the loss function. Despite the nonconvexity in this refine step, [76] guarantees the convergence of ADCG with a convergence rate of  $O(1/\epsilon)$  to reach  $\epsilon$ -accuracy in the function value, under some technical assumptions. In practice, the main computational benefits of ADCG are the absence of semidefinite constraints and small memory footprints, making it highly suitable for large-scale implementations.

TVSTORM [78] is a modification of ADCG tailored to solve (26) for 3D image reconstruction with some domain adaptations to speed up implementations. TVSTORM outperforms CSSTORM both in terms of computational time and reconstruction quality. Figure 11(a) shows the diffraction-limited imaging using conventional microscopy; in contrast, the 3D superresolution image reconstructed using TVSTORM in Figure 11(b) is much clearer, where the structure of 3D microtubules can be well resolved with the axial coordinate represented in different colors. Figure 11(c) and (d) compares the reconstruction quality of an enlarged region between TVSTORM and CSSTORM, where TVSTORM provides a visually more smooth reconstruction of the line structures in microtubules. Figure 12 shows that TVSTORM indeed has a higher detection rate and a lower false discovery rate than CSSTORM, while TVSTORM executes much faster.

## Final remarks

In this article, we presented an overview on how to leverage sparsity for continuous parameter estimation via the mathematical concept of atomic norms, which can be regarded as a generalization of the principle of  $\ell_1$  norms for discrete model selection. We showcased its application in superresolution from low-pass observations in single-molecule fluorescence microscopy. The appeal of the atomic norm approach is attributed to its elegant mathematical framework, strong performance guarantees, and promises of scalable numerical implementations.

The atomic norm is only one of many possible approaches to exploit sparsity over the continuum. One competitive alternative is structured low-rank matrix optimization [79]–[82]; see [83] for its connections and comparisons with the atomic norm approach.

Another line of work [84]–[86] generalizes the traditional CS to an infinite-dimensional Hilbert space. In addition, sampling theorems are developed for signals with a finite rate of innovations together with strategies for perfect reconstruction [87]. More recently, a sparse functional framework has been proposed as a

variational approach to handle sparsity over continuous and possibly nonlinear dictionaries. This category of estimators aims to recover functions with minimum support measure subject to the observation constraint [88], [89].

As a topic still under development, open problems abound for both theoretical and practical performance of optimization-based superresolution in general. We conclude by outlining some exciting future directions.

- *Tight performance analysis in noise:* Existing analyses of atomic norm denoising (21) typically only produce bounds that are tight up to some constant, making them less useful in practice. For example, attempts to benchmark the theoretical bounds against the CRB will be in vain due to the presence of large constants. It is therefore desirable to obtain tight performance bounds such as the one available for matrix denoising [90] that is asymptotically exact.
- *Adaptive selection of regularization parameters:* One benefit of ANM over traditional spectrum estimation approaches is that it can automatically select the model order. However, the choice of the regularization parameter (21) depends on the noise level, which may not be available in practice. How to optimally set the regularization parameters is another problem of great importance; see [91] for some recent developments.
- *Low-rank factorization for SDP formulations of atomic norms:* A popular heuristic to SDP with low-rank solutions is to apply low-rank matrix factorization and solve the corresponding nonconvex optimization problem [92], [93], with the premise of greatly reducing its computational cost. It will be interesting to see if this approach can be


## Algorithm 2. Alternating descent conditional gradient [76].

**Input:** Observation  $\mathbf{z}$ ; Parameter  $\eta > 0$ ;  
Initialize  $\mathcal{T}_0$  to an empty set,  $\mathbf{c}_0$  to  $\mathbf{0}$ , and  $j = 0$ ;  
**repeat** until stopping criteria  
  Localize the next spike:  

$$\theta_{j+1} \in \operatorname{argmax}_{\theta \in \Theta} \langle \mathbf{a}(\theta), \nabla_{\omega} \log p(\mathbf{z} \mid \mathcal{A}(\mathcal{T}_j)\mathbf{c}_j) \rangle$$
  
  Update support:  $\mathcal{T}_{j+1} \leftarrow \mathcal{T}_j \cup \{\theta_{j+1}\}$ ;  
  Refinement: **repeat**  
    1) Update the amplitudes:  

$$\mathbf{c}_{j+1} \leftarrow \operatorname{argmin}_{\|\mathbf{c}\|_1 \leq \eta} -\log p(\mathbf{z} \mid \mathcal{A}(\mathcal{T}_{j+1})\mathbf{c});$$
  
    2) Prune support:  $\mathcal{T}_{j+1} \leftarrow \operatorname{support}(\mathbf{c}_{j+1})$ ;  
    3) Local descent: improve  $\mathcal{T}_{j+1}$  by performing local descent on  $-\log p(\mathbf{z} \mid \mathcal{A}(\mathcal{T}_{j+1})\mathbf{c}_{j+1})$  holding the coefficient  $\mathbf{c}_{j+1}$  fixed;  
   $j \leftarrow j + 1$ ;  
**Output:**  $(\mathcal{T}_j, \mathbf{c}_j)$  and  $\mathbf{x}_j = \mathcal{A}(\mathcal{T}_j)\mathbf{c}_j$ .





**FIGURE 11.** (a) Diffraction-limited imaging of microtubules using conventional microscopy and (b) superresolution 3D image reconstruction using TVSTORM. A comparison of reconstruction quality between (c) TVSTORM and (d) CSSTORM. (bar:  $1.5 \mu\text{m}$ ). (Source: [78].)

applied to speed up the computation of atomic norms with performance guarantees.

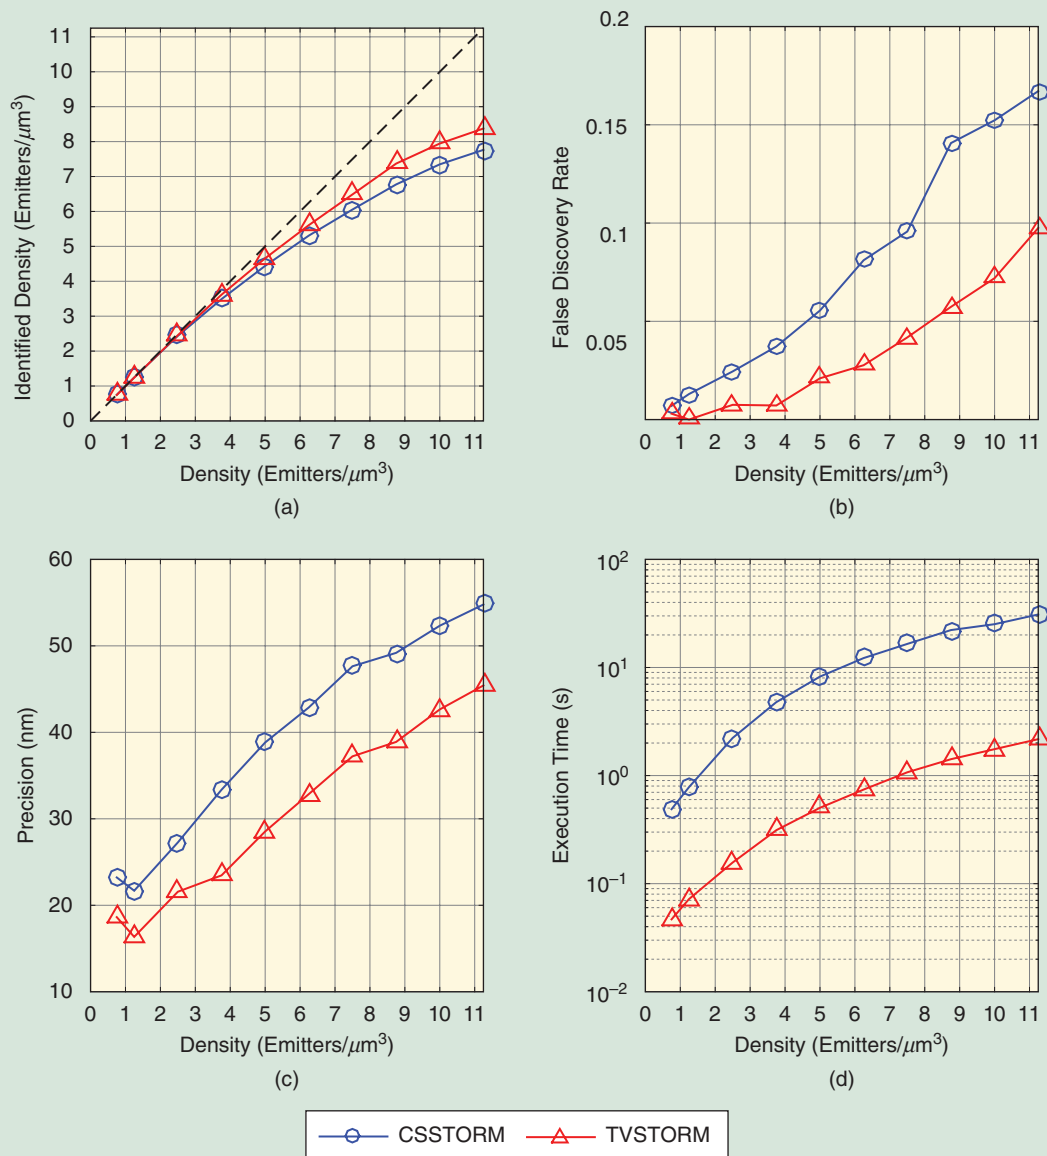
- *Bridging classical and optimization-based approaches:* There are deep connections between traditional (e.g. Prony, MUSIC, and so on) and optimization-based (e.g. ANM and nuclear norm minimization) approaches. Such connections have already been realized, for instance in the early works of Fuchs [94], where he provided an optimization interpretation of the Pisarenko method [95]. Another recent work [96] provided an optimization view to the MUSIC algorithm. It is hopeful that a confluence of past and current ideas will likely deepen our understandings and lead to further algorithmic improvements.
- *ANM for more general measurement models:* While there have been significant advances in the understanding of ANM for line spectrum estimation, its theory and

application to other measurement settings require further investigation.

- *Applications in communications, sensing, and imaging:* ANM has recently emerged as a popular approach for many practical applications, such as channel estimation in massive MIMO [97], [98], radar imaging [99], and nuclear magnetic resonance spectroscopy [100]. It is our hope that this article will stir more interest in applying the atomic norm in applications that call for high-resolution parameter estimation.

## Authors

**Yuejie Chi** (yuejiechi@cmu.edu) received her B.E. (Hon.) degree in electrical engineering from Tsinghua University, Beijing, China, in 2007 and her Ph.D. degree in electrical engineering from Princeton University, New Jersey, in 2012. She is currently an associate professor with the Department of



**FIGURE 12.** Comparisons between TVSTORM and CSSTORM for various performance metrics of 3D image reconstruction: (a) identified density, (b) false discovery rate, (c) precision, and (d) execution time with respect to the emitter density. (Source: [78].)

Electrical and Computer Engineering at Carnegie Mellon University, Pittsburgh, Pennsylvania, where she holds the Robert E. Doherty Early Career Development Professorship. She has received multiple awards, and she serves as an associate editor of *IEEE Transactions on Signal Processing*. Her research interests include statistical signal processing, machine learning, and large-scale optimization and their applications in data science, inverse problems, imaging, and sensing systems. She is a Senior Member of the IEEE.

**Maxime Ferreira Da Costa** (mferreira@cmu.edu) received his M.Sc. degree in signal processing from Imperial College London, United Kingdom, Diplôme d'Ingénieur from Supélec, France, in 2012, and Ph.D. degree in electrical and electronic engineering from Imperial College London in 2018. He is cur-

rently a research associate with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania. His research focuses on mathematical signal processing, inverse problems, and optimization. In 2018, he was shortlisted among the finalists for the Jack Keil Wolf Student Paper Award at the IEEE International Symposium on Information Theory. He is a Member of the IEEE.

### Acknowledgments

We would like to thank *IEEE Signal Processing Magazine's* editorial board for early feedback on our white paper and the associate editor and anonymous reviewers for their constructive suggestions that helped to improve the quality of this article. This work is supported in part by the Office of Naval

Research under grants N00014-18-1-2142 and N00014-19-1-2404 and by the National Science Foundation under grants CIF-1826519 and ECCS-1818571.

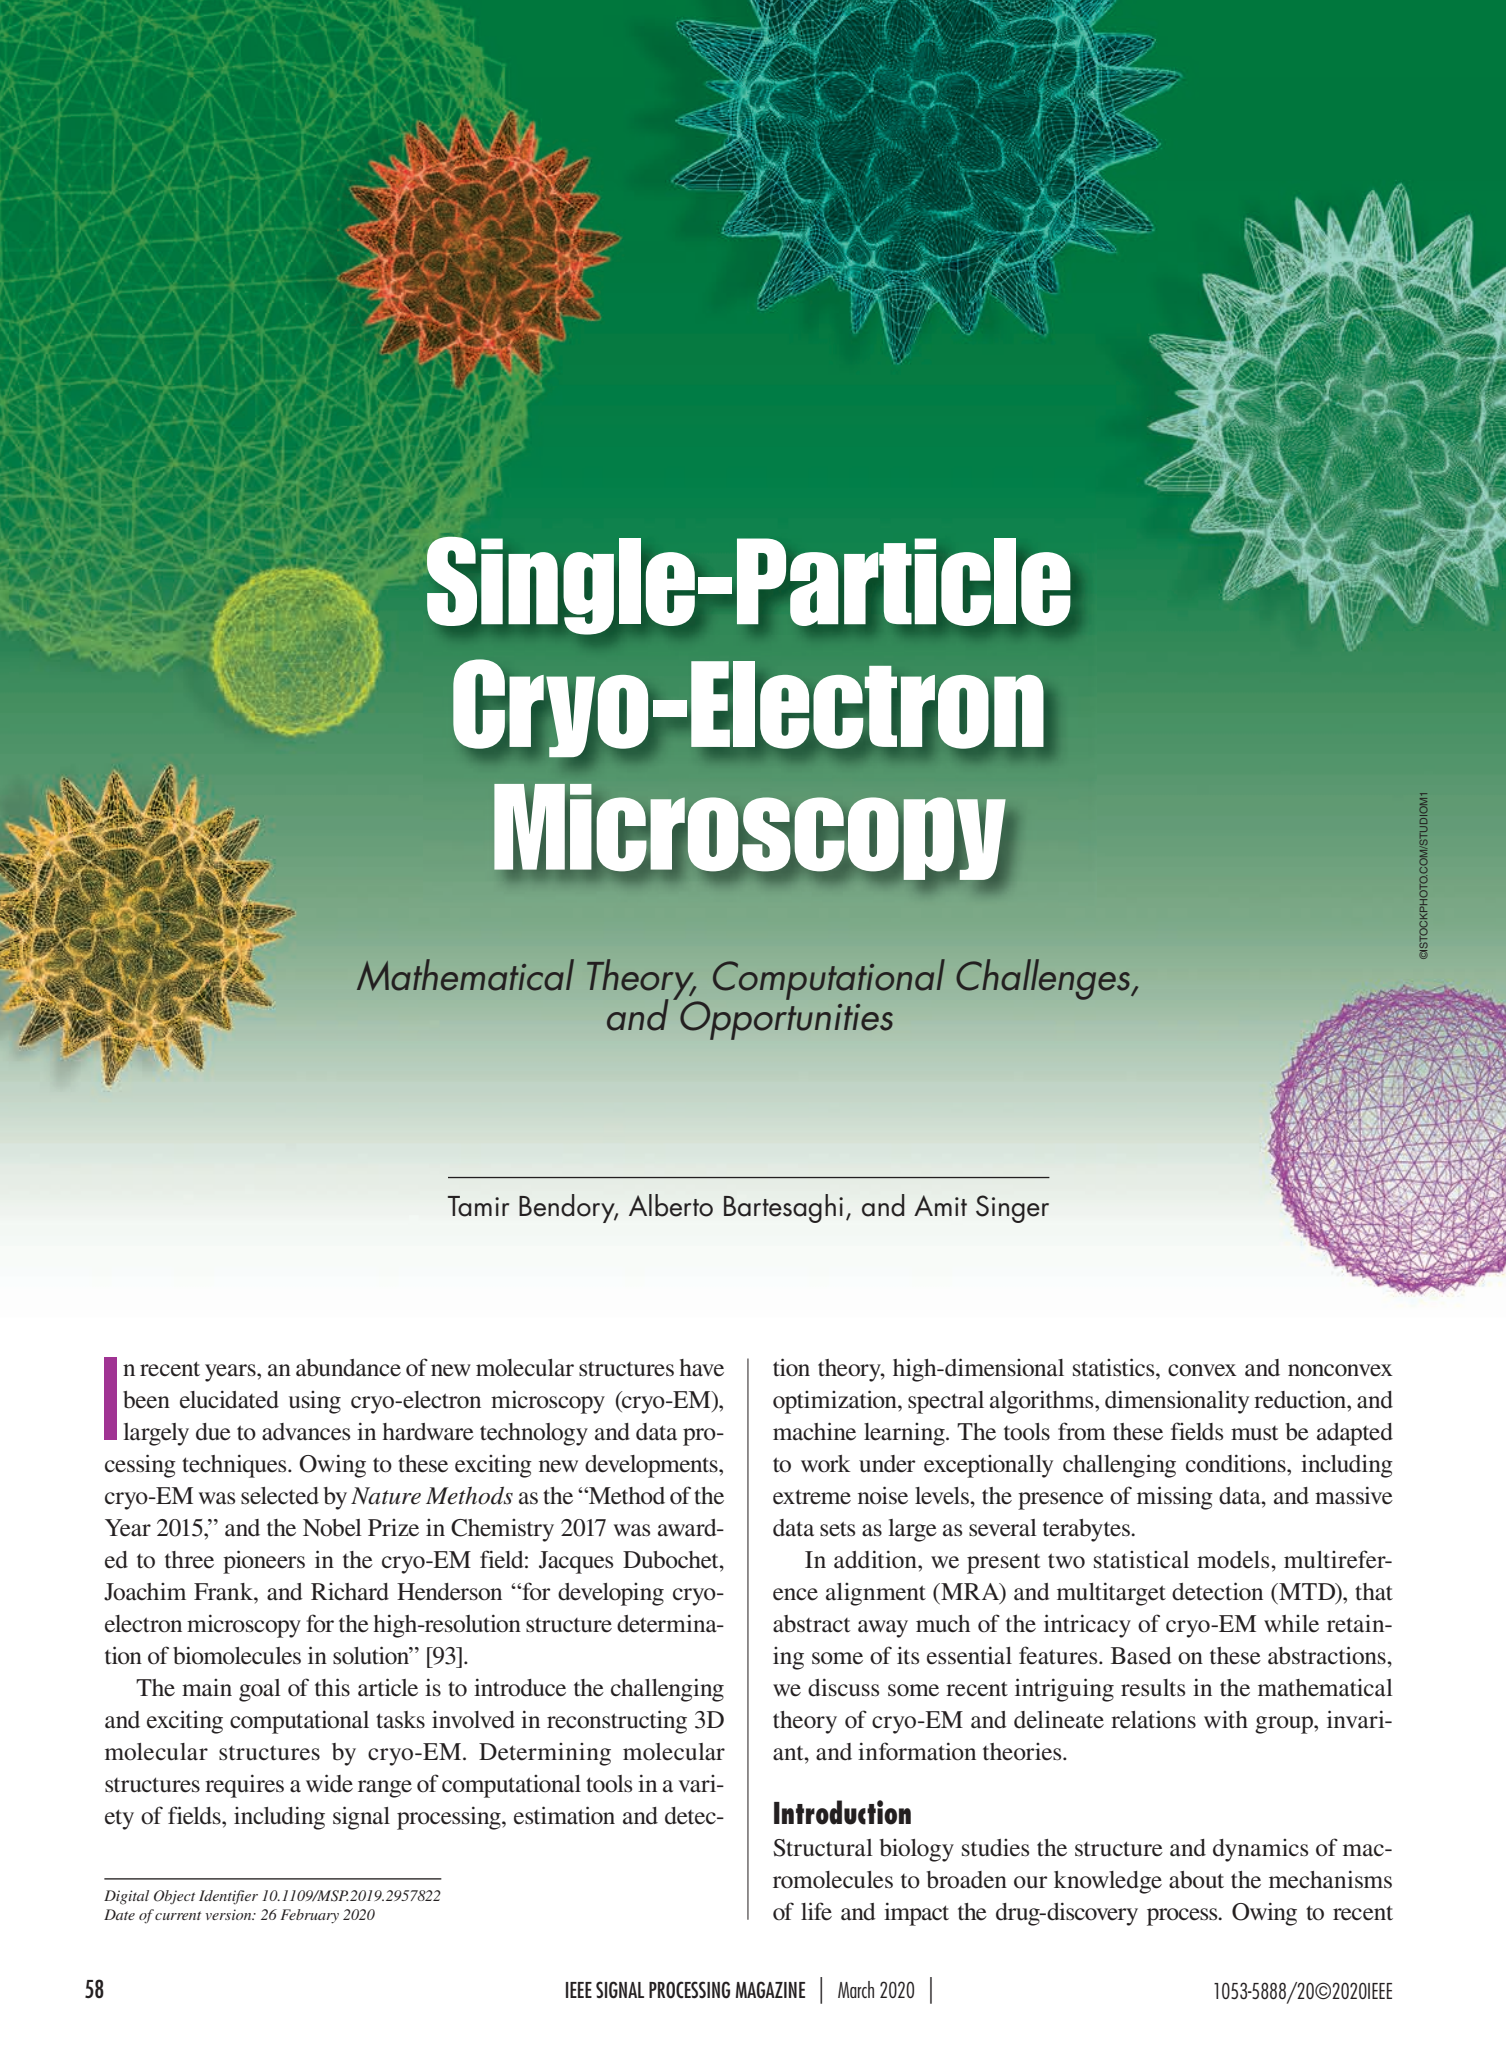
## References

- [1] P. Stoica and R. L. Moses, *Introduction to Spectral Analysis*, vol. 1. Englewood Cliffs, NJ: Prentice Hall, 1997.
- [2] S. M. Kay and S. L. Marple, "Spectrum analysis: A modern perspective," *Proc. IEEE*, vol. 69, no. 11, pp. 1380–1419, 1981. doi: 10.1109/PROC.1981.12184.
- [3] H. Krim and M. Viberg, "Two decades of array signal processing research: The parametric approach," *IEEE Signal Process. Mag.*, vol. 13, no. 4, pp. 67–94, 1996. doi: 10.1109/79.526899.
- [4] R. Prony, "Essai experimental et analytique, sur les lois de la dilatabilité des fluides de la vapeur de l'alkool, à différentes températures," *J. de l'Ecole Polytechnique Floréal et Plairial*, vol. 1, no. 2, pp. 24–76, 1795.
- [5] R. Kumaresan, D. Tufts, and L. L. Scharf, "A Prony method for noisy data: Choosing the signal components and selecting the order in exponential signal models," *Proc. IEEE*, vol. 72, no. 2, pp. 230–233, 1984. doi: 10.1109/PROC.1984.12849.
- [6] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, 1986. doi: 10.1109/TAP.1986.1143830.
- [7] R. Roy and T. Kailath, "ESPRIT-estimation of signal parameters via rotational invariance techniques," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 7, pp. 984–995, July 1989. doi: 10.1109/29.32276.
- [8] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 5, pp. 814–824, May 1990. doi: 10.1109/29.56027.
- [9] P. Stoica, R. L. Moses, B. Friedlander, and T. Soderstrom, "Maximum likelihood estimation of the parameters of multiple sinusoids from noisy measurements," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 3, pp. 378–392, 1989. doi: 10.1109/29.21705.
- [10] M. P. Clark and L. L. Scharf, "Two-dimensional modal analysis based on maximum likelihood," *IEEE Trans. Signal Process.*, vol. 42, no. 6, pp. 1443–1452, 1994. doi: 10.1109/78.286959.
- [11] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Commun. Pure Appl. Math.*, vol. 67, no. 6, pp. 906–956, 2014. doi: 10.1002/cpa.21455.
- [12] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Found. Comput. Math.*, vol. 12, no. 6, pp. 805–849, 2012. doi: 10.1007/s10208-012-9135-7.
- [13] S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, 2001. doi: 10.1137/S003614450037906X.
- [14] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010. doi: 10.1137/070697835.
- [15] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, "Compressed sensing off the grid," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7465–7490, 2013. doi: 10.1109/TIT.2013.2277451.
- [16] B. Huang, M. Bates, and X. Zhuang, "Super-resolution fluorescence microscopy," *Annu. Rev. Biochem.*, vol. 78, no. 1, pp. 993–1016, 2009. doi: 10.1146/annurev.biochem.77.061906.092014.
- [17] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, 2001. doi: 10.1109/18.959265.
- [18] Y. Chen and Y. Chi, "Harnessing structures in big data via guaranteed low-rank matrix estimation: Recent theory and fast algorithms via convex and nonconvex optimization," *IEEE Signal Process. Mag.*, vol. 35, no. 4, pp. 14–31, 2018. doi: 10.1109/MSP.2018.2821706.
- [19] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006. doi: 10.1002/cpa.20124.
- [20] D. Donoho, "For most large underdetermined systems of linear equations the minimal  $\ell_1$ -norm solution is also the sparsest solution," *Commun. Pure Appl. Math.*, vol. 59, no. 6, pp. 797–829, 2006. doi: 10.1002/cpa.20132.
- [21] C. Carathéodory, "Über den variabilitätsbereich der koeffizienten von potenzreihen, die gegebene werte nicht annehmen," *Math. Ann.*, vol. 64, no. 1, pp. 95–115, 1907. doi: 10.1007/BF01449883.
- [22] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Berlin: Springer-Verlag, 2013.
- [23] R. Sanyal, F. Sottile, and B. Sturmfels, "Orbitopes," *Mathematika*, vol. 57, no. 2, pp. 275–314, 2011. doi: 10.1112/S002557931100132X.
- [24] J.-B. Lasserre, *Moments, Positive Polynomials and Their Applications*, vol. 1. Singapore: World Scientific, 2010.
- [25] B. Dumitrescu, *Positive Trigonometric Polynomials and Signal Processing Applications*, vol. 103. Berlin: Springer-Verlag, 2007.
- [26] T. T. Georgiou, "The Carathéodory–Fejér–Pisarenko decomposition and its multivariable counterpart," *IEEE Trans. Autom. Control*, vol. 52, no. 2, pp. 212–228, 2007. doi: 10.1109/TAC.2006.890479.
- [27] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming*. (2018). [Online]. Available at <http://cvxr.com/cvx>
- [28] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Math.*, vol. 346, nos. 9–10, pp. 589–592, 2008. doi: 10.1016/j.crma.2008.03.014.
- [29] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007. doi: 10.1109/TIT.2007.909108.
- [30] C. Fernandez-Granda, "Super-resolution of point sources via convex programming," *Inform. Inference*, vol. 5, no. 3, pp. 251–303, 2016. doi: 10.1093/imaia/iaw005.
- [31] M. F. Da Costa and W. Dai, "A tight converse to the spectral resolution limit via convex programming," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, pp. 901–905. doi: 10.1109/ISIT.2018.8437490.
- [32] B. N. Bhaskar, G. Tang, and B. Recht, "Atomic norm denoising with applications to line spectral estimation," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5987–5999, 2013. doi: 10.1109/TSP.2013.2273443.
- [33] R. Tibshirani, "Regression shrinkage and selection via the Lasso," *J. Roy. Stat. Soc. B, Methodol.*, vol. 58, no. 1, pp. 267–288, 1996. doi: 10.1111/j.2517-6161.1996.tb02080.x.
- [34] G. Tang, B. N. Bhaskar, and B. Recht, "Near minimax line spectral estimation," *IEEE Trans. Inf. Theory*, vol. 61, no. 1, pp. 499–512, 2015. doi: 10.1109/TIT.2014.2368122.
- [35] C. Fernandez-Granda, "Support detection in super-resolution," in *Proc. 10th Int. Conf. Sampling Theory and Applications (SampTA 2013)*, 2013, pp. 145–148.
- [36] Q. Li and G. Tang, "Approximate support recovery of atomic line spectral estimation: A tale of resolution and precision," *Appl. Comput. Harmon. Anal.*, to be published. doi: 10.1016/j.acha.2018.09.005. [Online]. Available: <https://www.science-direct.com/science/article/abs/pii/S1063520318300824>
- [37] V. Duval and G. Peyré, "Exact support recovery for sparse spikes deconvolution," *Found. Comput. Math.*, vol. 15, no. 5, pp. 1315–1355, 2015. doi: 10.1007/s10208-014-9228-6.
- [38] M. F. Da Costa and Y. Chi, "On the stable resolution limit of total variation regularization for spike deconvolution. 2019. [Online]. Available: [arXiv:1910.01629](https://arxiv.org/abs/1910.01629).
- [39] T. Blu, P. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 31–40, Mar. 2008. doi: 10.1109/MSP.2007.914998.
- [40] B. D. Rao and K. S. Hari, "Performance analysis of root-MUSIC," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 12, pp. 1939–1949, 1989. doi: 10.1109/29.455540.
- [41] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2010. doi: 10.1561/22000000016.
- [42] G. H. Golub and H. A. van der Vorst, "Eigenvalue computation in the 20th century," *J. Comput. Appl. Math.*, vol. 123, nos. 1–2, pp. 35–65, 2000. doi: 10.1016/S0377-0427(00)00413-1.
- [43] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006. doi: 10.1109/TIT.2006.885507.
- [44] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006. doi: 10.1109/TIT.2006.871582.
- [45] G. Tang, B. N. Bhaskar, and B. Recht, "Sparse recovery over continuous dictionaries-just discretize," in *Proc. IEEE Asilomar Conf. Signals, Systems and Computers*, 2013, pp. 1043–1047. doi: 10.1109/ACSSC.2013.6810450.
- [46] V. Duval and G. Peyré, "Sparse regularization on thin grids I: The LASSO," *Inverse Probl.*, vol. 33, no. 5, p. 055008, 2017. doi: 10.1088/1361-6420/aa5e12.
- [47] V. Duval and G. Peyré, "Sparse spikes super-resolution on thin grids II: The continuous basis pursuit," *Inverse Probl.*, vol. 33, no. 9, p. 095008, 2017. doi: 10.1088/1361-6420/aa7fce.
- [48] Y. Chi, L. Scharf, A. Pezeshki, and A. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2182–2195, May 2011. doi: 10.1109/TSP.2011.2112650.
- [49] Z. Tan, P. Yang, and A. Nehorai, "Joint sparse recovery method for compressed sensing with structured dictionary mismatches," *IEEE Trans. Signal Process.*, vol. 62, no. 19, pp. 4997–5008, 2014. doi: 10.1109/TSP.2014.2343940.
- [50] B. Mamandipoor, D. Ramasamy, and U. Madhoo, "Newtonized orthogonal matching pursuit: Frequency estimation over the continuum," *IEEE Trans. Signal Process.*, vol. 64, no. 19, pp. 5066–5081, 2016. doi: 10.1109/TSP.2016.2580523.
- [51] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proc. Nat. Academy Sci. U. S. A.*, vol. 102, no. 27, pp. 9446–9451, 2005. doi: 10.1073/pnas.0502269102.



- [52] J.-J. Fuchs, "Sparsity and uniqueness for some specific under-determined linear systems," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, vol. 5. Philadelphia: IEEE, 2005, pp. v729–v732. doi: 10.1109/ICASSP.2005.1416407.
- [53] Q. Denoyelle, V. Duval, and G. Peyré, "Support recovery for sparse super-resolution of positive measures," *J. Fourier Anal. Appl.*, vol. 23, no. 5, pp. 1153–1194, 2017. doi: 10.1007/s00041-016-9502-x.
- [54] G. Schiebinger, E. Robeva, and B. Recht, "Superresolution without separation," *Inform. Inference*, vol. 7, no. 1, pp. 1–30, 2017. doi: 10.1093/imaia/iax006.
- [55] A. Eftekhari, J. Tanner, A. Thompson, B. Toader, and H. Tyagi, "Sparse non-negative super-resolution—simplified and stabilised," *Appl. Comput. Harmon. Anal.*, to be published. doi: 10.1016/j.acha.2019.08.004.
- [56] V. I. Morgenshtern and E. J. Candes, "Super-resolution of positive sources: The discrete setup," *SIAM J. Imag. Sci.*, vol. 9, no. 1, pp. 412–444, 2016. doi: 10.1137/15M1016552.
- [57] Y. Chi and Y. Chen, "Compressive two-dimensional harmonic retrieval via atomic norm minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1030–1042, 2015. doi: 10.1109/TSP.2014.2386283.
- [58] W. Xu, J.-F. Cai, K. V. Mishra, M. Cho, and A. Kruger, "Precise semidefinite programming formulation of atomic norm minimization for recovering d-dimensional ( $d \geq 2$ ) off-the-grid frequencies," in *Proc. IEEE Information Theory and Applications Workshop (ITA)*, 2014, pp. 1–4. doi: 10.1109/ITA.2014.6804267.
- [59] Y. Li and Y. Chi, "Off-the-grid line spectrum denoising and estimation with multiple measurement vectors," *IEEE Trans. Signal Process.*, vol. 64, no. 5, pp. 1257–1269, doi: 10.1109/TSP.2015.2496294.
- [60] Z. Tian and G. Giannakis, "Compressed sensing for wideband cognitive radios," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2007)*, vol. 4. Honolulu, HI: IEEE, 2007, pp. 1357–1360. doi: 10.1109/ICASSP.2007.367330.
- [61] E. J. Candès and Y. Plan, "A probabilistic and RIPless theory of compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7235–7254, 2011. doi: 10.1109/TIT.2011.2161794.
- [62] R. Heckel and M. Soltanolkotabi, "Generalized line spectral estimation via convex optimization," *IEEE Trans. Inf. Theory*, vol. 64, no. 6, pp. 4001–4023, 2018. doi: 10.1109/TIT.2017.2757003.
- [63] H. Fu and Y. Chi, "Quantized spectral compressed sensing: Cramer–Rao bounds and recovery algorithms," *IEEE Trans. Signal Process.*, vol. 66, no. 12, pp. 3268–3279, 2018. doi: 10.1109/TSP.2018.2827326.
- [64] D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM J. Appl. Math.*, vol. 49, no. 3, pp. 906–931, 1989. doi: 10.1137/0149053.
- [65] X. Li, "Compressed sensing and matrix completion with constant proportion of corruptions," *Constr. Approx.*, vol. 37, no. 1, pp. 73–99, 2013. doi: 10.1007/s00365-012-9176-9.
- [66] C. Fernandez-Granda, G. Tang, X. Wang, and L. Zheng, "Demixing sines and spikes: Robust spectral super-resolution in the presence of outliers," *Inform. Inference*, vol. 7, no. 1, pp. 105–168, 2017. doi: 10.1093/imaia/iax005.
- [67] Y. Li and Y. Chi, "Stable separation and super-resolution of mixture models," *Appl. Comput. Harmon. Anal.*, vol. 46, no. 1, pp. 1–39, 2019. doi: 10.1016/j.acha.2017.03.003.
- [68] Y. Chi, "Guaranteed blind sparse spikes deconvolution via lifting and convex optimization," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 782–794, 2016. doi: 10.1109/JSTSP.2016.2543462.
- [69] D. Yang, G. Tang, and M. B. Wakin, "Super-resolution of complex exponentials from modulations with unknown waveforms," *IEEE Trans. Inf. Theory*, vol. 62, no. 10, pp. 5809–5830, 2016. doi: 10.1109/TIT.2016.2586083.
- [70] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacio, M. W. Davidson, J. Lippincott-Schwartz et al., "Imaging intracellular fluorescent proteins at nanometer resolution," *Science*, vol. 313, no. 5793, pp. 1642–1645, 2006. doi: 10.1126/science.1127344.
- [71] M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)," *Nat. Meth.*, vol. 3, no. 10, pp. 793–796, 2006. doi: 10.1038/nmeth929.
- [72] S. T. Hess, T. P. Girirajan, and M. D. Mason, "Ultra-high resolution imaging by fluorescence photoactivation localization microscopy," *Biophys. J.*, vol. 91, no. 11, pp. 4258–4272, 2006. doi: 10.1529/biophysj.106.091116.
- [73] M. Ehrenberg, "Scientific background on the Nobel Prize in Chemistry 2014," The Nobel Prize, 2014. [Online]. Available: <https://www.nobelprize.org/uploads/2018/06/advanced-chemistryprize2014-1.pdf> or <https://www.nobelprize.org/prizes/chemistry/2014/advanced-information/>
- [74] B. Huang, W. Wang, M. Bates, and X. Zhuang, "Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy," *Science*, vol. 319, no. 5864, pp. 810–813, 2008. doi: 10.1126/science.1153529.
- [75] L. Zhu, W. Zhang, D. Elnatán, and B. Huang, "Faster STORM using compressed sensing," *Nat. Meth.*, vol. 9, no. 7, pp. 721–723, 2012. doi: 10.1038/nmeth.1978.
- [76] N. Boyd, G. Schiebinger, and B. Recht, "The alternating descent conditional gradient method for sparse inverse problems," *SIAM J. Optim.*, vol. 27, no. 2, pp. 616–639, 2017. doi: 10.1137/15M1035793.
- [77] N. Rao, P. Shah, and S. Wright, "Forward–backward greedy algorithms for atomic norm regularization," *IEEE Trans. Signal Process.*, vol. 63, no. 21, pp. 5798–5811, 2015. doi: 10.1109/TSP.2015.2461515.
- [78] J. Huang, M. Sun, J. Ma, and Y. Chi, "Super-resolution image reconstruction for high-density three-dimensional single-molecule microscopy," *IEEE Trans. Comput. Imag.*, vol. 3, no. 4, pp. 763–773, 2017. doi: 10.1109/TCI.2017.2699425.
- [79] Y. Chen and Y. Chi, "Robust spectral compressed sensing via structured matrix completion," *IEEE Trans. Inf. Theory*, vol. 60, no. 10, pp. 6576–6601, Oct. 2014. doi: 10.1109/TIT.2014.2343623.
- [80] K. H. Jin, D. Lee, and J. C. Ye, "A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix," *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 480–495, 2016. doi: 10.1109/TCI.2016.2601296.
- [81] G. Ongie, S. Biswas, and M. Jacob, "Convex recovery of continuous domain piecewise constant images from nonuniform Fourier samples," *IEEE Trans. Signal Process.*, vol. 66, no. 1, pp. 236–250, doi: 10.1109/TSP.2017.2750111.
- [82] J.-F. Cai, X. Qu, W. Xu, and G.-B. Ye, "Robust recovery of complex exponential signals from random Gaussian projections via low rank Hankel matrix reconstruction," *Appl. Comput. Harmon. Anal.*, vol. 41, no. 2, pp. 470–490, 2016. doi: 10.1016/j.acha.2016.02.003.
- [83] Y. Chi, "Convex relaxations of spectral sparsity for robust super-resolution and line spectrum estimation," in *Proc. Wavelets and Sparsity XVII*, vol. 10394. International Society for Optics and Photonics, 2017, p. 103941G.
- [84] G. Puy, M. E. Davies, and R. Gribonval, "Recipes for stable linear embeddings from Hilbert spaces to  $\mathbb{R}^m$ ," *IEEE Trans. Inf. Theory*, vol. 63, no. 4, pp. 2171–2187, 2017. doi: 10.1109/TIT.2017.2664858.
- [85] B. Adcock and A. C. Hansen, "Generalized sampling and infinite-dimensional compressed sensing," *Found. Comput. Math.*, vol. 16, no. 5, pp. 1263–1323, 2016. doi: 10.1007/s10208-015-9276-6.
- [86] B. Adcock, A. C. Hansen, C. Poon, and B. Roman, "Breaking the coherence barrier: A new theory for compressed sensing," in *Forum of Mathematics, Sigma*, vol. 5, no. e4, 84 pp., 2017. doi: 10.1017/fms.2016.32.
- [87] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, 2002. doi: 10.1109/TSP.2002.1003065.
- [88] L. F. Chamon, Y. C. Eldar, and A. Ribeiro, "Strong duality of sparse functional optimization," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4739–4743. doi: 10.1109/ICASSP.2018.8461311.
- [89] L. F. Chamon, Y. C. Eldar, and A. Ribeiro, "Functional nonlinear sparse models. 2018. [Online]. Available: [arXiv:1811.00577](https://arxiv.org/abs/1811.00577).
- [90] D. Donoho and M. Gavish, "Minimax risk of matrix denoising by singular value thresholding," *Ann. Statist.*, vol. 42, no. 6, pp. 2413–2440, 2014. doi: 10.1214/14-AOS1257.
- [91] C. Boyer, Y. De Castro, and J. Salmon, "Adapting to unknown noise level in sparse deconvolution," *Inform. Inference*, vol. 6, no. 3, pp. 310–348, 2017. doi: 10.1093/imaia/iaw024.
- [92] S. Burer and R. D. Monteiro, "A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization," *Math. Program.*, vol. 95, no. 2, pp. 329–357, 2003. doi: 10.1007/s10107-002-0352-8.
- [93] Y. Chi, Y. M. Lu, and Y. Chen, "Nonconvex optimization meets low-rank matrix factorization: An overview," *IEEE Trans. Signal Process.*, vol. 67, no. 20, pp. 5239–5269, 2019. doi: 10.1109/TSP.2019.2937282.
- [94] J.-J. Fuchs, "Extension of the Pisarenko method to sparse linear arrays," *IEEE Trans. Signal Process.*, vol. 45, no. 10, pp. 2413–2421, 1997. doi: 10.1109/78.640707.
- [95] V. F. Pisarenko, "The retrieval of harmonics from a covariance function," *Geophysical J. Int.*, vol. 33, no. 3, pp. 347–366, 1973. doi: 10.1111/j.1365-246X.1973.tb03424.x.
- [96] S. Li, H. Mansour, and M. B. Wakin, "An optimization view of MUSIC and its extension to missing data. 2018. Available: [arXiv:1806.03511](https://arxiv.org/abs/1806.03511).
- [97] J. Deng, O. Tirkkonen, and C. Studer, "MmWave channel estimation via atomic norm minimization for multi-user hybrid precoding," in *Proc. IEEE Wireless Communications and Networking Conf. (WCNC)*, 2018, pp. 1–6. doi: 10.1109/WCNC.2018.8377093.
- [98] Y. Tsai, L. Zheng, and X. Wang, "Millimeter-wave beamformed full-dimensional MIMO channel estimation based on atomic norm minimization," *IEEE Trans. Commun.*, vol. 66, no. 12, pp. 6150–6163, 2018. doi: 10.1109/TCOMM.2018.2864737.
- [99] Z. Zhu, G. Tang, P. Setlur, S. Gogineni, M. B. Wakin, and M. Rangaswamy, "Super-resolution in SAR imaging: Analysis with the atomic norm," in *Proc. IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2016, pp. 1–5. doi: 10.1109/SAM.2016.7569639.
- [100] J. Ying, H. Lu, Q. Wei, J.-F. Cai, D. Guo, J. Wu, Z. Chen, and X. Qu, "Hankel matrix nuclear norm regularized tensor completion for n-dimensional exponential signals," *IEEE Trans. Signal Process.*, vol. 65, no. 14, pp. 3702–3717, 2017. doi: 10.1109/TSP.2017.2695566.





# Single-Particle Cryo-Electron Microscopy

*Mathematical Theory, Computational Challenges,  
and Opportunities*

Tamir Bendory, Alberto Bartesaghi, and Amit Singer

In recent years, an abundance of new molecular structures have been elucidated using cryo-electron microscopy (cryo-EM), largely due to advances in hardware technology and data processing techniques. Owing to these exciting new developments, cryo-EM was selected by *Nature Methods* as the “Method of the Year 2015,” and the Nobel Prize in Chemistry 2017 was awarded to three pioneers in the cryo-EM field: Jacques Dubochet, Joachim Frank, and Richard Henderson “for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution” [93].

The main goal of this article is to introduce the challenging and exciting computational tasks involved in reconstructing 3D molecular structures by cryo-EM. Determining molecular structures requires a wide range of computational tools in a variety of fields, including signal processing, estimation and detec-

tion theory, high-dimensional statistics, convex and nonconvex optimization, spectral algorithms, dimensionality reduction, and machine learning. The tools from these fields must be adapted to work under exceptionally challenging conditions, including extreme noise levels, the presence of missing data, and massive data sets as large as several terabytes.

In addition, we present two statistical models, multireference alignment (MRA) and multitarget detection (MTD), that abstract away much of the intricacy of cryo-EM while retaining some of its essential features. Based on these abstractions, we discuss some recent intriguing results in the mathematical theory of cryo-EM and delineate relations with group, invariant, and information theories.

## Introduction

Structural biology studies the structure and dynamics of macromolecules to broaden our knowledge about the mechanisms of life and impact the drug-discovery process. Owing to recent

groundbreaking developments, chiefly in hardware technologies and data processing techniques, many new molecular structures have been elucidated to near-atomic resolutions using cryo-EM [8], [44], [47], [52], [81].

In a cryo-EM experiment, biological macromolecules suspended in a liquid solution are rapidly frozen into a thin ice layer. The 3D location and orientation of particles within the ice are random and unknown. An electron beam then passes through the sample, and a 2D tomographic projection, called a *micrograph*, is recorded. The goal is to reconstruct a high-resolution estimate of the 3D electrostatic potential of the molecule (particularly its atomic structure) from a set of micrographs.

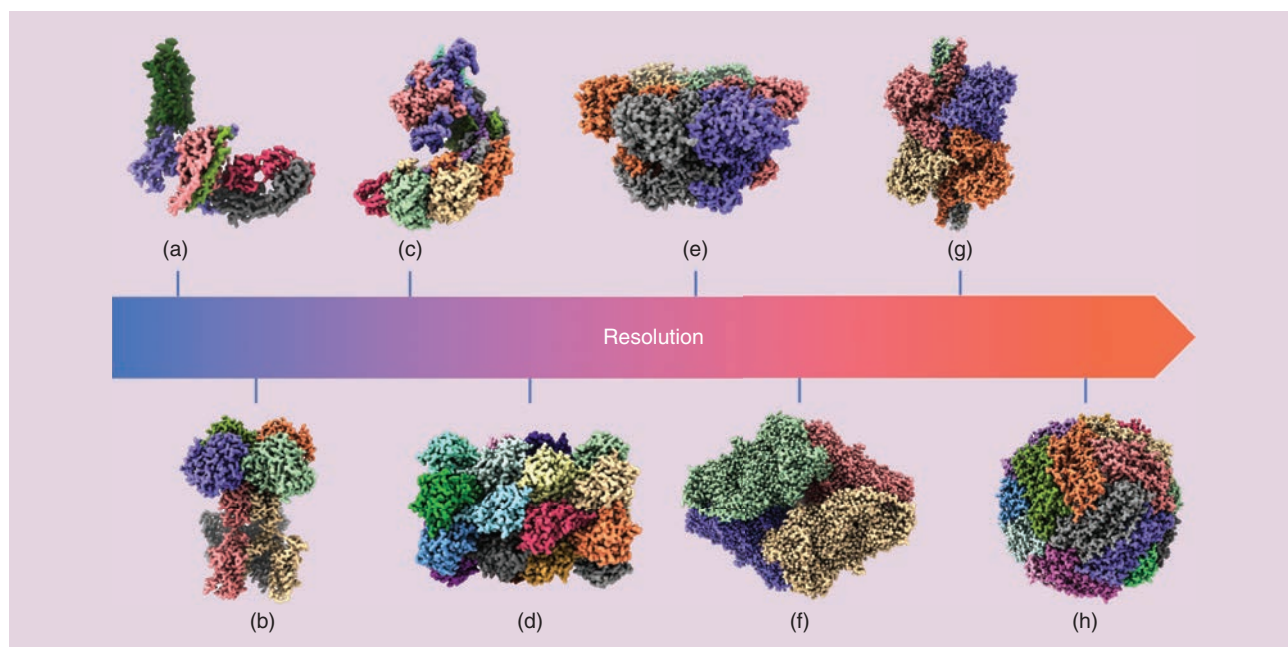
The resolution measures the smallest detail that is distinguishable in a recovered 3D structure; structures with better resolutions resolve finer features. For example, at resolutions of 9 Å,  $\alpha$ -helices are resolved, at resolutions of 4.8 Å, individual  $\beta$ -strands are resolved, and, at resolutions of 3.5 Å, many amino acid side chains are resolved [8]. Figure 1 shows a gallery of important biomedical structures solved by cryo-EM at increasingly higher resolutions. Figure 2 presents an example of a micrograph of the enzyme  $\beta$ -galactosidase and the corresponding high-resolution 3D reconstruction [15].

The signal-to-noise ratio (SNR) of cryo-EM data is very low due to two compounding reasons. On the one hand, the micrograph's contrast is low due to the absence of contrast

enhancement agents, such as heavy-metal stains. On the other hand, the noise level is high because the electron doses must be kept low to prevent damage to the radiation-sensitive biological molecules. The difficulty of estimating the 3D structure in this low-SNR regime, when the orientation and location of the particles are unknown, is the crux of the cryo-EM problem.


Forty years ago, Dubochet and colleagues [27] devised a new technique to preserve biological samples within a thin layer of an amorphous solid form of water called *vitreous ice*. In contrast to “regular” ice, vitreous ice lacks a crystalline molecular arrangement and, therefore, allows preservation of biological structures. As the vitreous ice is maintained at liquid-nitrogen or liquid-helium temperatures, the technique was named *cryo-EM*. In the following decades, the successful application of cryo-EM was limited to the study of large and highly symmetric structures, such as ribosomes and different types of viruses. Before 2013, only a few structures were resolved at resolutions better than 7 Å, and the field was dubbed *blob-ology* due to the blobby appearance of the structures at these resolutions. We refer the reader to [8], [30], [52], and [81] for a more detailed historical account of the development of the technology.

Since 2013, single-particle reconstruction using cryo-EM has been undergoing fast transformations, leading to an abundance of new high-resolution structures and reaching close to



**FIGURE 1.** A gallery of important biomedical structures solved by single-particle cryo-EM at increasing resolutions. (a) The 4.5-Å structure of the human rhodopsin receptor bound to an inhibitory G protein [43], a member of the family of G-protein-coupled receptors, which are the target of approximately 35% of drugs approved by the U.S. Food and Drug Administration. (b) A 3.3-Å map of a voltage-activated potassium channel, an integral membrane protein responsible for potassium ion transport [49]. (c) The 2.9-Å cryo-EM structure of a clustered regularly interspaced short palindromic repeats (CRISPR)/CRISPR-associated protein implicated in gene editing [36]. (d) A 2.4-Å map of the T20s proteasome, a complex that degrades unnecessary or damaged proteins by proteolysis [obtained using data available from Electron Microscopy Public Image Archive (EMPIAR)-10025] [1]. (e) The 2.3-Å structure of human p97 adenosine triphosphate (ATP)/ATPase associated with diverse cellular activities, a key mediator of several protein homeostasis processes and a target for cancer [12]. (f) The 1.9-Å structure of the  $\beta$ -galactosidase enzyme in complex with a cell-permeant inhibitor [15]. (g) The 1.8-Å structure of the conformationally dynamic enzyme glutamate dehydrogenase [50]. (h) A 1.6-Å map of human apoferritin, a critical intracellular iron-storage protein (obtained using data available from EMPIAR-10200).





**FIGURE 2.** Examples of high-resolution cryo-EM imaging of the  $\beta$ -galactosidase enzyme in complex with a cell-permeant inhibitor. (a) A micrograph of  $\beta$ -galactosidase showing individual particle projections (indicated with white circle). (b) The power spectra of the image shown in (a) and estimated CTF matching the characteristic Thon ring oscillations (see the “CTF Estimation and Correction” section). (c) A 1.9-Å resolution map obtained from approximately 150,000 individual particle projections extracted from the publicly available data set EMPIAR-10061 [15], [16]. (d) A close-up view of reconstruction shown in (c), highlighting high-resolution features of the map at the individual amino acid level.

atomic resolution [15]. Figure 3 presents the historical growth of the number of high-resolution structures produced by cryo-EM. Advances in camera technology and data processing contributed to the “resolution revolution” in cryo-EM [44].

First and foremost, a new generation of detectors was developed, called *direct electron detectors (DEDs)*, that—in contrast to charge-coupled device cameras—do not convert electrons to photons. DEDs dramatically improved image quality and SNR, thereby increasing the attainable resolution of cryo-EM. These detectors have high-output frame rates that allow the recording of multiple frames per micrograph (“movies”) rather than the integration of individual exposures [58]. These movies can be used to compensate for motion induced by the electron beam to the sample; see the “Motion Correction” section. In addition, recent hardware developments have enabled

the acquisition and storage of huge amounts of data, which, combined with ready access to CPU and GPU resources, have helped propel the field forward.

The prevalent technique in structural biology in the last half century was X-ray crystallography. This technique suffers from three intrinsic weaknesses that can be mitigated by using cryo-EM imaging. First, many molecules, among them different types of membrane proteins, were notoriously difficult to crystallize. In contrast, the sample preparation procedure for a cryo-EM experiment is significantly simpler, does not require crystallization, and needs smaller amounts of sample. Second, crystal contacts may alter the structure of proteins, making it difficult to recover their physiologically relevant conformation. Cryo-EM samples, instead, are rapidly frozen into vitreous ice, which preserves the molecules in a near-physiological

state. Third, the X-ray beam aggregates the information from all molecules simultaneously and, thus, hinders visualization of structural variability.

In stark contrast, the projection of each particle in a cryo-EM experiment is recorded independently, and, thus, multiple structures—associated with different functional states—can be estimated [74]. A third technique, nuclear magnetic resonance (NMR), can be used to elucidate molecular structures in physiological conditions (water solution at room temperature) but is restricted to small structures of up to approximately 50 kDa.

The main goal of this article is to introduce the unique and exciting algorithmic challenges and cutting-edge mathematical problems arising from cryo-EM research. The estimation of molecular structures involves developing and adopting computational tools in signal processing, estimation and detection theory, high-dimensional statistics, convex and nonconvex optimization, spectral algorithms, dimensionality reduction, and machine learning as well as knowledge in group theory, invariant theory, and information theory.


All tools from the aforementioned fields should be adapted to exceptional conditions: an extremely low-SNR environment and the presence of missing data (for instance, 2D location and 3D orientation of samples in the micrograph). In addition, the devised algorithms should be efficient when run on massively large data sets on the order of several terabytes. This article provides an account of the leading software packages in the field and discusses their underlying mathematical, statistical, and algorithmic principles [33], [57], [63], [76].

Before moving on, we want to mention that topics of great importance to practitioners, such as the physics and optics of an electron microscope, sample preparation, and data acquisition, are not discussed in this article. These topics are thoroughly covered by biologically oriented surveys, such as [30], [52], and [81] and the references therein.

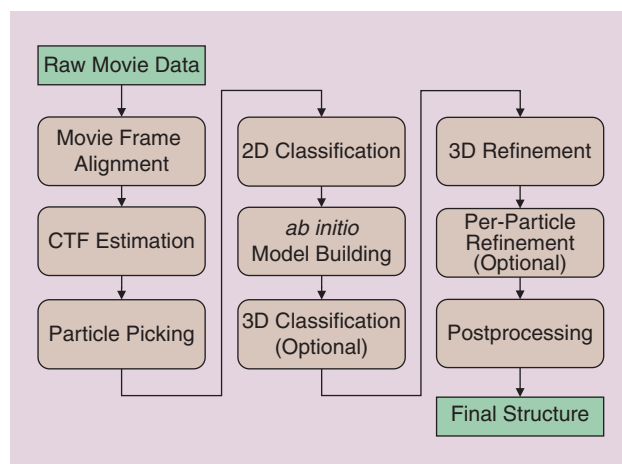
## The cryo-EM reconstruction problem

Modern electron microscopes produce multiple micrographs, each composed of a series of frames (a “movie”). The first stage of any contemporary algorithmic pipeline is to align and average the frames to mitigate the effects of movement induced by the electron beam and, thus, to improve the SNR. This process is called *motion correction* or *movie frame alignment*. The next step, termed *particle picking*, consists of detecting and extracting the particles’ tomographic projections from the micrographs.

Perfect detection requires finding the particles’ center of mass, which is difficult to estimate due to the characteristics of the Fourier transform of the microscope’s point spread function (PSF), called the *contrast transfer function (CTF)*; see the “CTF Estimation and Correction” section. The output of the particle-picking stage is a series of images  $I_1, \dots, I_N$  from which we wish to estimate the 3D structure; this section focuses on the problem of 3D reconstruction from picked particles, which is the heart of the computational pipeline of single-particle reconstruction using cryo-EM. Motion correction and



**FIGURE 3.** The recent growth in the number of high-resolution structures produced by single-particle cryo-EM. (a) The cumulative number of structures solved by single-particle cryo-EM in the last 17 years, as recorded in the Electron Microscopy Data Bank public repository [2]. (b) The corresponding values for the average resolution of maps deposited in the database showing an inflection point after the year 2013, coinciding with the introduction of DED technology in the cryo-EM field.



**FIGURE 4.** A flowchart diagram showing the computational pipeline required to convert raw movie data into high-resolution structures by single-particle cryo-EM. (Adapted from [90].)

particle picking are discussed in more detail in the “Building Blocks in the Computational Pipeline” section.

Figure 4 shows the complete cryo-EM imaging pipeline. Briefly, raw data are first preprocessed at the movie frame alignment (“Motion Correction” section) and CTF estimation (“CTF Estimation and Correction” section) steps, followed by particle picking to detect and extract the individual projections from micrographs. Occasionally, particle picking is followed by a particle-pruning stage to remove noninformative picked particles (“Particle Picking” section). The output of this stage is a set of 2D images; each (ideally) contains a noisy tomographic projection taken from an unknown viewing direction.



Particle images are then classified (“2D Classification” section): the 2D classes may be used to construct *ab initio* models, as templates for particle picking, to provide a quick assessment of the particles and for symmetry detection. Then, an *ab initio* model is built using the 2D classes or by alternative techniques (“Ab Initio Modeling” section). If the data contain structural variability or a mix of structures (see the “The Cryo-EM Reconstruction Problem” and “Conformational Heterogeneity: Modeling and Recovery” sections), then a 3D classification step is applied to cluster the projection images into the different structural conformations.

Initial models are subjected to 3D high-resolution refinement (“High-Resolution Refinement” section), and an additional per-particle refinement may be applied. Finally, a postprocessing stage is employed to facilitate interpretation of structures in terms of atomic models. Different software packages may use slightly different workflows and, occasionally, some of the steps are applied iteratively. For instance, one can use the 2D classes to repeat particle picking with more reliable templates.

Let  $\phi: \mathbb{R}^3 \rightarrow \mathbb{R}$  represent the 3D molecular structure to be estimated. Under the assumption that the particle picking is executed well (i.e., each particle is detected up to a small translation; we assume for simplicity no false detection), each image  $I_1, \dots, I_N$  is formed by rotating  $\phi$  by a 3D rotation  $R_{\omega_i}$ , integrating along the  $z$ -axis (tomographic projection), shifting by a 2D shift  $T_i$ , convolving with the PSF of the microscope  $h_i$ , and adding noise:

$$I_i = h_i * T_i \left( \int_{-\infty}^{\infty} \phi(R_{\omega_i}^T r) dz \right) + \text{“noise,”} \quad r = (x, y, z),$$

$$i = 1, \dots, N, \quad (1)$$

where  $*$  denotes convolution. The model can be written more concisely as

$$I_i = h_i * T_i P R_{\omega_i} \phi + \text{“noise,”} \quad i = 1, \dots, N, \quad (2)$$

where the tomographic projection is denoted by  $P$ , and  $(R\phi)(r) = \phi(R^T r)$ . The image is then sampled on a Cartesian grid. The integration along the  $z$ -axis is called the *X-ray transform* (not to be confused with the radon transform in which the integration is over hyperplanes). The rotations  $R_{\omega}$  describe the unknown 3D orientation of the particles embedded in the ice, and they can be thought of as random elements of the special orthogonal group  $\omega \in SO(3)$ . These rotations can be represented as  $3 \times 3$  orthogonal matrices with determinant one or by quaternions. The 2D shifts  $t \in \mathbb{R}^2$  result from detection inaccuracies, which are usually small. The PSFs  $h_i$  are assumed to be known, and their Fourier transforms suffer from many zero crossings, making deconvolution challenging; see the “CTF Estimation and Correction” section for further discussion.

The cryo-EM inverse problem of reconstruction consists of estimating the 3D structure  $\phi$  from the 2D images  $I_1, \dots, I_N$ . Importantly, the 3D rotations  $\omega_1, \dots, \omega_N$  and the 2D shifts  $t_1, \dots, t_N$  are called *nuisance variables*; although the rotations and shifts are unknown a priori, their estimation

is not an aim by itself. Figure 5 shows an example of the cryo-EM problem in its most simplified version, without noise, CTF, and shifts.

The reconstruction of  $\phi$  is possible with up to three intrinsic ambiguities: a global 3D rotation, the position of the center of the molecule (3D location), and handedness. This last symmetry, also called *chirality*, means that it is impossible to distinguish whether the molecule was reflected about a 2D plane through the origin. The handedness of the structure cannot be determined from cryo-EM images alone because the original 3D object and its reflection give rise to identical sets of projections related by the following conjugation:  $\tilde{R}_{\omega_i} = J R_{\omega_i} J^{-1}$ , where  $J = \text{diag}(1, 1, -1)$ .

In the presence of structural variability,  $\phi$  may be thought of as a random signal with an unknown distribution defined over a space of possible structures (which might be unknown as well). In this case, the task is more ambitious and involves estimating the whole distribution of conformations. Usually, the distribution is assumed to be discrete (i.e., in each measurement, we observe one of a few possible structures or conformational states) or to lie in a low-dimensional subspace or manifold. This subject is further discussed in the “Conformational Heterogeneity: Modeling and Recovery” section and in a recent survey [74].

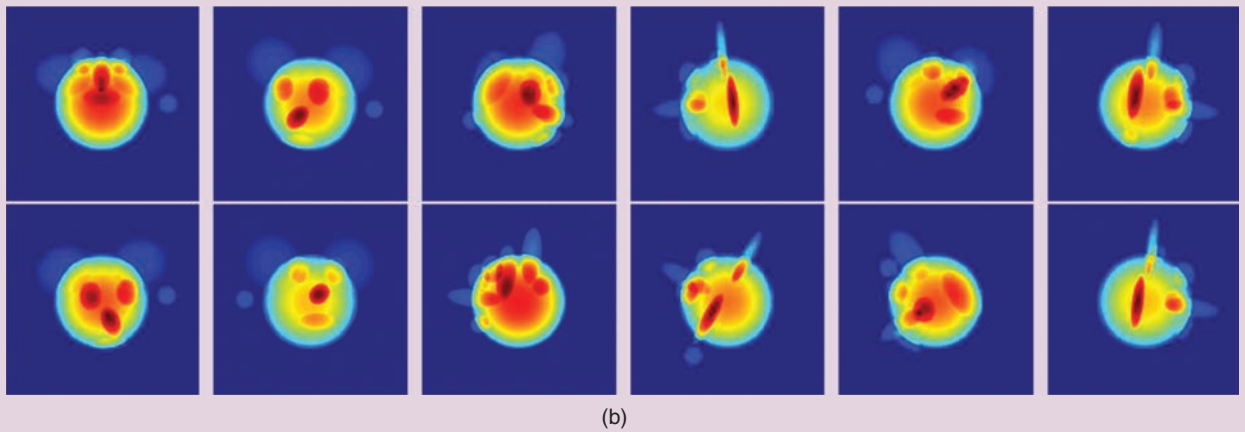

The chief noise source in cryo-EM at the frame level (before motion correction) is shot noise, which follows a Poisson distribution. After movie frames are averaged to produce micrographs, it is customary to assume that the noise is characterized by a Gaussian distribution. Indeed, all current algorithms build on—implicitly or explicitly—the speculated Gaussianity of the noise. While the spectrum of the noise is not white—mainly due to inelastic scattering, variations in the thickness of ice, and the PSF of the microscope—it is assumed to be a 1D radial function (i.e., constant along the angular direction). The common practice is to estimate the parameters of the noise power spectrum during the 3D reconstruction process or from regions in the micrographs that, presumably, contain no signal.

## Main computational challenges

The difficulty in determining high-resolution molecular structures using cryo-EM hinges on three characteristic features of the cryo-EM data: high noise level, missing data, and its massive size. This section elaborates on these unique features, while the next sections dive into the different tasks and algorithms involved in cryo-EM data processing.

### High noise level

In a cryo-EM experiment, the electron doses must be kept low to mitigate radiation damage due to electron illumination. The low doses induce high noise levels on the acquired raw data frames. Together with the image’s low contrast, this results in SNR levels that are usually well below 0 dB and might be as low as −20 dB (i.e., the power of the noise is 100 times greater than the power of the signal). Under such low-SNR conditions, standard tasks, such as aligning, detecting, or clustering signals, become very challenging.



**FIGURE 5.** (a) A simulated 3D structure and (b) a dozen of its noise-free tomographic projections from different viewing directions. The most simplified version of the cryo-EM problem is estimating the 3D structure from the 2D projection images when the viewing directions are unknown. In practice, the projections are highly noisy, slightly shifted, and convolved with the microscope's PSF. (Used with permission from [71].)

To comprehend the difficulty of performing estimation tasks in a low-SNR environment, let us consider a simple rotation-estimation problem. Let us denote by  $x \in \mathbb{R}^L$  the samples of a periodic signal on the unit circle;  $x$  is assumed to be known. We wish to estimate a rotation  $\theta \in [0, 2\pi)$  from a noisy measurement

$$y = R_\theta x + \varepsilon, \quad (3)$$

where  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$ , and  $R_\theta$  denotes the rotation operator, that is,  $(R_\theta x)(t) = x(t - \theta)$ . To estimate  $\theta$ , we correlate the signal  $x$  with rotated versions of  $y$  and choose the inverse of the rotation that maximizes the correlation as the estimator; this technique is called *template matching*. In the absence of noise, template matching simply correlates the signal with its rotated versions; the maximum is attained when  $y$  is rotated by  $-\theta$ .

However, in the presence of noise, we get an additional stochastic term due to the correlation of the noise with the signal. Consequently, if  $\sigma^2$  is large, it is likely that the peak of the correlation will not be close to  $-\theta$ . In particular, when  $\sigma \rightarrow \infty$ , the location of the peak is distributed uniformly on the circle. This result can be derived formally using the Neyman–Pear-

son lemma and holds true for any estimation technique, not necessarily template matching [18]. Even if we collect  $N$  measurements (each with a different rotation), it is impossible to estimate the  $N$  rotations accurately. In [6], it was shown that the Cramér–Rao bound of this problem is proportional to  $\sigma^2$  and independent of  $N$ . Therefore, if the noise level is high, the variance of any estimator will be high as well.

The same conclusions that were derived for the simple rotation estimation problem (3) remain true for cryo-EM. For example, even if the 3D structure is known, aligning a noisy raw image against multiple noiseless projection templates that correspond to rotations sampled from  $SO(3)$  will produce no salient peak in the correlation if the noise level is high. In particular, the higher the noise level, the flatter the distribution over  $SO(3)$ .

#### Missing data—Unknown viewing directions and locations

The viewing direction and location associated with each particle in a micrograph are unknown a priori. If they were known, estimating the structure  $\phi$  would be a linear inverse problem, similar to the reconstruction problem in computerized tomography

(CT). The recovery in this case is based on the Fourier slice theorem, which states that the 2D Fourier transform of a tomographic projection is the restriction of the 3D Fourier transform of  $\phi$  to a 2D plane. Mathematically, it can be written succinctly as

$$SRF_3\phi = F_2PR\phi, \quad (4)$$

where  $F_2$ ,  $F_3$ ,  $P$ , and  $S$  are, respectively, the 2D Fourier, 3D Fourier, tomographic projection, and restriction operators; this theorem is a direct corollary of the fact that the Fourier operator  $F$  and the rotation operator  $R$  commute, that is,  $RF = FR$ .

The Fourier slice theorem implies that acquiring tomographic projections from known viewing directions is equivalent to sampling the 3D Fourier space. Hence, given enough projections, one can estimate  $\phi$  to a certain resolution. This is the underlying principle behind many CT imaging algorithms, such as the classical filtered-back projection (FBP). However, FBP is not optimal for cryo-EM (even if viewing directions are known) since not all viewing directions are necessarily represented in the data, and the sampling in the Fourier domain is nonuniform; the coverage of viewing directions affects the quality of the solution.

An alternative is the algebraic reconstruction technique (ART), which solves the linear inverse problem by iterative projections; this algorithm is called the *Kaczmarz method* in numerical linear algebra. However, ART is rarely used in cryo-EM because it is slow: it does not exploit the Fourier slice theorem and, thus, cannot be accelerated using fast Fourier transforms (FFTs). Modern approaches were developed to exploit the Fourier slice theorem and account for nonuniform sampling; popular algorithms are based on efficient gridding methods (that compute a uniformly sampled version of a function from a nonuniformly sampled version by choosing proper weights) [55] and using nonuniform FFT packages to harness the structure of the linear system [14], [35].

While the viewing directions in cryo-EM are unknown, there is a rigorous technique to estimate them based on common lines; see the “Estimating Viewing Directions Using Common Lines” section as well as [72] and [79]. Unfortunately, any method for estimating the viewing directions is destined to fail when the SNR is low, for the same reasons that estimating  $\theta$  in (3) would fail. Bearing in mind that the ultimate goal is to estimate the 3D structure—not the viewing directions—it is essential to consider statistical methods that circumvent rotation estimation, such as the maximum likelihood (ML) and the method of moments.

### Massive data sets and high dimensionality

A single session of data collection in a typical cryo-EM experiment produces a few thousand micrographs, each containing several hundred individual particle projections. Depending on the type of detector used during acquisition, micrographs can range in size from a few tens of megapixels up to 100 megapixels for the newest detectors. Moreover, the new generation of detectors can record each micrograph as a rapid burst of frames, producing large movie files that result in data sets of

several terabytes in size. The sheer volume of data must be taken into account early on in the algorithmic design process, and in addition to storage considerations, steps must be taken to ensure the efficient use of computational resources to keep up with the ever-increasing throughput of data produced by modern cameras.

Another issue is the dimensionality of the reconstruction problem. The number of voxels of a typical  $200 \times 200 \times 200$  volume is 8 million. Estimating so many parameters poses a challenge, both from the computational complexity side (for example, how to find the maximum of the likelihood function of 8 million variables) and also on the statistical-estimation front (see the “Denoising and Dimensionality Reduction Techniques” section). The problem is even more severe when multiple structures, or even a continuum of structures, need to be estimated (see the “Remaining Computational and Theoretical Challenges” section). Getting to high resolution is a major bottleneck due to a compounding computational burden effect: as more parameters need to be estimated, more data are required for their estimation, and the computation becomes ever more expensive and challenging.

## 3D reconstruction from projections

### High-resolution refinement

The reconstruction procedure of high-resolution 3D structures is usually split into two stages: constructing an initial low-resolution model, which is later refined by applying an iterative algorithm (“refinement algorithm”); see Figure 4. This section is devoted to refinement techniques, while different approaches to constitute low-resolution estimates using *ab initio* modeling—initialization-free models—are discussed in the next section.

Refinement techniques for cryo-EM can be broadly classified into two categories: hard and soft angular-assignment methods. The hard-assignment approach is based on template matching. At each iteration, multiple projections are generated from the current estimate of the structure; the projections should, ideally, densely cover  $SO(3)$ . Then, a single viewing direction is assigned to each experimental image based on the projection with which it correlates best. The angular assignment can be performed either in real or Fourier space.

The advantage of working in Fourier space is that it is not necessary to rotate the molecule: projections can be computed fast using off-the-shelf nonuniform FFT packages (e.g., [14] and [35]), as implied by the Fourier slice theorem. Moreover, with this representation, the CTF is simply a diagonal operator. Once the viewing directions of all experimental images are assigned, the 3D structure is constructed using standard linear-inversion techniques; see the discussion in the “Main Computational Challenges” section. The algorithm iterates between hard angular assignment and structure construction until convergence. Although the quality of hard angular assignments may be influenced by the high noise levels, there are several examples of packages that follow this type of approach, including EMAN2 [76] and cisTEM [33], among others.

A second class of strategies comprises the soft-assignment methods. Similar to hard-assignment methods, in each iteration, the experimental images are correlated with multiple template projections. However, instead the best match being chosen, a similarity score is given to each pair of experimental image and projection. These scores, also called *weights* or *responsibilities*, are computed according to the generative statistical model of the experimental images and used for reconstruction, as explained later in this section.

The soft assignment methods are instances or variants of the well-known expectation-maximization (ML-EM) algorithm [25]. Notably, this relation classifies the 3D reconstruction problem as a problem of ML estimation or, more generally, a problem in Bayesian statistics and, thus, provides a solid theoretical and algorithmic framework. In particular, it enables the incorporation of priors that, essentially, act as regularizers in the reconstruction process. In the context of cryo-EM, ML-EM was first applied to 2D cryo-EM images by Sigworth [68] and was later implemented for 3D reconstruction by the software packages RELION [63] and cryoSPARC [57].

In what follows, we describe the ML-EM algorithm for a more general statistical model and then particularize it to the special case of cryo-EM. Suppose we collect  $N$  observations from the model:

$$y_i = L_{\theta_i} x + \varepsilon_i, \quad i = 1, \dots, N, \quad (5)$$

where  $L_{\theta}$  is a linear transformation acting on the signal  $x$ , parameterized by a random variable  $\theta$ , and  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$ . The goal is to estimate  $x$ , while  $\theta_1, \dots, \theta_N$  are nuisance variables. A typical assumption is that  $x$  can be represented by a finite number of coefficients (for instance, its Fourier expansion is finite) and that these coefficients were drawn from a Gaussian distribution with mean  $\mu$  and covariance matrix  $\Sigma$ . The key for reliable estimation is to marginalize over the nuisance variables  $\theta_1, \dots, \theta_N$ . Without marginalization (i.e., when the goal is to jointly estimate  $x$  and  $\theta_1, \dots, \theta_N$ ), the number of parameters grows with the number of measurements indefinitely, and, thus, the ML estimator may be inconsistent. This phenomenon is exemplified by the Neyman–Scott “paradox.”

In cryo-EM, the transformation  $L_{\theta}$  is the operator described in (2). This operator rotates the volume, computes its 2D tomographic projection, and applies a 2D translation and convolution with the PSF;  $\theta$  is drawn from a distribution defined over the 5D space of 3D rotations and 2D translations. The distribution of  $\theta$  is generally unknown and should be estimated as part of the ML-EM algorithm.

Let us denote  $\mathbf{y} = (y_1, \dots, y_N)$ . The posterior distribution  $p(x|\mathbf{y})$  of (5) is proportional to the product of the prior  $p(x)$  with the likelihood function

$$\begin{aligned} \mathcal{L}(x; \mathbf{y}) &= p(\mathbf{y}|x) = \prod_{i=1}^N \sum_{\theta_i \in \Theta} p(y_i|x, \theta_i) p(\theta_i) \\ &= \frac{1}{(2\pi\sigma^2)^{M/2}} \prod_{i=1}^N \sum_{\theta_i \in \Theta} p(\theta_i) e^{-\frac{1}{2\sigma^2} \|y_i - L_{\theta_i} x\|^2}, \end{aligned} \quad (6)$$

where  $\Theta$  is a discrete space on which  $\theta$  is defined, and  $M$  is the length of each observation. If the entire posterior distribution could be computed, then one could obtain far more statistical information about  $x$  than just maximizing the likelihood. For instance, the best estimator in the minimum mean square error (MMSE) sense is obtained by marginalizing over the posterior:

$$\hat{x}_{\text{MMSE}}(\mathbf{y}) = \int_x p(x|\mathbf{y}) x dx = \mathbb{E}\{x|\mathbf{y}\}. \quad (7)$$

Unfortunately, the entire posterior can rarely be computed, especially in big data problems, such as cryo-EM. The ML-EM framework provides a simple, yet frequently very effective, iterative method that tries to compute the maximum a posteriori estimator (MAP), that is, the maximal value of  $p(x|\mathbf{y})$ .

Each iteration of the ML-EM algorithm consists of two steps. The first step (E-step) computes the expected value of the log of the posterior with respect to the nuisance variables, conditioned on the current estimate of  $x$ , denoted here by  $x_t$ , and the data

$$\begin{aligned} Q(x|x_t) &= \mathbb{E}_{\theta|\mathbf{y}, x_t} \{\log p(x|\mathbf{y}, \theta)\} \\ &= \text{constant} - \frac{1}{2\sigma^2} \sum_{i=1}^N \sum_{\theta_i \in \Theta} w_{i,\ell} \|y_i - L_{\theta_i} x\|^2 + \log p(x), \end{aligned} \quad (8)$$

where

$$w_{i,\ell} = p(\theta_i = \theta_{\ell} | y_i, x_t) = \frac{e^{-\frac{1}{2\sigma^2} \|y_i - L_{\theta_{\ell}} x_t\|^2}}{\sum_{\theta_i \in \Theta} e^{-\frac{1}{2\sigma^2} \|y_i - L_{\theta_i} x_t\|^2}}. \quad (9)$$

If  $x \sim \mathcal{N}(0, \Sigma)$ , then  $\log p(x) = -(1/2)x^T \Sigma^{-1}x$ .

The second step (M-step) updates  $x$  by

$$x_{t+1} = \arg\max_x Q(x|x_t), \quad (10)$$

which is usually performed by setting the gradient of  $Q$  to zero. If  $x$  is assumed to be Gaussian, then the M-step reduces to solving a linear system of equations with respect to  $x$ . The E- and M-steps are applied iteratively. It is well known that  $p(x_{t+1}|\mathbf{y}) \geq p(x_t|\mathbf{y})$ ; however, since the landscape of  $p(x|\mathbf{y})$  is usually nonconvex, the iterations are not guaranteed to converge to the MAP estimator [25]. Usually, the ML-EM iterations halt when  $(p(x_{t+1}|\mathbf{y}) - p(x_t|\mathbf{y}))/p(x_t|\mathbf{y})$  is smaller than some tolerance (but other criteria can be employed as well). The posterior distribution at each iteration can be evaluated according to (6).

The implementation details of the ML-EM algorithm for cryo-EM vary across different software packages. The popular package RELION, for example, incorporates a prior of uniform distribution of rotations over  $SO(3)$  (although the distribution itself is usually nonuniform), and each Fourier coefficient of the 3D volume was drawn independently from a normal distribution [63]. The variance of the Fourier coefficients' prior is updated at each iteration by averaging over



concentric frequency shells of the current estimate (resulting in a 1D radial function). If the ML-EM algorithm is initialized with a smooth volume (which is the common practice), the variance is large only at low frequencies; hence, high frequencies are severely regularized. As a result, the first ML-EM iterations mostly update the low frequencies, while the high frequencies are only slightly affected.

As the algorithm proceeds, the structure includes more high-frequency content; thus, the variance of these frequency shells increases, and the effective resolution of the 3D map gradually improves. This strategy of initializing with a low-resolution structure and gradually increasing the resolution is called *frequency marching* and is shared by many packages in the field; see further discussion in the “Frequency Marching” section. Another popular package, cryoSPARC, also employs an ML-EM algorithm with the prior that each voxel in real space was drawn independently from a Poisson distribution with a constant parameter [57].

We mention that such statistical priors (e.g., each voxel or frequency is drawn from Poisson or normal distributions) are usually chosen out of mathematical and computational convenience and may bias the reconstruction process. It is an open challenge to replace those statistical priors by priors that are based on or inspired by biological knowledge, that is, previously reconstructed structures.

The main drawback of the ML-EM approach, especially at high resolution, is the high computational load of the E-step in each iteration. Specifically, each ML-EM iteration requires correlating each experimental image (typically a few hundred thousand) with multiple synthetic projections of the current structure estimate [sampled densely over  $SO(3)$ ] and their 2D translations. To alleviate the computational burden, a variety of methods are employed to narrow the 5D search space. Popular techniques are based on multiscale approaches, whereby an initial search is done on a coarse grid, followed by local searches on finer grids. A more sophisticated idea was suggested in [57], based on the “branch-and-bound” methodology, that rules out regions of the search space that are not likely to contain the optimum of the objective function.

### Ab initio modeling

In this section, we describe some of the intriguing ideas that were proposed to construct *ab initio* models, that is, models that do not require an initial guess. These methods usually result in low-resolution estimates that can be later refined as described in the previous section. To emphasize the importance of robust *ab initio* techniques, we begin by discussing model bias—a phenomenon of crucial importance in cryo-EM and statistics in general.

#### Model bias and validation

The cryo-EM inverse problem is nonconvex and, thus, the output of 3D reconstruction algorithms may depend on their initializations. This raises the validation problem: how can we verify that a given estimate is a faithful representation of the underlying data? Currently, a 3D model is treated as valid if

its structural features meet the common biological knowledge (primary structure, secondary structure, and so on) and if it passes some computational tests based on different heuristics. For instance, the reconstruction algorithm can be initialized from many different points. If the algorithm always converges to the same or similar structures, it strengthens the confidence in the attained solution. Moreover, since the data are usually uploaded to public repositories, different researchers can examine it and compare the results against each other [61]; see a more detailed discussion in the “Verification” section.

Despite all precautions taken by researchers in the field, the verification methods are not immunized against systematic errors—this pitfall is dubbed *model bias*. One important example concerns the particle-picking stage, when one aims to detect and extract particle projections from noisy micrographs. The majority of the detection algorithms in the field are based on template-matching techniques, despite their intrinsic flaws. Specifically, choosing improper templates can lead to erroneous detection, which, in turn, biases the 3D reconstruction algorithm toward the chosen templates; see [39] and further discussion on particle-picking techniques in the “Particle Picking” section.

The model bias phenomenon is exemplified by the “Einstein from noise” experiment [67]. In this experiment,  $N$  images of pure independently identically distributed (i.i.d.) Gaussian noise are correlated with a reference image. (In the original article, the authors chose an image of Einstein as the reference and, thus, the name.) Then, each pure-noise image is shifted to best align with the reference image (based on the peak of their cross correlation), and, finally, all images are averaged.

Without bias, one would expect that averaging pure-noise images would converge toward an image of zeros as  $N$  diverges. However, in practice, the resulting image is similar to the reference image, that is, the algorithm is biased toward the reference image; see [67, Fig. 2]. In the context of cryo-EM, the lesson is that, without prudent algorithmic design, the reconstructed molecular structure may reflect the scientist’s pre-sumptions rather than the structure that best explains the data. We now turn our attention to some of the existing techniques for *ab initio* modeling.

#### Stochastic gradient descent over the likelihood function

Stochastic gradient descent (SGD) has gained popularity in recent years, especially due to its invaluable role in the field of deep learning [21]. The underlying idea is very simple. Suppose that we wish to minimize an objective function of the form

$$f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x). \quad (11)$$

A gradient descent algorithm is an optimization technique used to minimize the objective function by iteratively moving in the direction opposite to its gradient. Its  $t$ th iteration takes on the form

$$x_{t+1} = x_t - \eta_t \frac{1}{N} \sum_{i=1}^N \nabla f_i(x_t), \quad (12)$$

where  $\nabla$  denotes a gradient and  $\eta_t$  is called the *step size* or *learning rate*. In SGD, we simply replace the full sum by a random element:

$$x_{t+1} = x_t - \eta_t \nabla f_i(x_t), \quad (13)$$

where  $i$  is chosen uniformly at random from  $\{1, \dots, N\}$ . In expectation, the direction of each SGD step is the direction of the full gradient (12). The main advantage of SGD is that each iteration is significantly cheaper to compute than the full gradient (12). It is also easy to modify the algorithm to sum over a few random elements at each iteration, rather than just one, to improve robustness at the cost of higher computational load per iteration. In addition, it is commonly believed that the inherent randomness of SGD helps to escape local minima in some nonconvex problems.

This strategy was proven to be effective for constructing *ab initio* models using cryo-EM, achieving resolutions of up to 10 Å [57]. The objective function is the negative log of the posterior distribution. As the log posterior involves a sum over the experimental images (6), the SGD scheme chooses, at each iteration, one random image (or a subset of images) to approximate the gradient direction.

#### Estimating viewing directions using common lines

If the orientations and locations of the particles in the ice are known [equivalently, the 3D rotations and 2D translations in (2)], then recovering the 3D structure is a linear problem that can be solved using one of the many solutions developed for CT imaging. In this section, we survey an analytical method for estimating the viewing directions using the common-lines property.

Due to the Fourier slice theorem, any pair of projection images has a pair of central lines (one in each image) on which their Fourier transforms agree. For generic molecular structures (with no symmetry), it is possible to uniquely identify this common line, for example, by cross-correlating all possible central lines in one image with all possible central lines in the other image and choosing the pair of lines with maximum cross correlation.

The common line pins down two out of the three Euler angles associated with the relative rotation  $R_i^{-1}R_j$  between images  $I_i$  and  $I_j$ . A third image is required to determine the third angle: the three common-line pairs between the three images uniquely determine their relative rotations. This technique is called *angular reconstitution*, and it was suggested, independently, by Vainshtein and Goncharov [78] and Van Heel [79]. The main drawback of this procedure is its sensitivity to noise; it requires the three pairs of common lines to be accurately identified. Moreover, estimating the rotations of additional images sequentially (using their common lines with the previously rotationally assigned images) can quickly accumulate errors.

As a robust alternative, it was proposed that the rotations  $R_1, \dots, R_N$  be estimated from the common lines of all pairs  $I_i, I_j$  simultaneously; this framework is called *synchronization*. Since this strategy takes all pairwise information into

account, it has better tolerance to noise. The rotation assignment can be done using a spectral algorithm or semidefinite programming [72] and enjoys some theoretical guarantees; see, for instance, [70]. Nevertheless, as discussed in the “Main Computational Challenges” section, in a low-SNR regime, any method for estimating the rotations would fail. Therefore, the method requires as input high-SNR images that can be obtained using a procedure called *2D classification*; see further discussion in the “2D Classification” section. However, 2D classification blurs the fine details of the images, and, therefore, the attained resolution of the method is limited [34].

Three last comments are in order. First, if the structure possesses a nontrivial symmetry, there are multiple common lines between pairs of images that should be considered [56]. Second, interestingly, this technique cannot work when the underlying object is a 2D image (rather than 3D, as in cryo-EM). In this case, the Fourier transform of the tomographic projection (which is a 1D function in that case) is a line that goes through the origin of the 2D Fourier space of the underlying object. Therefore, the single common point of any two projections taken from different angles is the origin, and, thus, there is no way to find the relative angle between them. Finally, the common-lines method is effective for additional signal processing applications, such as the study of specimen populations [46].

#### The method of moments

Suppose that a set of parameters  $x$  characterizes a distribution  $p(y|x)$  of a random variable  $y$  with  $L$  entries. We observe  $N$  i.i.d. samples of  $y$  and wish to estimate  $x$ . In the method of moments, the underlying idea is to relate the moments of the observed data with  $x$ . Those moments can be estimated from the data by averaging over the observations

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N y_i &\approx \mathbb{E}\{p(y|x)\} = M_1(x) \\ \frac{1}{N} \sum_{i=1}^N y_i y_i^T &\approx \mathbb{E}\{p(y y^T | x)\} = M_2(x) \\ &\vdots \\ \frac{1}{N} \sum_{i=1}^N y_i^{\otimes k} &\approx \mathbb{E}\{p(y^{\otimes k} | x)\} = M_k(x), \end{aligned} \quad (14)$$

where  $y^{\otimes k}$  is a tensor with  $L^k$  entries, and the entry indexed by  $n = (n_1, \dots, n_k) \in \mathbb{Z}_L^k$  is given by  $\prod_{i=1}^k y(n_i)$ . By the law of large numbers, when  $N \rightarrow \infty$ , the average converges almost surely to the expectation. The right-hand side underscores that the expectations are solely functions of  $x$ ; the moment tensors  $M_1, M_2, \dots, M_k$  can be derived analytically. For instance, in cryo-EM, the expectation is taken against the random rotations, translations, and noise; the moments are functions of the 3D volume, the 5D distribution over the space of 3D rotations and 2D translations, and some parameters of the noise statistics (due to bias terms). The final step of the method of moments is to solve the system of equations (14), which might be highly nontrivial. Importantly, since estimating the moments requires only one pass over the observations, it can be done potentially

on the fly during data acquisition. This is especially important for cryo-EM, in which hundreds of thousands of images must be processed.

Computing the  $q$ th moment involves the product of  $q$  noisy terms and, thus, the variance of its estimation scales as  $\sigma^{2q}/N$ . As a result, when invoking the method of moments, it is crucial to determine the lowest-order moment that identifies the parameters  $x$  uniquely: this moment determines the *estimation rate* of the problem, that is, how many samples are required to accurately estimate the parameters. Remarkably, it was shown that under a general statistical model that is tightly related to the cryo-EM problem (as will be described in the “MRA” section), the first moment that distinguishes between different parameters determines the sample complexity of the problem (i.e., how many observations are necessary to attain an accurate estimate) in the low-SNR regime [4]. In addition, the number of equations in (14) increases with the order of moments. Thus, using lower-order moments reduces the computational complexity.

The first to suggest the method of moments for single particle reconstruction, 40 years ago, was Zvi Kam [42]. Under the simplifying assumptions that the particles are centered, the rotations are uniformly distributed, and, by ignoring the CTF, he showed that the second moment of the experimental images determines the 3D structure, up to a set of orthogonal matrices. To determine those matrices, he suggested computing a subset of the third-order moment. Recently, it was shown that under the same assumptions, the structure is indeed determined uniquely from the third moment [9]. Remarkably, if the distribution of rotations is nonuniform, then the second moment suffices [66]. Therefore, in light of [4], the sample complexity of the cryo-EM problem in the low-SNR regime, under the specified conditions, scales as  $\sigma^4$  and  $\sigma^6$  for nonuniform and uniform distribution of rotations, respectively. It is still an open question as to how many images are required to solve the full cryo-EM problem; see further discussion in the “Remaining Computational and Theoretical Challenges” section.

Interest in the method of moments has been recently revived, largely due to its potential application to X-ray free-electron lasers [48], [82]. We believe that the method of moments has been largely overlooked by the cryo-EM community since Kam’s article for three main reasons. First, Kam’s original formulation required uniform distribution of viewing directions, an assumption that typically does not hold in practice; a recent article extends the method to any distribution over  $SO(3)$  [66]. Second, estimating the second- and third-order statistics accurately requires a large amount of data that was not available until recent years. Third, accurate estimation of high-order moments for high-dimensional problems requires modern statistical techniques, such as eigenvalue shrinkage in the spiked covariance model, that have been introduced only in the last decade [26].

### Frequency marching

Most of the cryo-EM reconstruction algorithms start with a low-resolution estimate of the structure, which is then gradually refined to higher resolutions. This process is dubbed *frequency marching*, and it can be done explicitly by constructing

a low-resolution model using an *ab initio* technique that is later refined by ML-EM [57] or, more implicitly, by an iteration-dependent regularization [63].

In [13], a deterministic and mathematically rigorous method to gradually increase the resolution was proposed. The 3D Fourier transform of the object is expanded by concentric shells. At each iteration, each experimental image is compared with many simulated projections of the current low-resolution estimate, and a viewing direction is assigned to each experimental image by template matching (hard angular assignment). Given the angular assignments, an updated structure with one more frequency shells is computed by solving a linear least squares problem, thus progressively increasing the resolution.

## Building blocks in the computational pipeline

This section elaborates on some of the building blocks in the algorithmic pipeline of single-particle reconstruction using cryo-EM. For each task, we introduce the problem and discuss the underlying principles behind some of the most commonly used solutions.

### Motion correction

During data collection, the electron beam induces sample motion that mitigates high-resolution information. The modern detectors (DEDs) acquire multiple frames per micrograph, allowing partial correction of the motion blur by aligning and averaging the frames. In essence, motion correction is an alignment problem (also referred to as *registration* or *synchronization*) that shares many similarities with classical tasks in signal processing and computer vision. The main challenge in alignment is the high noise levels, which hamper the precise estimation of relative shifts between frames.


Several solutions were proposed to the motion-correction problem based on a variety of methods, such as pairwise alignment of all frames and optical flow; see a survey on the subject in [58] and the references therein. The first-generation solutions aimed to estimate the movement of the entire micrograph; however, the drift is not homogeneous across the entire field of view, motivating the development of more accurate local techniques. One such strategy is implemented by the software MotionCor2 [88], which describes the motion as a block-based deformation that varies smoothly throughout the exposure.

The micrograph is first divided into patches, and motions within each patch are estimated based on cross correlation. Then, the local motions are fitted to a time-varying 2D polynomial, which is quadratic in the 2D frame coordinates and cubic in the time axis. Finally, the frames are summed, with or without dose weighting, which was determined according to a radiation-damage analysis. More recently, strategies for per-particle motion correction have been proposed that yield significant improvements in resolution [15], [92]. Figure 6 shows examples of different motion-correction strategies.

### CTF estimation and correction

Each cryo-EM image is affected by the PSF of the microscope through convolution with a kernel  $h_i$ ; see (2). The Fourier





**FIGURE 6.** Examples of beam-induced motion correction. All movement trajectories in this figure are color coded, with cyan representing the position of the first frame and magenta indicating the position of the last frame in the sequence. (a) Strategies for global motion correction compensate for movement of the specimen across the entire field of view containing multiple particles. (Used with permission from [32].) (b) Semilocal strategies for motion correction align frames across subregions defined on a discrete grid. (Used with permission from [88].) (c) Per-particle or local drift correction allows accurate tracking of individual particles throughout the exposure to electrons. (Used with permission from [15].)

transform of the PSF—the CTF—suffers from multiple zero crossings, making its inversion (i.e., deconvolution) challenging. To compensate for the missing frequencies, the standard routine is to apply different PSFs to different micrographs. Thus, spectral information that was suppressed in one micrograph may appear in another.

The CTF cannot be specified precisely by the user. Instead, the number and location of the zero crossings is given by the specific experimental parameters of the microscope, which are hard to control. Thus, a preceding step entails estimating the CTF from the acquired data. Specifically, the CTF is modeled as

$$\text{CTF}(k, \lambda, \Delta f, C_s) = -\sin\left(\pi\lambda|k|^2\Delta f - \frac{\pi\lambda^3|k|^4 C_s}{2} + \alpha\right) \cdot E(|k|), \quad (15)$$

where  $k$  is the frequency index,  $\alpha$  is a small phase-shift term,  $\lambda$  is the electron wavelength,  $\Delta f$  is the objective defocus,  $C_s$  is the spherical aberration, and  $E(|k|)$  is an exponentially decaying envelope function—specified by a parameter called *B-factor*—due to the beam energy spread, beam coherence, and sample drift [59]. The CTF mitigates the very low frequencies, and, therefore, centering projection images (i.e., finding their center of mass) is challenging. In fact, the center of mass depends on the CTF, and, thus, two projection images with the same viewing angles but different CTFs (e.g., different defocus values) will have different centers in real space.

The CTF is approximately constant along concentric rings, although, in practice, those rings might be slightly deformed; this deformation can be modeled by an additional parameter called *astigmatism*. The defocus value of the microscope has a pivotal role: high defocus values enhance low-resolution features and improve the contrast of the image (increase the SNR), whereas low defocus values enhance high-resolution features (fine details) at the cost of lower contrast (lower SNR).

CTF estimation begins by computing the power spectrum of the whole micrograph. The power spectrum exhibits “rings,” called *Thon rings*, that correspond to the oscillations of the sine function (15); those rings are used to fit the CTF’s parameters [see Figure 2(b)]. Two popular software packages, CTFIND4 [59] and Gctf [85], estimate those parameters by generating multiple templates according to the postulated model (15). The template that best fits the measured Thon rings is used as the estimated CTF.

Once the CTF is estimated, the goal is to invert its action. The challenge stems from the structure of the CTF (15) because it has many small values and, thus, direct inversion is impossible. The most popular technique for CTF correction, “phase flipping,” is embarrassingly simple: it disregards the information about amplitude changes and corrects the data only for the sign of the CTF, thereby obtaining the correct phases in Fourier space. Namely, if a micrograph is represented in the Fourier domain by

$$I_{\text{micrograph}}(k) = \text{CTF}(k)I(k), \quad (16)$$

where  $I(k)$  is the micrograph before the CTF action, then the corrected image would be

$$I_{\text{corrected}}(k) = \text{sign}(\text{CTF}(k))I_{\text{micrograph}}(k) = |\text{CTF}(k)|I(k). \quad (17)$$

Another approach is to apply Wiener filtering that attempts to correct both the amplitudes and the phases [29], [69]. However, this approach requires knowledge of the SNR (which is not trivial to estimate in a low-SNR environment) and cannot correct for missing frequencies.

Instead of inverting the CTF explicitly, an attractive alternative is to incorporate the CTF into the forward model of the reconstruction algorithm. This can be done naturally in



the ML-EM approach, the method of moments, and the SGD framework. Because different CTFs are applied to different micrographs, in principle, full CTF correction is feasible.

### Particle picking

A micrograph consists of regions that contain only noise; those with noisy 2D projections; and those with other contaminants, such as carbon film. In particle picking, the goal is to detect and extract the 2D projections (particles) from the noisy micrographs. High-resolution reconstruction typically requires hundreds of thousands of particles, and, thus, manual picking is time consuming and tedious. In addition, it may introduce subjective bias into the procedure.

Many solutions have been proposed in the literature for the particle-picking problem, based on standard edge-detection techniques, machine learning [83], [84], [91], and template matching. For the last of these, a popular strategy is the one implemented in RELION. In this framework, the user manually selects a few hundred particles. These particle images are then 2D classified (see the next section), and the 2D classes are used as templates for an automatic particle picking based on template matching [64].

As discussed in the “Model Bias and Validation” section, a major concern of any particle-picking algorithm is a systematic bias. To alleviate the risk for model bias, a fully automated technique was proposed in [37]. Rather than templates, the algorithm automatically selects a set of reference windows that include both particle and noise windows. The selection is based on the local mean and variance: regions with particles typically have lower mean and higher variance than regions without particles. Then, regions in the micrograph are correlated with the reference windows. The regions that are most likely to contain a particle (high correlation) and those that are least likely (low correlation with all reference windows) are used to train a support-vector machine classifier. The output of this linear classifier is used to pick the particles.

In practice, a considerable number of picked particles are usually noninformative, containing adjacent particles that are too close to each other, only part of a particle, or just pure noise. Consequently, it is common to try to prune out these outliers. In [62], it was proposed to employ several particle pickers simultaneously and compute a consensus between them using a deep neural network. In [90], a simple pruning strategy is devised by viewing the output of the particle-picking algorithm as a mixture of Gaussians.

### 2D classification

As experimental images corresponding to similar viewing directions tend to be very much alike (due to the smoothness of the molecular structure), it is common to divide the images into several classes (i.e., clustering) and average them to increase the SNR. This process is called 2D classification, and the averaged images are referred to as *class averages*. Since the global in-plane rotation of each micrograph is arbitrary, the clustering should be invariant under in-plane rotations: two images

from a similar viewing direction but with different in-plane rotations are supposed to be grouped together after appropriate alignment. The class averages are used for a variety of tasks: to construct *ab initio* models, as templates for particle picking, to provide a quick assessment of the particles, to remove picked particles that are associated with noninformative classes, and for symmetry detection.

There are three main computational aspects that make the 2D classification task quite challenging. First, as already mentioned, low SNR impedes accurate clustering and alignment. Second, the high computational complexity of finding the optimal clustering among hundreds of thousands of images may be prohibitive unless designed carefully. Third, it is not clear what is the appropriate similarity metric to accurately compare images.

Many different strategies for 2D classification were proposed; see, for instance, [73] and [80]. A popular 2D classification algorithm, implemented in RELION [65], is based on the ML-EM scheme. In this algorithm, each observed image is modeled as a sample from the statistical model

$$y_i = h_i * T_{t_i} R_{\theta_i} x_{k_i} + \varepsilon_i, \quad i = 1, \dots, N, \quad (18)$$

where  $k$  is distributed uniformly over  $\{1, \dots, K\}$  (these are the  $K$  class averages to be estimated),  $\theta$  is distributed uniformly over  $\theta \in [0, 2\pi)$ ,  $t$  is drawn from an isotropic 2D Gaussian,  $h_i$  is the estimated PSF, and  $\varepsilon_i \sim \mathcal{N}(0, \sigma^2 I)$ . Given this model, it is straightforward to implement an ML-EM algorithm to estimate  $x_1, \dots, x_K$  following the guidelines of the “3D Reconstruction From Projections” section.

This algorithm, although popular, suffers from three imperative weaknesses. First, the computational burden of running ML-EM is a major hurdle as the algorithm needs to go through all experimental images at each iteration. Second, it assumes that each experimental image is originated from only  $K$  possible class averages. However, this is an inaccurate model, as the orientations of particles in the ice layer are distributed continuously over  $SO(3)$  (in addition to structural variability that may be present). Finally, ML-EM typically suffers from the winner-takes-all phenomenon: most experimental images would correlate well with—and, thus, be assigned to—the class averages that enjoy higher SNRs. As a result, ML-EM tends to output only a few, low-resolution classes. This phenomenon was already recognized in [73] and is also present for ML-EM-based 3D classification and refinement.

An alternative solution was proposed in [87]. In this framework, each image is averaged with its few nearest neighbors, after a proper in-plane alignment. The nearest neighbor search is executed efficiently over the bispectra of the images. The bispectrum is a rotationally invariant feature of an image—that is, it remains unchanged under an in-plane rotation. To reduce the computational complexity and denoise the data, each image is first compressed using a dimensionality-reduction technique, called *steerable principal component analysis (PCA)*, which is the main topic of the next section.

### Denoising and dimensionality-reduction techniques

PCA is a widely used tool for linear dimensionality reduction and denoising, dating back to Pearson [41], [54]. PCA computes the best (in the least squares sense) low-dimensional affine space that approximates the data by projecting the images into the space spanned by the leading eigenvectors of their covariance matrix. In cryo-EM, covariance estimation was introduced by Kam [42] and was used for the first time for dimensionality reduction and image classification by van Heel and Frank [80]. PCA is often used to denoise the experimental images and as part of the 2D classification stage.

Cryo-EM images are equally likely to appear in any in-plane rotation. Consequently, when performing PCA, it makes sense to take all in-plane rotations of each image into account. Luckily, there is a simple way to account for all in-plane rotations without rotating the images explicitly. This can be done by expanding the images in a steerable basis: a basis formed as an outer product of radial functions with an angular Fourier basis; examples are Fourier Bessel, 2D prolate spheroidal wave functions, and Zernike polynomials.

When integrating over all possible in-plane rotations, the covariance matrix of the expansion coefficients enjoys a block-diagonal structure, reducing the computational load significantly: while the covariance matrix of images of size  $L \times L$  has  $L^4$  entries, the block-diagonal structure guarantees that merely  $O(L^3)$  entries are nonzero. This, in turn, reduces the computational complexity of steerable PCA from  $O(NL^4 + L^6)$  to  $O(NL^3 + L^4)$ : the first term is the cost of computing the sample covariance over  $N$  images, and the second term is the cost of the eigendecomposition over all blocks [86].

Classical covariance-estimation techniques usually assume that the number of samples is considerably larger than the signal's dimension. However, this is not the typical case in cryo-EM, in which the dimensionality of the molecule is of the same order as the number of measurements. In this regime, one can take advantage of recent developments in high-dimensional statistics under the “spiked covariance model” [26]. These techniques, based on eigenvalue shrinkage, were recently applied successfully to denoise cryo-EM images [20].

### Mathematical frameworks for cryo-EM data analysis

Inspired by the cryo-EM problem, researchers have studied a couple of abstract mathematical models in recent years. These models provide a general framework for analyzing cryo-EM from theoretical and statistical perspectives while removing some of its complications. In what follows, we introduce the models and succinctly describe some intriguing results.

#### MRA

The MRA model reads [10]

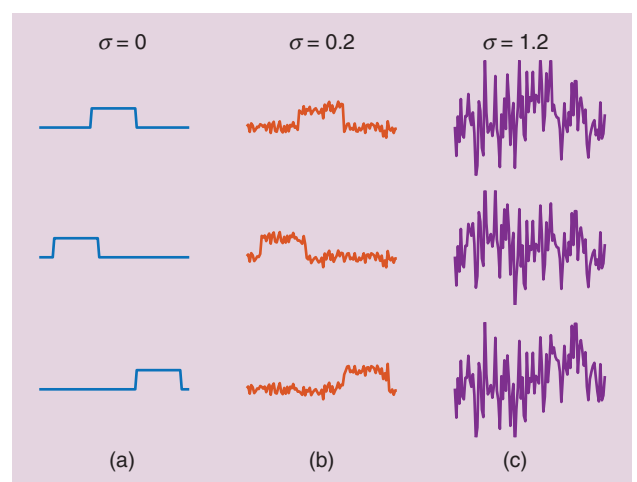
$$y_i = T_i(g_i \circ x) + \varepsilon_i, \quad g_i \in G, \quad (19)$$

where  $G$  and  $T_i$  are, respectively, a known compact group and linear operators, and  $\varepsilon_i \sim \mathcal{N}(0, \sigma^2 I)$ . The signal  $x \in \mathbb{R}^L$  is as-

sumed to lie in a known space (say, the space of signals with a finite spectral expansion) on which random elements of  $G$  act. The task is to estimate the signal  $x$  from the observations  $y_1, \dots, y_N$  while the group elements  $g_1, \dots, g_N$  are unknown. Since the statistics of  $y$  are invariant under the action of a constant  $g$  on  $x$ , the recovery is possible up to left multiplication by some group element  $g \in G$ —that is, we wish to estimate the orbit of  $x$  under the group action.

Figure 7 shows an example of three discrete 1D MRA observations under the group of cyclic shifts  $\mathbb{Z}/L$  at different noise levels. If we assume perfect particle picking, then the cryo-EM model (2) is a special case of (19) when  $G$  is the group of 3D rotations  $SO(3)$  and  $T$  is the linear operator that takes the rotated structure, integrates along the  $z$ -axis (tomographic projection), convolves with the PSF, and samples it on a Cartesian grid. The MRA model formulates many additional applications, including structure from motion in computer vision [5], localization and mapping in robotics [60], study of specimen populations [46], optical and acoustical trapping [28], and denoising of permuted data [53].

While surveying all recent results about the MRA model is beyond the scope of this article, we wish to present a remarkable information/theoretic result about the sample complexity of the problem (and, therefore, also for cryo-EM under the stated assumptions). In many cases, among them when  $T = I$  and the cryo-EM setup [without shifts so that  $G = SO(3)$ ], one can estimate the group elements  $g_i$  from the measurements  $y_i$  in the high-SNR regime using one of many synchronization techniques (based on spectral methods, semidefinite programming, or nonconvex programming) [70], [72]. Once the group elements are identified, one can estimate the signal by aligning all observations (undo the group actions) and averaging.




**FIGURE 7.** An example of MRA observations at different noise levels  $\sigma$ . Each column consists of three different cyclic shifts (the group actions) of a 1D periodic discrete signal (the linear operator  $T$  is the identity operator): (a)  $\sigma = 0$ , (b)  $\sigma = 0.2$ , and (c)  $\sigma = 1.2$ . Clearly, if the noise level is low, estimating the signal is easy: one can align the observations (i.e., undo the group action) and average out the noise. The challenge is to estimate the signal when high noise levels hinder alignment, as in (c).

The variance of averaging over i.i.d. Gaussian variables is proportional to  $1/(N \cdot \text{SNR})$ —that is, the number of measurements should be proportional to  $1/\text{SNR}$  to achieve an accurate estimate. In this case, we say that the sample complexity of the problem in the high-SNR regime is  $1/\text{SNR}$ : this is a common scenario in many data processing tasks. In the low-SNR regime, the situation is radically different. Specifically, it was shown that, when  $\sigma \rightarrow \infty$ , the sample complexity of the MRA problem is determined by the first moment  $\bar{q}$  that distinguishes between different signals. More precisely, in the asymptotic regime where  $N$  and  $\sigma$  diverge, the estimation error of any method is bounded away from zero if  $N \cdot \text{SNR}^{\bar{q}}$  is bounded from above. We then say that the sample complexity in the low-SNR regime is proportional to  $1/\text{SNR}^{\bar{q}}$  [4].

For the cryo-EM setup, assuming perfect particle picking and no CTF, it was shown that the sample complexity depends on the distribution of rotations: for uniform distribution the sample complexity scales as  $1/\text{SNR}^3$  (i.e., the structure is determined by the third moment), whereas for generic nonuniform distribution, the second moment suffices, and, thus, the sample complexity scales as  $1/\text{SNR}^2$  [9], [66]. Similar results were derived for simpler setups; for instance, when a discrete signal is acted on by cyclic shifts (as shown in Figure 7) [3], [9], [19].

The pivotal role of moments in sample complexity analysis led naturally to the design of algorithms based on the method of moments. Some of these algorithms hinge on the tightly related notion of invariant polynomials. A polynomial  $p$  is called *invariant* under the action of a group  $G$  if, for any signal  $x$  in the specified space and any  $g \in G$ , it satisfies  $p(g \circ x) = p(x)$ . One particularly important invariant is the third-order polynomial, the bispectrum—originally proposed by John Tukey [77]—that was used for 2D classification [87] and *ab initio* modeling [42]. We refer readers to [71] and the references therein for a more detailed account of the MRA problem.



**FIGURE 8.** An example of an MTD observation with five signal occurrences at different noise levels (20): (a)  $\sigma = 0$ , (b)  $\sigma = 0.2$ , and (c)  $\sigma = 1.2$ . When (a) the noise level is low, it is easy to detect the signal occurrences and estimate the underlying signal by averaging out the noise. However, (c) when the noise level is high, reliable detection is rendered challenging. The task is, then, to estimate the underlying signal directly, without intermediate detection of its occurrences.

## MTD

MTD is the problem of estimating a signal that occurs multiple times at unknown locations in a noisy measurement. In its simplest form, the MTD problem is an instance of the 1D blind deconvolution problem and can be written as

$$y = x * s + \varepsilon, \quad (20)$$

where  $x \in \mathbb{R}^L$  is the target signal,  $s \in \{0, 1\}^{N-L+1}$  is a binary signal whose ones indicate the location of the signal copies in the measurement  $y \in \mathbb{R}^N$ , and  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$ ; see Figure 8. Detecting the signal occurrences in the data (i.e., estimating the signal  $s$ ) is the analog of particle picking in cryo-EM. Clearly, if the noise level is low, one can estimate  $s$  (and, analogously, detect the particle projections in the micrograph), extract the signal occurrences, and average them to suppress the noise.

However, as mentioned previously, low SNR precludes detection (and particle picking), and, therefore, one must estimate  $x$  directly, without explicit estimation of  $s$ . In [18], it was shown that under certain generative models of  $s$ , the signal  $x$  can be estimated provably, at any noise level, from the bispectrum (third-order statistics). More ambitiously, numerical experiments suggest that the bispectrum suffices to estimate multiple signals  $x_1, \dots, x_K$  from a mix of blind deconvolution problems

$$y = \sum_{i=1}^K x_i * s_i + \varepsilon, \quad (21)$$

without explicit estimation of the binary signals  $s_1, \dots, s_K$  that indicate the location of the corresponding signals.

The MTD model can be extended to formulate a generative model of a micrograph: the key is to treat  $x$  as a random signal that represents the tomographic projections of the molecule  $\phi$ , rather than a deterministic signal as in (20) and (21). Specifically, locations are chosen in the 2D plane: these are the positions of the embedded samples in the micrograph. For each location, a signal is drawn from a probability distribution described by the model

$$x = h * PR_\omega \phi, \quad (22)$$

where the 3D rotation  $R_\omega$  is applied to the volume  $\phi$  according to a (possibly unknown) distribution of  $\omega$  over  $SO(3)$ ,  $P$  is a tomographic projection, and  $h$  is the microscope's PSF; the goal is to estimate  $\phi$ . Therefore, the MTD model paves the way toward fully modeling the cryo-EM problem, including most of its important features [18]. In particular, MTD provides a mathematical and computational framework for reconstructing  $\phi$  directly from micrographs, without intermediate particle picking [17]. A full analysis of this model is still lacking.

## Remaining computational and theoretical challenges

Some interesting and important computational and theoretical challenges that lie ahead for single-particle reconstruction using cryo-EM are given next.

### Conformational heterogeneity: Modeling and recovery

One of the important opportunities offered by cryo-EM is its ability to analyze different functional and conformational

states of macromolecules. Mathematically, it entails estimating multiple 3D structures simultaneously; we refer to this as the *heterogeneity problem*. There is no consensus about the proper way to model the heterogeneity problem, and the computational tools are not well established. Therefore, we believe that it offers an opportunity for researchers with strong mathematical and computational backgrounds to make a profound impact on the field of structural biology.

We now briefly survey different approaches to handling the heterogeneity problem. To this end, we extend the basic cryo-EM model (2) to account for conformational variability:

$$y_i = h_i * T_{t_i} P R_{\omega_i} \phi_i + \text{"noise,"} \quad i = 1, \dots, N, \quad (23)$$

where the goal is to estimate the distribution from which the structures  $\phi_1, \dots, \phi_N$  are sampled. With this formulation, the problem is ill posed: there is not enough information in  $N$  2D images to recover  $N$  3D structures. Hence, additional assumptions about the structures must be made. The different techniques in the field can be broadly classified into two categories: discrete and continuous models. The discrete heterogeneous model assumes that the measurements stem from a few independent volumes (i.e., a discrete distribution). Under this model, each projection image can be written as

$$y_i = h_i * T_{t_i} P R_{\omega_i} \phi_{k_i} + \text{"noise,"} \quad i = 1, \dots, N, \quad (24)$$

where  $\phi_k$ ,  $k = 1, \dots, K$ , represent the  $K$  different volumes. The advantage of this approach is evident: it is easy to extend the ML-EM framework to incorporate several volumes simultaneously. Alternatively, some software packages apply a preliminary 3D classification stage, in which the 2D experimental images are being classified into different structures. CryoSPARC runs multiple trials of an SGD algorithm from different initializations. This procedure occasionally leads to multiple low-resolution estimates, and each of those is refined by an ML-EM algorithm [57]. Although the simplicity of the discrete model is a big advantage, its drawbacks are apparent: it does not scale well for large  $K$  and ignores the correlation between different functional states of the molecule and, thus, overlooks important information.

The second approach, referred to as *continuous heterogeneity*, assumes that  $\phi_1, \dots, \phi_N$  can be embedded in a low-dimensional space. For instance, one approach is to assume that the set of conformations lies in a linear subspace (that can be learned using PCA) or in more intricate low-dimensional manifolds (that could be learned by other spectral methods, e.g., diffusion maps). An alternative suggestion was to model the structure as a set of rigid domains that can move with respect to each other. We refer readers to a thorough survey on the subject in [74].

### Verification

Given a 3D reconstruction, how do we verify that it is a reliable and faithful representation of the underlying molecule? This is a question of crucial importance for any scientific field. Several validation techniques were proposed in the

cryo-EM literature. For example, it is possible to check the consistency of the 3D map by recording pairs of images of the same particles at different tilt angles and comparing the relative angle between orientations assigned to each projection: ideally, it should agree with the relative rotation angle of the microscope's specimen holder used during the experiment.

This approach is called the *tilt-pair validation technique* and is useful only for intermediate-resolution structures [40], [61]. In practice, structure validation is based on a set of heuristics and the experts' knowledge and experience. For instance, it is common to initialize 3D reconstruction algorithms from multiple random points; if all instances attain similar structures, it serves as a validation fidelity. If the same molecule were reconstituted by other technologies, such as X-ray crystallography and NMR, then structures can be compared. In addition, experimental data are usually uploaded to public repositories, and, thus, other researchers can process the same data, experiment with different computational techniques, and compare the results [61].

Verification is also related to the question of determining the resolution of a recovered structure. The current convention is to reconstruct two structures independently, each from one half of the data. The two subsets are chosen at random. The highest frequency for which the two structures agree (up to some tolerance) determines the resolution [75]. This process can be understood as an indication for the confidence we have in the structure at a given resolution.

However, this method is susceptible to systematic flaws; if the same refinement procedure is applied to both halves of the data, it can induce correlated blunders (although the data are independent), and the resolution determination would be unsound. In [24], a simple computational procedure was proposed to validate a structure by assessing the amount of overfitting that is present in the 3D map. Establishing computational tools that provide confidence intervals to estimated structures and are immunized against systematic errors is one of the remaining challenges in the field.

### Theoretical foundations of cryo-EM

Many mathematical and statistical properties of the cryo-EM problem are still unexplained (and even unexplored). A prominent example is the sample complexity of the full cryo-EM problem: given a fixed SNR level (that might be very low) and a fixed setup, how many particles are required to achieve a desired resolution? An initial analysis was conducted in [9] and [66]; see the "MRA" section.

Another interesting question concerns computational complexity and information-computational gaps [11]. The fact that there is enough information to solve a problem (sample complexity) does not immediately imply that there is an efficient (polynomial time) algorithm to solve it. For example, it might be that for nonuniform distribution of rotations, the second moment would suffice to estimate a 3D structure from an information-theoretic perspective [66], but, at the same time, it could be computationally hard in that regime. In that case, we might need to consider the third-order moment to design a computationally efficient algorithm. Information-computational gaps



have been observed empirically in variants of the MRA and MTD models [18], [22]. Related questions—that might be even more challenging—deal with the properties of specific algorithms, such as ML-EM, the method of moments, and SGD, that are not well understood.

Another interesting research thread regards the size limit of molecular structures that can be elucidated by cryo-EM. The common belief in the community is that very small molecules cannot be visualized using cryo-EM. The logic is simple: small molecules induce low contrast, and, thus, low SNR on the micrograph, which, in turn, hinders detection (particle picking) [38]. A recent article contends this belief and suggests that it is possible, at least in principle, to reconstruct structures directly from the micrograph, without particle picking and at any SNR level, given enough data [17].

### Machine learning

The groundbreaking advances in machine learning in the last decade have dramatically reshaped many computational fields and penetrated into some scientific applications. Naturally, learning techniques based on deep neural networks have been applied to cryo-EM as well. Examples are particle picking [83], [84], [91] (as discussed in the “Particle Picking” section), validation [7], 3D reconstruction [89], and particle pruning [62]. In addition, manifold-learning techniques were designed for 3D heterogeneity analysis [31], [51] and denoising [45]. Modern-learning techniques were also implemented in other cryo-EM applications that do not involve single-particle reconstruction; see, for instance, an application to feature extraction in cellular electron cryotomography [23].

Deep learning gained its popularity in applications with low noise levels. The performance of these techniques in more challenging environments, such as with highly contaminated data, is not clear yet. In addition, supervised-learning techniques are susceptible to model bias—the reconstruction will depend heavily on the training data rather than on the experimental images. This explains why the impact of deep learning on the cryo-EM field—especially on more involved tasks, such as 3D reconstruction and the heterogeneity problem—is limited at the moment. We expect that more efforts in this direction will be made in the coming years. In particular, it is still to be clarified whether this set of computational tools can outperform current tools in the field, which are based on more classical statistics.

### Perspective

Single-particle reconstruction using cryo-EM is an alluring research area for investigators interested in developing modern computational tools and sophisticated mathematical models for an emerging scientific field. In this article, we introduced the problem of constituting 3D molecular structures using cryo-EM and described its unique computational characteristics and challenges. We delineated relations between the cryo-EM inverse problem and a variety of disciplines at the core of signal processing, information theory, statistics, machine learning, and group theory. We believe that contributions from

these areas have the potential to drive the field forward. New ideas and solutions can, and should, be tested on experimental cryo-EM data sets publicly available online [1], [2]. We also reviewed two abstract frameworks to conveniently study the cryo-EM inverse problem from computational, statistical, and theoretical perspectives.

We believe that the challenges arising in cryo-EM research provide ample opportunities to investigate and test novel algorithms and advanced mathematical techniques to impact a task of paramount importance: to broaden our understanding of the fundamental mechanisms of life.

### Acknowledgments

We thank the anonymous reviewers and Jong Chul Ye for their valuable comments and suggestions. Amit Singer was supported in part by award R01GM090200 from the National Institute of General Medical Sciences, FA9550-17-1-0291 from the Air Force Office of Scientific Research, Simons Foundation Math+X Investigator Award, the Moore Foundation Data-Driven Discovery Investigator Award, and National Science Foundation BIGDATA Award IIS-1837992.

### Authors

**Tamir Bendory** (bendory@tauex.tau.ac.il) received his B.Sc. degree in biomedical engineering in 2010 and his M.Sc. and Ph.D. degrees in electrical engineering in 2012 and 2015, respectively, all from the Technion-Israel Institute of Technology. He then joined the program in applied and computational mathematics at Princeton University, New Jersey, as a postdoctoral research associate. In 2019, he joined Tel Aviv University as a senior lecturer (assistant professor) of electrical engineering. His research interests include theoretical and computational aspects of signal processing, data science, statistics, and optimization.

**Alberto Bartesaghi** (alberto@cs.duke.edu) received his B.Sc. and M.Sc. degrees in electrical engineering from the Universidad de la Republica, Montevideo, Uruguay, in 1999 and 2001 respectively, and his Ph.D. degree in electrical and computer engineering from the University of Minnesota, Minneapolis, in 2005. He then joined the Biophysics Section of the Laboratory of Cell Biology at the National Cancer Institute/National Institutes of Health, Bethesda, Maryland, to conduct his postdoctoral studies and later became an associate scientist with the Center for Cancer Research. In 2018, he joined Duke University as an associate professor of computer science, biochemistry, and electrical and computer engineering.

**Amit Singer** (amits@math.princeton.edu) received his B.Sc. degree in physics and mathematics and his Ph.D. degree in applied mathematics from Tel Aviv University, Israel, in 1997 and 2005, respectively. He is a professor of mathematics and core member of the Program in Applied and Computational Mathematics and of the Center for Statistics and Machine Learning at Princeton University, New Jersey. He joined Princeton in 2008 as an assistant professor. From 2005 to 2008, he was a Gibbs Assistant Professor in Applied Mathematics at Yale University. His current research interests in applied mathematics

include theoretical and computational aspects of data science and developing computational methods for structural biology.

## References

- [1] "EMPIAR: Electron Microscopy Public Image Archive," European Bioinformatics Institute, Cambridgeshire, UK. [Online]. Available: <https://www.ebi.ac.uk/pdbe/emdb/empiar/>
- [2] "The Electron Microscopy Data Bank (EMDB) at PDB," European Bioinformatics Institute, Cambridgeshire, UK. [Online]. Available: <http://www.ebi.ac.uk/pdbe/emdb/>
- [3] E. Abbe, T. Bendory, W. Leeb, J. M. Pereira, N. Sharon, and A. Singer, "Multireference alignment is easier with an aperiodic translation distribution," *IEEE Trans. Inf. Theory*, vol. 65, no. 6, pp. 3565–3584, 2018. doi: 10.1109/TIT.2018.2889674.
- [4] E. Abbe, J. M. Pereira, and A. Singer, "Estimation in the group action channel," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2018, pp. 561–565. doi: 10.1109/ISIT.2018.8437646.
- [5] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, "Building Rome in a day," *Commun. ACM*, vol. 54, no. 10, pp. 105–112, 2011. doi: 10.1145/2001269.2001293.
- [6] C. Aguerreberre, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5707–5722, 2016. doi: 10.1109/TSP.2016.2600517.
- [7] T. K. Avramov, D. Vyeneli, J. Gomez-Blanco, S. Adinarayanan, J. Vargas, and D. Si, "Deep learning for validating and estimating resolution of cryo-electron microscopy density maps," *Molecules*, vol. 24, no. 6, p. 1181, 2019. doi: 10.3390/molecules24061181.
- [8] X.-C. Bai, G. McMullan, and S. H. Scheres, "How cryo-EM is revolutionizing structural biology," *Trends Biochem. Sci.*, vol. 40, no. 1, pp. 49–57, 2015. doi: 10.1016/j.tibs.2014.10.005.
- [9] A. S. Bandeira, B. Blum-Smith, J. Kileel, A. Perry, J. Weed, and A. S. Wein, Estimation under group actions: Recovering orbits from invariants. 2017. [Online]. Available: [arXiv:1712.10163](https://arxiv.org/abs/1712.10163)
- [10] A. S. Bandeira, Y. Chen, and A. Singer, Non-unique games over compact groups and orientation estimation in cryo-EM. 2015. [Online]. Available: [arXiv:1505.03840](https://arxiv.org/abs/1505.03840)
- [11] A. S. Bandeira, A. Perry, and A. S. Wein, "Notes on computational-to-statistical gaps: Predictions using statistical physics," *Port. Math.*, vol. 75, no. 2, pp. 159–186, 2018. doi: 10.4171/PM/2014.
- [12] S. Banerjee, A. Bartesaghi, A. Merk, P. Rao, S. L. Bulfer, Y. Yan, N. Green, B. Mroczkowski et al., "2.3 Å resolution cryo-EM structure of human p97 and mechanism of allosteric inhibition," *Science*, vol. 351, no. 6275, pp. 871–875, 2016. doi: 10.1126/science.aad7974.
- [13] A. Barnett, L. Greengard, A. Pataki, and M. Spivak, "Rapid solution of the cryo-EM reconstruction problem by frequency marching," *SIAM J. Imaging Sci.*, vol. 10, no. 3, pp. 1170–1195, 2017. doi: 10.1137/16M1097171.
- [14] A. H. Barnett, J. Magland, and L. af Klinteberg, "A parallel nonuniform fast Fourier transform library based on an 'exponential of semicircle' kernel," *SIAM J. Sci. Comput.*, vol. 41, no. 5, pp. C479–C504, 2019. doi: 10.1137/18M120885X.
- [15] A. Bartesaghi, C. Aguerreberre, V. Falconieri, S. Banerjee, L. A. Earl, X. Zhu, N. Grigorieff, J. L. Milne et al., "Atomic resolution cryo-EM structure of  $\beta$ -galactosidase," *Structure*, vol. 26, no. 6, pp. 848–856, 2018. doi: 10.1016/j.str.2018.04.004.
- [16] A. Bartesaghi, A. Merk, S. Banerjee, D. Matthies, X. Wu, J. L. Milne, and S. Subramaniam, "2.2 Å resolution cryo-EM structure of  $\beta$ -galactosidase in complex with a cell-permeant inhibitor," *Science*, vol. 348, no. 6239, pp. 1147–1151, 2015. doi: 10.1126/science.aab1576.
- [17] T. Bendory, N. Boumal, W. Leeb, E. Levin, and A. Singer, Toward single particle reconstruction without particle picking: Breaking the detection limit. 2018. [Online]. Available: [arXiv:1810.00226](https://arxiv.org/abs/1810.00226)
- [18] Bendory, N. Boumal, W. Leeb, E. Levin, and A. Singer, "Multi-target detection with application to cryo-electron microscopy," *Inverse Problems*, vol. 35, no. 10, 2019. doi: 10.1088/1361-6420/ab2aec. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6420/ab2aec/meta>
- [19] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer, "Bispectrum inversion with application to multireference alignment," *IEEE Trans. Signal Process.*, vol. 66, no. 4, pp. 1037–1050, 2018. doi: 10.1109/TSP.2017.2775591.
- [20] T. Bhamre, T. Zhang, and A. Singer, "Denoising and covariance estimation of single particle cryo-EM images," *J. Struct. Biol.*, vol. 195, no. 1, pp. 72–81, 2016. doi: 10.1016/j.jsb.2016.04.013.
- [21] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *Siam Rev.*, vol. 60, no. 2, pp. 223–311, 2018. doi: 10.1137/16M1080173.
- [22] N. Boumal, T. Bendory, R. R. Lederman, and A. Singer, "Heterogeneous multi-reference alignment: A single pass approach," in *Proc. 52nd IEEE Annu. Conf. Information Sciences and Systems (CISS)*, 2018, pp. 1–6. doi: 10.1109/CISS.2018.8362313.
- [23] M. Chen, W. Dai, S. Y. Sun, D. Jonasch, C. Y. He, M. F. Schmid, W. Chiu, and S. J. Ludtke, "Convolutional neural networks for automated annotation of cellular cryo-electron tomograms," *Nat. Meth.*, vol. 14, no. 10, pp. 983–985, 2017. doi: 10.1038/nmeth.4405.
- [24] S. Chen, G. McMullan, A. R. Faruqi, G. N. Murshudov, J. M. Short, S. H. Scheres, and R. Henderson, "High-resolution noise substitution to measure overfitting and validate resolution in 3D structure determination by single particle electron cryomicroscopy," *Ultramicroscopy*, vol. 135, pp. 24–35, Dec. 2013. doi: 10.1016/j.ultramicro.2013.06.004.
- [25] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc. B, Methodol.*, vol. 39, no. 1, pp. 1–22, 1977. doi: 10.1111/j.2517-6161.1977.tb01600.x.
- [26] D. L. Donoho, M. Gavish, and I. M. Johnstone, "Optimal shrinkage of eigenvalues in the spiked covariance model," *Ann. Stat.*, vol. 46, no. 4, pp. 1742, 2018. doi: 10.1214/17-AOS1601.
- [27] J. Dubochet, J. Lepault, R. Freeman, J. Berriman, and J.-C. Homo, "Electron microscopy of frozen water and aqueous solutions," *J. Microsc.*, vol. 128, no. 3, pp. 219–237, 1982. doi: 10.1111/j.1365-2818.1982.tb04625.x.
- [28] P. Elbau, M. Ritsch-Marte, O. Scherzer, and D. Schmutz, Inverse problems of trapped objects. 2019. [Online]. Available: [arXiv:1907.01387](https://arxiv.org/abs/1907.01387)
- [29] J. Frank and P. Penczek, "On the correction of the contrast transfer function in biological electron microscopy," *Optik*, vol. 98, no. 3, pp. 125–129, 1995.
- [30] J. Frank, "Advances in the field of single-particle cryo-electron microscopy over the last decade," *Nat. Protoc.*, vol. 12, no. 2, pp. 209–212, 2017. doi: 10.1038/nprot.2017.004.
- [31] J. Frank and A. Ourmazd, "Continuous changes in structure mapped by manifold embedding of single-particle data in cryo-EM," *Methods*, vol. 100, pp. 61–67, May 2016. doi: 10.1016/j.jymeth.2016.02.007.
- [32] T. Grant and N. Grigorieff, "Measuring the optimal exposure for single particle cryo-EM using a 2.6 Å reconstruction of rotavirus VP6," *Elife*, vol. 4, p. e06980, May 2015. doi: 10.7554/eLife.06980.
- [33] T. Grant, A. Rohou, and N. Grigorieff, "cisTEM, user-friendly software for single-particle image processing," *Elife*, vol. 7, p. e35383, 2018. doi: 10.7554/eLife.35383.
- [34] I. Greenberg and Y. Shkolnisky, "Common lines modeling for reference free ab-initio reconstruction in cryo-EM," *J. Struct. Biol.*, vol. 200, no. 2, pp. 106–117, 2017. doi: 10.1016/j.jsb.2017.09.007.
- [35] L. Greengard and J.-Y. Lee, "Accelerating the nonuniform fast Fourier transform," *SIAM Rev.*, vol. 46, no. 3, pp. 443–454, 2004. doi: 10.1137/S003614450343200X.
- [36] T. W. Guo, A. Bartesaghi, H. Yang, V. Falconieri, P. Rao, A. Merk, E. T. Eng, A. M. Raczkowski et al., "Cryo-EM structures reveal mechanism and inhibition of DNA targeting by a CRISPR-Cas surveillance complex," *Cell*, vol. 171, no. 2, pp. 414–426, 2017. doi: 10.1016/j.cell.2017.09.006.
- [37] A. Heimowitz, J. Adén, and A. Singer, "APPLE picker: Automatic particle picking, a low-effort cryo-EM framework," *J. Struct. Biol.*, vol. 204, no. 2, pp. 215–227, 2018. doi: 10.1016/j.jsb.2018.08.012.
- [38] R. Henderson, "The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological molecules," *Q. Rev. Biophys.*, vol. 28, no. 2, pp. 171–193, 1995. doi: 10.1017/S003358350000305X.
- [39] R. Henderson, "Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 45, pp. 18,037–18,041, 2013. doi: 10.1073/pnas.1314449110.
- [40] R. Henderson, S. Chen, J. Z. Chen, N. Grigorieff, L. A. Passmore, L. Ciccarelli, J. L. Rubinstein, R. A. Crowther et al., "Tilt-pair analysis of images from a range of different specimens in single-particle electron cryomicroscopy," *J. Mol. Biol.*, vol. 413, no. 5, pp. 1028–1046, 2011. doi: 10.1016/j.jmb.2011.09.008.
- [41] I. Jolliffe, *Principal Component Analysis*. Berlin: Springer-Verlag, 2011.
- [42] Z. Kam, "The reconstruction of structure from electron micrographs of randomly oriented particles," *J. Theor. Biol.*, vol. 82, no. 1, pp. 15–39, 1980. doi: 10.1016/0022-5193(80)90088-0.
- [43] Y. Kang, O. Kuybeda, P. W. de Waal, S. Mukherjee, N. Van Eps, P. Dutka, X. E. Zhou, A. Bartesaghi et al., "Cryo-EM structure of human rhodopsin bound to an inhibitory G protein," *Nature*, vol. 558, no. 7711, pp. 553–558, 2018. doi: 10.1038/s41586-018-0215-y.
- [44] W. Kühlbrandt, "The resolution revolution," *Science*, vol. 343, no. 6178, pp. 1443–1444, 2014. doi: 10.1126/science.1251652.
- [45] B. Landa and Y. Shkolnisky, "The steerable graph Laplacian and its application to filtering image datasets," *SIAM J. Imaging Sci.*, vol. 11, no. 4, pp. 2254–2304, 2018. doi: 10.1137/18M1169394.
- [46] A. Levis, Y. Y. Schechner, and R. Talmon, "Statistical tomography of microscopic life," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2018, pp. 6411–6420. doi: 10.1109/CVPR.2018.00671.
- [47] M. Liao, E. Cao, D. Julius, and Y. Cheng, "Structure of the TRPV1 ion channel determined by electron cryo-microscopy," *Nature*, vol. 504, no. 7478, pp. 107, 2013. doi: 10.1038/nature12822.

- [48] H. Liu, B. K. Poon, D. K. Saldin, J. C. Spence, and P. H. Zwart, "Three-dimensional single-particle imaging using angular correlations from X-ray laser data," *Acta Crystallogr. A, Found. Crystallogr.*, vol. 69, no. 4, pp. 365–373, 2013. doi: 10.1107/S0108767313006016.
- [49] D. Matthies, C. Bae, G. E. Toombes, T. Fox, A. Bartesaghi, S. Subramaniam, and K. J. Swartz, "Single-particle cryo-EM structure of a voltage-activated potassium channel in lipid nanodiscs," *Elife*, vol. 7, p. e37558, 2018. doi: 10.7554/eLife.37558.
- [50] A. Merk, A. Bartesaghi, S. Banerjee, V. Falconieri, P. Rao, M. I. Davis, R. Prangani, M. B. Boxer et al., "Breaking cryo-EM resolution barriers to facilitate drug discovery," *Cell*, vol. 165, no. 7, pp. 1698–1707, 2016. doi: 10.1016/j.cell.2016.05.040.
- [51] A. Moscovich, A. Halevi, J. Andén, and A. Singer, "Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes," *Inverse Probl.*, to be published. doi: 10.1088/1361-6420/ab4f55. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6420/ab4f55/meta>
- [52] E. Nogales and S. H. Scheres, "Cryo-EM: A unique tool for the visualization of macromolecular complexity," *Mol. Cell*, vol. 58, no. 4, pp. 677–689, 2015. doi: 10.1016/j.molcel.2015.02.019.
- [53] A. Pananjady, M. J. Wainwright, and T. A. Courtade, "Denoising linear models with permuted data," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2017, pp. 446–450. doi: 10.1109/ISIT.2017.8006567.
- [54] K. Pearson, "On lines and planes of closest fit to systems of points in space," *London, Edinburgh, Dublin Philosoph. Mag. J. Sci.*, vol. 2, no. 11, pp. 559–572, 1901. doi: 10.1080/14786440109462720.
- [55] P. A. Penczek, R. Renka, and H. Schomberg, "Gridding-based direct Fourier inversion of the three-dimensional ray transform," *J. Opt. Soc. Am. A*, vol. 21, no. 4, pp. 499–509, 2004. doi: 10.1364/JOSA.21.000499. [Online]. Available: <https://www.osapublishing.org/josaa/abstract.cfm?uri=JOSAA-21-4-499>
- [56] G. Pragier and Y. Skhkolnisky, "A common lines approach for ab initio modeling of cyclically symmetric molecules," *Inverse Problems*, Nov. 19, 2019. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6420/ab2fb2>
- [57] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, "cryoSPARC: Algorithms for rapid unsupervised cryo-EM structure determination," *Nat. Meth.*, vol. 14, no. 3, pp. 290–296, 2017. doi: 10.1038/nmeth.4169.
- [58] Z. Ripstein and J. Rubinstein, "Processing of cryo-EM movie data," *Meth. Enzymol.*, vol. 579, pp. 103–124, 2016. doi: 10.1016/bs.mie.2016.04.009.
- [59] A. Rohou and N. Grigorieff, "CTFFIND4: Fast and accurate defocus estimation from electron micrographs," *J. Struct. Biol.*, vol. 192, no. 2, pp. 216–221, 2015. doi: 10.1016/j.jsb.2015.08.008.
- [60] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard, "SE-sync: A certifiably correct algorithm for synchronization over the special Euclidean group," *Int. J. Robot. Res.*, vol. 38, nos. 2–3, pp. 95–125, 2019. doi: 10.1177/0278364918784361.
- [61] P. B. Rosenthal and J. L. Rubinstein, "Validating maps from single particle electron cryomicroscopy," *Curr. Opin. Struct. Biol.*, vol. 34, pp. 135–144, 2015. doi: 10.1016/j.sbi.2015.07.002.
- [62] R. Sanchez-Garcia, J. Segura, D. Maluenda, J. M. Carazo, and C. O. S. Sorzano, "Deep consensus, a deep learning-based approach for particle pruning in cryo-electron microscopy," *IUCrJ*, vol. 5, no. 6, pp. 854–865, 2018. doi: 10.1107/S2052252518014392.
- [63] S. H. Scheres, "RELION: Implementation of a Bayesian approach to cryo-EM structure determination," *J. Struct. Biol.*, vol. 180, no. 3, pp. 519–530, 2012. doi: 10.1016/j.jsb.2012.09.006.
- [64] S. H. Scheres, "Semi-automated selection of cryo-EM particles in RELION-1.3," *J. Struct. Biol.*, vol. 189, no. 2, pp. 114–122, 2015. doi: 10.1016/j.jsb.2014.11.010.
- [65] S. H. Scheres, M. Valle, R. Nuñez, C. O. Sorzano, R. Marabini, G. T. Herman, and J.-M. Carazo, "Maximum-likelihood multi-reference refinement for electron microscopy images," *J. Mol. Biol.*, vol. 348, no. 1, pp. 139–149, 2005. doi: 10.1016/j.jmb.2005.02.031.
- [66] N. Sharon, J. Kileel, Y. Khoo, B. Landa, and A. Singer, "Method of moments for 3-D single particle ab initio modeling with non-uniform distribution of viewing angles," *Inverse Probl.*, vol. 34, no. 7, 2019. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6420/ab6139/meta>
- [67] M. Shatsky, R. J. Hall, S. E. Brenner, and R. M. Glaeser, "A method for the alignment of heterogeneous macromolecules from electron microscopy," *J. Struct. Biol.*, vol. 166, no. 1, pp. 67–78, 2009. doi: 10.1016/j.jsb.2008.12.008.
- [68] F. Sigworth, "A maximum-likelihood approach to single-particle image refinement," *J. Struct. Biol.*, vol. 122, no. 3, pp. 328–339, 1998. doi: 10.1006/jsbi.1998.4014.
- [69] C. V. Sindelar and N. Grigorieff, "An adaptation of the Wiener filter suitable for analyzing images of isolated single particles," *J. Struct. Biol.*, vol. 176, no. 1, pp. 60–74, 2011. doi: 10.1016/j.jsb.2011.06.010.
- [70] A. Singer, "Angular synchronization by eigenvectors and semidefinite programming," *Appl. Comput. Harmon. Anal.*, vol. 30, no. 1, pp. 20–36, 2011. doi: 10.1016/j.acha.2010.02.001.
- [71] A. Singer, "Mathematics for cryo-electron microscopy," in *Proc. Int. Congr. Mathematicians*, 2018, pp. 3995–4014.
- [72] A. Singer and Y. Skhkolnisky, "Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming," *SIAM J. Imaging Sci.*, vol. 4, no. 2, pp. 543–572, 2011. doi: 10.1137/090767777.
- [73] C. Sorzano, J. Bilbao-Castro, Y. Skhkolnisky, M. Alcorlo, R. Melero, G. Caffarena-Fernández, M. Li, G. Xu et al., "A clustering approach to multireference alignment of single-particle projections in electron microscopy," *J. Struct. Biol.*, vol. 171, no. 2, pp. 197–206, 2010. doi: 10.1016/j.jsb.2010.03.011.
- [74] C. Sorzano, A. Jiménez, J. Mota, J. Vilas, D. Maluenda, M. Martínez, E. Ramírez-Aportela, T. Majtner et al., "Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy," *Acta Crystallogr. F, Struct. Biol. Commun.*, vol. 75, no. 1, pp. 19–32, 2019. doi: 10.1107/S2053230X18015108.
- [75] C. Sorzano, J. Vargas, J. Otón, V. Abrishami, J. de la Rosa-Trevín, J. Gómez-Blanco, J. Vilas, R. Marabini et al., "A review of resolution measures and related aspects in 3D electron microscopy," *Prog. Biophys. Mol. Biol.*, vol. 124, pp. 1–30, Mar. 2017. doi: 10.1016/j.phiomolbio.2016.09.005.
- [76] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, "EMAN2: An extensible image processing suite for electron microscopy," *J. Struct. Biol.*, vol. 157, no. 1, pp. 38–46, 2007. doi: 10.1016/j.jsb.2006.05.009.
- [77] J. Tukey, "The spectral representation and transformation properties of the higher moments of stationary time series," in D. R. Brillinger, *The Collected Works of John W. Tukey*, vol. 1. Belmont, CA: Wadsworth, 1953, pp. 165–184.
- [78] B. Vainshtein and A. Goncharov, "Determination of the spatial orientation of arbitrarily arranged identical particles of unknown structure from their projections," *Sov. Phys. Dokl.*, vol. 31, p. 278, Apr. 1986.
- [79] M. Van Heel, "Angular reconstruction: A posteriori assignment of projection directions for 3D reconstruction," *Ultramicroscopy*, vol. 21, no. 2, pp. 111–123, 1987. doi: 10.1016/0304-3991(87)90078-7.
- [80] M. Van Heel and J. Frank, "Use of multivariate statistics in analysing the images of biological macromolecules," *Ultramicroscopy*, vol. 6, no. 2, pp. 187–194, 1981. doi: 10.1016/0304-3991(81)90059-0.
- [81] K. R. Vinothkumar and R. Henderson, "Single particle electron cryomicroscopy: Trends, issues and future perspective," *Q. Rev. Biophys.*, vol. 49, p. e13, Jan. 2016. doi: 10.1017/S0033583516000068.
- [82] B. von Ardenne, M. Mechelke, and H. Grubmüller, "Structure determination from single molecule X-ray scattering with three photons per image," *Nat. Commun.*, vol. 9, p. 2375, June 2018. doi: 10.1038/s41467-018-04830-4.
- [83] T. Wagner, F. Merino, M. Stabrin, T. Moriya, C. Antoni, A. Apelbaum, P. Hagel, O. Sitsel et al., "SPHIRE-crYOLO is a fast and accurate fully automated particle picker for cryo-EM," *Commun. Biol.*, vol. 2, p. 218, June 2019. doi: 10.1038/s42003-019-0437-z.
- [84] F. Wang, H. Gong, G. Liu, M. Li, C. Yan, T. Xia, X. Li, and J. Zeng, "DeepPicker: A deep learning approach for fully automated particle picking in cryo-EM," *J. Struct. Biol.*, vol. 195, no. 3, pp. 325–336, 2016. doi: 10.1016/j.jsb.2016.07.006.
- [85] K. Zhang, "Gctf: Real-time CTF determination and correction," *J. Struct. Biol.*, vol. 193, no. 1, pp. 1–12, 2016. doi: 10.1016/j.jsb.2015.11.003.
- [86] Z. Zhao, Y. Skhkolnisky, and A. Singer, "Fast steerable principal component analysis," *IEEE Trans. Comput. Imaging*, vol. 2, no. 1, pp. 1–12, 2016. doi: 10.1109/TCI.2016.2514700.
- [87] Z. Zhao and A. Singer, "Rotationally invariant image representation for viewing direction classification in cryo-EM," *J. Struct. Biol.*, vol. 186, no. 1, pp. 153–166, 2014. doi: 10.1016/j.jsb.2014.03.003.
- [88] S. Q. Zheng, E. Palovcak, J.-P. Armache, K. A. Verba, Y. Cheng, and D. A. Agard, "MotionCor2: Anisotropic correction of beam-induced motion for improved cryo-electron microscopy," *Nat. Meth.*, vol. 14, no. 4, pp. 331–332, 2017. doi: 10.1038/nmeth.4193.
- [89] E. D. Zhong, T. Bepler, J. H. Davis, and B. Berger, "Reconstructing continuously heterogeneous structures from single particle cryo-EM with deep generative models. 2019. [Online]. Available: [arXiv:1909.05215](https://arxiv.org/abs/1909.05215)
- [90] Y. Zhou, A. Moscovich, T. Bendory, and A. Bartesaghi, "Unsupervised particle sorting for high-resolution single-particle cryo-EM," *Inverse Probl.*, to be published. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6420/ab5ec8/meta>
- [91] Y. Zhu, Q. Ouyang, and Y. Mao, "A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy," *BMC Bioinformatics*, vol. 18, no. 1, p. 348, 2017. doi: 10.1186/s12859-017-1757-y.
- [92] J. Zivanov, T. Nakane, and S. H. Scheres, "A Bayesian approach to beam-induced motion correction in cryo-EM single-particle analysis," *IUCrJ*, vol. 6, no. 1, pp. 5–17, 2019. doi: 10.1107/S205225251801463X.
- [93] "Press release: The Nobel Prize in Chemistry 2017," Oct. 4, 2017. [Online]. Available: <https://www.nobelprize.org/prizes/chemistry/2017/press-release/>



# The Importance of Continuity for Linear Time-Invariant Systems

Linear time-invariant (LTI) systems play a fundamental role in signal processing. Continuity is an important property of LTI systems, without which many conclusions about LTI systems, such as convolution formula and commutative law, are not true in general. However, this concept does not receive as much attention as it should in the literature of signal processing. In this “Lecture Notes” article, we provide a detailed discussion on the importance of continuity for LTI systems. For simplicity, we mainly consider the discrete-time signals and systems.

## Relevance

It is widely known that an LTI system  $T$  can be completely characterized by its impulse response, and the output signal of  $T$  can be computed by the convolution of the input signal with the impulse response of  $T$ . In most signal processing literature, e.g., [1]–[3], the continuity is not required in the derivation of this result, whereas in some mathematical literature, e.g., [4]–[6], the continuity is required. To our best knowledge, however, there are no examples in the literature to show that the convolution formula does not hold for discontinuous LTI systems.

In this note, we construct two discontinuous LTI systems that have the same impulse response but behave differently. Moreover, we show that the convolution formula does not hold for these systems. We also prove that, for a

linear system, the continuity implies the bounded-input, bounded-output (BIBO) stability, whereas the BIBO stability does not imply the continuity.

## Prerequisites

A discrete-time signal is usually expressed as a sequence of real or complex numbers that conveys information [7]. Because all signals in real life have finite energy, we restrict our discussion in the  $\ell^2(\mathbb{Z})$  space defined by

$$\ell^2(\mathbb{Z}) = \left\{ x \mid \sum_{n=-\infty}^{\infty} |x(n)|^2 < \infty \right\}, \quad (1)$$

where  $\mathbb{Z}$  is the set of integers,  $x$  is a doubly infinite sequence of complex numbers whose  $n$ th component is  $x(n)$ , and  $|x(n)|$  is the absolute value of  $x(n)$ . Note that, in most signal processing literature, both the signal  $x$  and its  $n$ th component are denoted by  $x(n)$ . Although the exact meaning of  $x(n)$  will be clear from the context, we use  $x$  and  $x(n)$  to denote the signal and its value at time instant  $n$ , respectively, to avoid the potential ambiguity.

A system is formally defined as a device that manipulates signals. Therefore, it can be described by an operator between some given signal spaces. In this article, we denote the output signal of a system  $T$  with an input signal  $x$  by  $y = Tx$  or  $y = T[x]$ . We assume that the readers are familiar with the concepts of vector space, normed space, linear independent, and convergence; refer to [8] and [9] for more details. The concept of continuous linear operator is given next

because we mainly study the continuity of LTI systems in this article.

**Definition 1** (continuous linear operator [9])

Let  $(X, \|\cdot\|_X)$  and  $(Y, \|\cdot\|_Y)$  be two normed spaces over the complex number field  $\mathbb{C}$ , where  $\|\cdot\|_X$  and  $\|\cdot\|_Y$  are the norms on  $X$  and  $Y$ , respectively. A mapping  $T: X \rightarrow Y$  is called a *linear operator* from  $X$  to  $Y$  if it satisfies the following conditions:

- 1)  $T[x + y] = Tx + Ty$  for all  $x, y \in X$ .
- 2)  $T[\alpha x] = \alpha(Tx)$  for all  $\alpha \in \mathbb{C}$  and all  $x \in X$ .

If  $T: X \rightarrow Y$  is a linear operator and there exists a real number  $\gamma \in \mathbb{R}$  such that  $\|Tx\|_Y \leq \gamma \|x\|_X$  for all  $x \in X$ , then  $T$  is called a *bounded linear operator*. The smallest  $\gamma$  such that the inequality holds for all nonzero  $x \in X$  is called the *norm* of  $T$ , denoted by  $\|T\|$ . If  $T: X \rightarrow Y$  is a linear operator and the sequence  $\{Tx_k\}_{k=1}^{\infty}$  converges to  $Tx_0 \in Y$  whenever the sequence  $\{x_k\}_{k=1}^{\infty}$  converges to  $x_0 \in X$ , then  $T$  is said to be *continuous* at  $x_0$ . If  $T$  is linear and continuous at every point of  $X$ , then it is called a *continuous linear operator* on  $X$ .  $\square$

A key property of linear operators is given by Theorem 1, given next.

## Theorem 1

Let  $T$  be a linear operator defined on a normed space  $X$ . Then:

- 1)  $T$  is continuous if and only if  $T$  is continuous at a single point of  $X$ .
- 2)  $T$  is continuous if and only if  $T$  is bounded.  $\square$

*Proof 1* [9, p. 97]

Obviously, for any  $m \in \mathbb{Z}$ , the translation operator  $S_m$  defined by

$$(S_m x)(n) = x(n - m) \quad (2)$$

is a continuous linear operator. If a system  $T$  commutes with translations, i.e.,  $TS_m = S_m T$ , then  $T$  is said to be time invariant. Besides linearity and time invariance, another critical property of practical systems is the stability. A system is called *BIBO stable* if every bounded input produces a bounded output [3, p. 198]. Finally, we review the concept of direct sum that will be used in Theorem 2 and Example 1.

Definition 2 (direct sum [8])

If  $\mathcal{V}$  and  $\mathcal{W}$  are two subspaces of a vector space  $\mathcal{Y}$  such that  $\mathcal{V} \cap \mathcal{W} = \{0\}$ , then  $\mathcal{V}$  and  $\mathcal{W}$  are said to be disjoint. Moreover, the set

$$\begin{aligned} \mathcal{X} &= \mathcal{V} \oplus \mathcal{W} \\ &= \{v + w \mid v \in \mathcal{V}, w \in \mathcal{W}\} \end{aligned} \quad (3)$$

is called the *direct sum* of  $\mathcal{V}$  and  $\mathcal{W}$ . Obviously,  $\mathcal{X}$  is also a subspace of  $\mathcal{Y}$ . Furthermore, every  $x \in \mathcal{X}$  has a unique representation  $x = v + w$  with  $v \in \mathcal{V}$  and  $w \in \mathcal{W}$ .  $\square$

### Problem statement and solution

We first prove two preparation theorems in the section “Two Preparation Theorems,” based on which we construct a counterexample to illustrate the importance of continuity for LTI systems in the section “The Importance of Continuity.” We also discuss the relation between continuity and BIBO stability in the “Continuity

and BIBO Stability” section, and finally, the conditions of continuity for linear systems are given in the “Conditions for Continuity” section.

### Two preparation theorems

Because an LTI system defined on a finite-dimensional space is automatically continuous [9, p. 96], the phenomenon that the convolution formula does not hold can occur only in infinite-dimensional spaces. Therefore, constructing the counterexamples is not trivial, and some preliminary results are required. Hence, we first prove two theorems in this section that will be used in the next section to construct the counterexample. Theorem 2 states that we can construct a new LTI system from two given LTI systems defined on disjoint spaces. Theorem 3 states that we can construct numerous LTI systems between spaces spanned by the translations of some basic signals.

#### Theorem 2

Let  $\mathcal{V}$  and  $\mathcal{W}$  be disjoint subspaces and  $\mathcal{X} = \mathcal{V} \oplus \mathcal{W}$  be the direct sum of  $\mathcal{V}$  and  $\mathcal{W}$ . If  $L_1: \mathcal{V} \rightarrow \mathcal{W}$  and  $L_2: \mathcal{W} \rightarrow \mathcal{V}$  are LTI systems, then  $T: \mathcal{X} \rightarrow \mathcal{X}$  defined by

$$T[x] = T[v + w] = L_2[w] + L_1[v] \quad (4)$$

is an LTI system on  $\mathcal{X}$ , where  $v \in \mathcal{V}$ ,  $w \in \mathcal{W}$ , and  $x \in \mathcal{X}$ .  $\square$

#### Proof 2

This theorem can be illustrated by Figure 1. Note that because  $L_1$  and  $L_2$  are time invariant,  $\mathcal{V}$  and  $\mathcal{W}$  must be closed under translations, i.e.,  $S_m[v] \in \mathcal{V}$  and  $S_m[w] \in \mathcal{W}$  for all




FIGURE 1. The linearity and time invariance of  $L_1$  and  $L_2$  imply the linearity and time invariance of  $T$ .

$v \in \mathcal{V}$ , all  $w \in \mathcal{W}$ , and all  $m \in \mathbb{Z}$ . We first show that  $T$  is well defined. Because  $\mathcal{X}$  is the direct sum of  $\mathcal{V}$  and  $\mathcal{W}$ , it follows from Definition 2 that the decomposition  $x = v + w$  is unique. Therefore,  $L_2[w] \in \mathcal{V}$  and  $L_1[v] \in \mathcal{W}$  are uniquely determined by  $x$ . Consequently,  $T[x] \in \mathcal{X}$  is uniquely determined by  $x$ , which shows that  $T$  is well defined on  $\mathcal{X}$ . Next, we prove that  $T$  is an LTI system. Let  $x_1 = v_1 + w_1$  and  $x_2 = v_2 + w_2$ , where  $v_1, v_2 \in \mathcal{V}$  and  $w_1, w_2 \in \mathcal{W}$ . Then

$$\begin{aligned} T[x_1 + x_2] &\stackrel{(a)}{=} T[(v_1 + v_2) + (w_1 + w_2)] \\ &\stackrel{(b)}{=} L_2[w_1 + w_2] + L_1[v_1 + v_2] \\ &\stackrel{(c)}{=} L_2[w_1] + L_2[w_2] + L_1[v_1] + L_1[v_2] \\ &\stackrel{(d)}{=} (L_2[w_1] + L_1[v_1]) + (L_2[w_2] + L_1[v_2]) \\ &\stackrel{(e)}{=} T[x_1] + T[x_2], \end{aligned} \quad (5)$$

where (a) follows from the definition of  $\mathcal{X}$ , (b) from the definition of  $T$ , (c) from the linearity of  $L_2$  and  $L_1$ , (d) from the commutative law and associative law of addition, and (e) from the definition of  $T$ . Let  $x = v + w$  with  $v \in \mathcal{V}$  and  $w \in \mathcal{W}$ . Then for any  $\alpha \in \mathbb{C}$ , we have

$$\begin{aligned} T[\alpha x] &\stackrel{(a)}{=} T[\alpha v + \alpha w] \\ &\stackrel{(b)}{=} L_2[\alpha w] + L_1[\alpha v] \\ &\stackrel{(c)}{=} \alpha L_2[w] + \alpha L_1[v] \\ &\stackrel{(d)}{=} \alpha (L_2[w] + L_1[v]) \\ &\stackrel{(e)}{=} \alpha T[x], \end{aligned} \quad (6)$$

where (a) follows from the definition of  $\mathcal{X}$ , (b) from the definition of  $T$ , (c) from the linearity of  $L_2$  and  $L_1$ , (d) from the distributive law between addition and multiplication, and (e) from the definition of  $T$ . Finally, because  $S_m$  is an LTI system,  $S_m x = S_m v + S_m w$ . Hence,

$$\begin{aligned} T[S_m x] &\stackrel{(a)}{=} T[S_m v + S_m w] \\ &\stackrel{(b)}{=} L_2[S_m w] + L_1[S_m v] \\ &\stackrel{(c)}{=} S_m L_2[w] + S_m L_1[v] \\ &\stackrel{(d)}{=} S_m (L_2[w] + L_1[v]) \\ &\stackrel{(e)}{=} S_m [T[x]], \end{aligned} \quad (7)$$

where (a) follows from the linearity of  $S_m$ , (b) from the definition of  $T$ , (c) from the time invariance of  $L_2$  and  $L_1$ ,

(d) from the linearity of  $S_m$ , and (e) from the definition of  $T$ . Equations (5)–(7) indicate that  $T$  is an LTI system on  $\mathcal{X}$ . ■

### Theorem 3

Let  $v_0$  be a nonzero signal in  $\ell^2(\mathbb{Z})$  and  $\mathcal{V} = \text{span}\{S_m[v_0]\}_{m \in \mathbb{Z}}$  be the linear space spanned by  $v_0$  and its shifted versions, i.e.,

$$\mathcal{V} = \left\{ \sum_{m=M_1}^{M_2} \alpha_m S_m[v_0] \mid M_1, M_2 \in \mathbb{Z}, \alpha_m \in \mathbb{C} \right\}. \quad (8)$$

Similarly, let  $w_0$  be an arbitrary signal in  $\ell^2(\mathbb{Z})$  and  $\mathcal{W}$  be the linear space spanned by  $w_0$  and its shifted versions. Suppose the set  $\{S_m[v_0]\}_{m \in \mathbb{Z}}$  is linearly independent; then every nonzero  $v \in \mathcal{V}$  has a unique expansion  $v = \sum_{m=M_1}^{M_2} \alpha_m S_m[v_0]$ . In addition, for any  $l \in \mathbb{Z}$ ,

$$L[v] = \sum_{m=M_1}^{M_2} \alpha_m S_m S_l[w_0] \quad (9)$$

defines an LTI system from  $\mathcal{V}$  to  $\mathcal{W}$ , as shown in Figure 2. □

### Proof 3

The linearity of  $L$  is implied by its definition; thus, the only thing we must prove is the time invariance of  $L$ . Since  $S_k$  is an LTI system,

$$\begin{aligned} S_k[v] &= \sum_{m=M_1}^{M_2} \alpha_m S_m S_k[v_0] \\ &= \sum_{m=M_1}^{M_2} \alpha_m S_{m+k}[v_0]. \end{aligned} \quad (10)$$

Let  $p = m + k$  and  $\beta_p = \alpha_m$ ; then  $S_k[v]$  can be expressed as

$$S_k[v] = \sum_{p=M_1+k}^{M_2+k} \beta_p S_p[v_0]. \quad (11)$$

Thus, we have

$$\begin{aligned} L[S_k[v]] &\stackrel{(a)}{=} \sum_{p=M_1+k}^{M_2+k} \beta_p S_p S_l[w_0] \\ &\stackrel{(b)}{=} \sum_{m=M_1}^{M_2} \alpha_m S_k S_m S_l[w_0] \\ &\stackrel{(c)}{=} S_k \sum_{m=M_1}^{M_2} \alpha_m S_m S_l[w_0] \\ &\stackrel{(d)}{=} S_k[L[v]], \end{aligned} \quad (12)$$

where (a) follows from the definition of  $L$ , (b) from the property of

$S_{m+k} = S_k S_m$ , (c) from the linearity of  $S_k$ , and (d) from the definition of  $L$ . Therefore,  $L$  is time invariant, and the proof is completed. ■

### The importance of continuity

In this section, we will construct an example based on Theorems 2 and 3 to illustrate the importance of continuity for LTI systems. More specifically, we will show that without the condition of continuity, an LTI system may not be completely specified by its impulse response. Moreover, the convolution formula and the commutative law may not hold either.

### Example 1

Let  $v_0$  be the Kronecker delta signal,

$$v_0(n) = \delta(n) = \begin{cases} 1, & n = 0, \\ 0, & n \neq 0. \end{cases} \quad (13)$$

Then the linear space  $\mathcal{V} = \text{span}\{S_m[v_0]\}_{m \in \mathbb{Z}}$  is closed under translations. Because the set  $\{S_m[v_0]\}_{m \in \mathbb{Z}}$  is linearly independent, every nonzero signal  $v \in \mathcal{V}$  can be uniquely expressed as

$$v = \sum_{m=M_1}^{M_2} \alpha_m S_m[v_0], \quad (14)$$

where  $M_1, M_2 \in \mathbb{Z}$  and  $\alpha_m \in \mathbb{C}$  depend on  $v$ . Similarly, let  $w_0$  be a right-sided sequence,

$$w_0(n) = \begin{cases} \frac{1}{n}, & n > 0, \\ 0, & n \leq 0. \end{cases} \quad (15)$$

Then every nonzero signal  $w$  in  $\mathcal{W} = \text{span}\{S_m[w_0]\}_{m \in \mathbb{Z}}$  can be uniquely expressed as

$$w = \sum_{m=N_1}^{N_2} \beta_m S_m[w_0], \quad (16)$$

where  $N_1, N_2 \in \mathbb{Z}$  and  $\beta_m \in \mathbb{C}$  depend on  $w$ . By Theorem 3, we can define three LTI systems  $L_1: \mathcal{W} \mapsto \mathcal{V}$ ,  $L_2: \mathcal{W} \mapsto \mathcal{V}$ , and  $L_3: \mathcal{V} \mapsto \mathcal{W}$  as

$$L_1[w] = \sum_{m=N_1}^{N_2} \beta_m S_m[v_0] \quad (17)$$

$$L_2[w] = \sum_{m=N_1}^{N_2} \beta_m S_{m+1}[v_0] \quad (18)$$

$$L_3[v] = \sum_{m=M_1}^{M_2} \alpha_m S_m[w_0]. \quad (19)$$

It is clear that every  $v \in \mathcal{V}$  and  $w \in \mathcal{W}$  has finite energy (recall that  $\sum_{m=1}^{\infty} 1/m^2$  converges [10, p. 62]). In addition, because  $\mathcal{V}$  consists of finite-length signals and  $\mathcal{W}$  consists of infinite-length signals except for the zero signal,  $\mathcal{V}$  and  $\mathcal{W}$  are disjoint subspaces of  $\ell^2(\mathbb{Z})$ . Let  $\mathcal{X} = \mathcal{V} \oplus \mathcal{W}$  and  $x = v + w$  with  $v \in \mathcal{V}$  and  $w \in \mathcal{W}$ ; then it follows from Theorem 2 that  $T_1: \mathcal{X} \mapsto \mathcal{X}$  defined by

$$T_1[x] = L_1[w] + L_3[v], \quad (20)$$

and  $T_2: \mathcal{X} \mapsto \mathcal{X}$  defined by

$$T_2[x] = L_2[w] + L_3[v], \quad (21)$$

are LTI systems. There are some interesting results about  $T_1$  and  $T_2$  as shown in Figure 3.


1)  $T_1$  and  $T_2$  have the same impulse response  $h = w_0$ . This is because  $\delta = v_0 \in \mathcal{X}$ ,

$$T_1[\delta] = T_2[\delta] = L_3[v_0] = w_0. \quad (22)$$

2)  $T_1$  and  $T_2$  are not equal. This is because


$$T_1[w_0] = L_1[w_0] = v_0, \quad (23)$$

whereas



**FIGURE 2.**  $L$  defines an LTI system from  $\mathcal{V}$  to  $\mathcal{W}$  for any  $l \in \mathbb{Z}$ , where  $\mathcal{V} = \text{span}\{S_m[v_0]\}_{m \in \mathbb{Z}}$  and  $\mathcal{W} = \text{span}\{S_m[w_0]\}_{m \in \mathbb{Z}}$ .





**FIGURE 3.**  $T_1$  and  $T_2$  are LTI systems with the same impulse response  $h$ . However,  $T_1 \neq T_2$ ,  $T_1[w_0] \neq h * w_0$ , and  $T_2 T_1 \neq T_1 T_2$ .

$$T_2[w_0] = L_2[w_0] = S_1[v_0] \neq v_0. \quad (24)$$

- 3)  $T_1[w_0] \neq h * w_0$ , where  $*$  is the convolution operator. This is because  $h = w_0$ , whereas

$$T_1[w_0] = v_0 \neq w_0 * w_0. \quad (25)$$

Similarly,  $T_2[w_0] \neq h * w_0$ .

- 4)  $T_1$  and  $T_2$  do not obey the commutative law. This is because

$$T_2 T_1[v_0] = T_2[w_0] = S_1[v_0], \quad (26)$$

whereas

$$T_1 T_2[v_0] = T_1[w_0] = v_0 \neq S_1[v_0]. \quad (27)$$

Items 1) and 2) indicate that two LTI systems may behave differently although they have the same impulse response. Item 3) indicates that not every LTI system has the convolution relation between the input and output. Item 4) indicates that two LTI systems may not be commutative.  $\square$

Because a continuous LTI system can be expressed by convolution (see [5, Sec. 1.3] or [6, Sec. 2.5]), the reason why  $T_1$  and  $T_2$  in Example 1 have these peculiar behaviors is because they are discontinuous. The proof is given next. Recall that the norm on  $\ell^2(\mathbb{Z})$  is defined by

$$\|x\|_2 = \left( \sum_{n=-\infty}^{\infty} |x(n)|^2 \right)^{1/2}. \quad (28)$$

Let

$$x_k = \sum_{m=1}^k \frac{1}{m} S_m[\delta]. \quad (29)$$

Then the sequence  $\{x_k\}_{k=1}^{\infty}$  converges to  $w_0$  because

$$\lim_{k \rightarrow \infty} \|x_k - w_0\|_2^2 = \lim_{k \rightarrow \infty} \sum_{m=k+1}^{\infty} \frac{1}{m^2} = 0. \quad (30)$$

Since  $T_1$  is an LTI system, we have

$$\begin{aligned} T_1[x_k] &= \sum_{m=1}^k \frac{1}{m} S_m T_1[\delta] \\ &= \sum_{m=1}^k \frac{1}{m} S_m[w_0]. \end{aligned} \quad (31)$$

However,

$$\begin{aligned} T_1[\lim_{k \rightarrow \infty} x_k] &= T_1[w_0] \\ &= v_0 \neq \lim_{k \rightarrow \infty} T_1[x_k]. \end{aligned} \quad (32)$$

Therefore,  $T_1$  is discontinuous at  $w_0$ . It follows from Theorem 1 that  $T_1$  is discontinuous at every point of  $\mathcal{X}$ . Similarly, it can be proved that  $T_2$  is discontinuous too.

### Continuity and BIBO stability

The relation between continuity and BIBO stability is given by Theorem 4.

#### Theorem 4

Let  $T$  be a linear system from a normed space  $\mathcal{X}$  to a normed space  $\mathcal{Y}$ . If  $T$  is continuous, then it is BIBO stable. However, the converse is not true.  $\square$

#### Proof 4

Because  $T$  is a continuous linear operator, it is bounded (see Theorem 1). Thus, it follows from Definition 1 that for all  $x \in \mathcal{X}$ , we have  $\|Tx\|_{\mathcal{Y}} \leq \|T\| \|x\|_{\mathcal{X}}$ . Consequently, every bounded input  $x \in \mathcal{X}$  produces a bounded output  $Tx$ , indicating that  $T$  is BIBO stable. However, the follow-

ing example shows that BIBO stability does not imply continuity.  $\blacksquare$

### Example 2

Let  $\mathcal{X}$  be the space of all polynomials defined on  $[0, 1]$  with a norm given by

$$\|x\| = \max_{0 \leq t \leq 1} |x(t)|. \quad (33)$$

Let  $T = d/dt$  be the differential operator. Obviously,  $T$  is linear on  $\mathcal{X}$ , and the range of  $T$  is contained in  $\mathcal{X}$ . Because every continuous function on a closed interval attains its maximum and minimum [10, p. 89],  $\|Tx\|$  is bounded for all  $x \in \mathcal{X}$ . Thus,  $T$  is BIBO stable. However, let  $x_k(t) = t^k$ , then

$$\frac{\|T[x_k]\|}{\|x_k\|} = \frac{\max_{0 \leq t \leq 1} |kt^{k-1}|}{\max_{0 \leq t \leq 1} |t^k|} = k. \quad (34)$$

Thus there does not exist any real number  $\gamma$  such that  $\|Tx\| \leq \gamma \|x\|$  for all  $x \in \mathcal{X}$ . That is,  $T$  is unbounded, or equivalently,  $T$  is discontinuous.  $\square$

Stability means that small changes in the input generate small changes in the output. In other words, the output must depend continuously on the input [11, p. 2]. Hence, continuity implies stability for linear systems, and it must be considered when designing practical linear systems. Example 2 shows that BIBO stability does not always guarantee the stability of a linear system, however. Therefore, continuity is more stringent and important than BIBO stability for linear systems.

### Conditions for continuity

There are three ways to determine whether a linear system  $T$  is continuous or not.

- 1) If the input space of  $T$  is finite dimensional, then  $T$  is continuous [9, p. 96].
- 2) If the input-output relation of  $T$  is expressed by convolution and the impulse response  $h$  of  $T$  is absolutely summable, i.e.,  $\sum_{n \in \mathbb{Z}} |h(n)| < \infty$ , then  $T$  is continuous [3, p. 329].
- 3) If  $T$  satisfies the superposition principle with infinitely many terms, i.e.,

$$T\left[\sum_{k=1}^{\infty} \alpha_k x_k\right] = \sum_{k=1}^{\infty} \alpha_k T[x_k] \quad (35)$$

holds for every convergent series  $\sum_{k=1}^{\infty} \alpha_k x_k$ , then  $T$  is continuous [12].

(continued on page 100)

## Delivering Standards to Industries: The MPEG Case

**D**elivering products, services, and applications to the masses is a problem that more companies face every day in a globalized economy. Imagine what the problem was like 30 years ago when the delivery of media content was a reality but had been achieved by providing solutions that applied to specific media types, delivery methods, countries, regions, industries, and companies. Telecommunication companies distributed music; cable operators distributed TV via cable; terrestrial and satellite broadcasters did the same through terrestrial and satellite networks; and different types of businesses distributed all sorts of recorded media on physical support (film, laser discs, compact cassettes, Video Home System/Betamax cassettes, and so forth).

Figure 1 depicts the vertical businesses of the analog world when media distribution was a collection of industry-dependent systems, each using its own

technologies for the baseband signal. The figure is an extremely simplified representation of the situation of those years because it does not take into account the region, country, or company-based differences within each industry.

It took 30 years of progress in electronics to convert Nyquist's intuition into the practice of his industry. During the 1960s, the International Telecommunication Union (ITU) developed G.711, the standard for digital speech. It required 50 years to do the same for the entertainment industry (Philips and Sony's compact disc) and television industry [ITU-Radio Communication Sector (ITU-R) Recommendation 601].

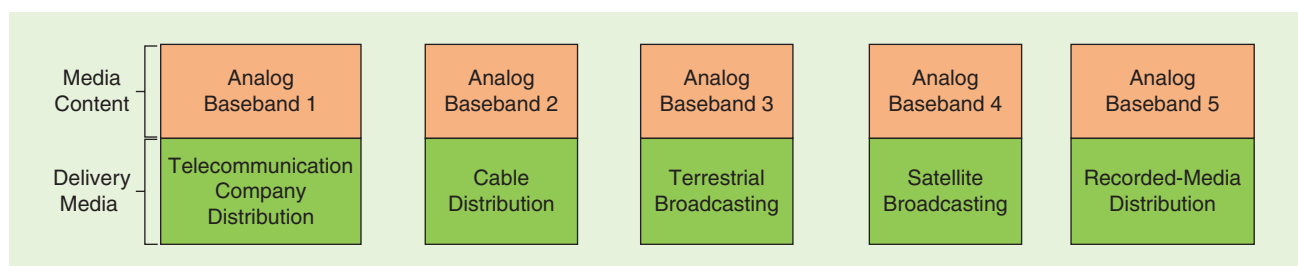
The telecommunication network could carry 64 kilobits/s of digital speech, but no consumer-level media delivery of that time could carry the 1.41 megabits/s of digital audio, much less the 216 megabits/s of digital video. Therefore, during the 1960s, studies on the compression of digitized media began in earnest. Twenty or so years of research started paying off. In 1982, the

ITU approved Recommendation H.100 for the transmission of videoconferences at 1.5–2 megabits/s, and two years later, it began the work that would produce Recommendation H.261 for video coding at  $p \times 64$  kilobits/s, approved in 1988. Other industries, countries, and companies toyed with digital-media compression technologies, including the

- European Digital Video Recorder project
- Philips and RCA independent interactive video on compact-disc projects
- Committee for Mixed Telephone and Television studies on digital television transmission for “primary contribution”
- Italian broadcaster RAI and Telettra's HDTV codec for satellite broadcasting
- Advisory Committee on Advanced Television Science plans to introduce HDTV in the United States (originally analog but later turned to digital).

Companies, industries, and standards committees were jockeying for a position in the upcoming digital world, with

Digital Object Identifier 10.1109/MSP.2019.2953637  
Date of current version: 26 February 2020



**FIGURE 1.** The analog-media distribution.

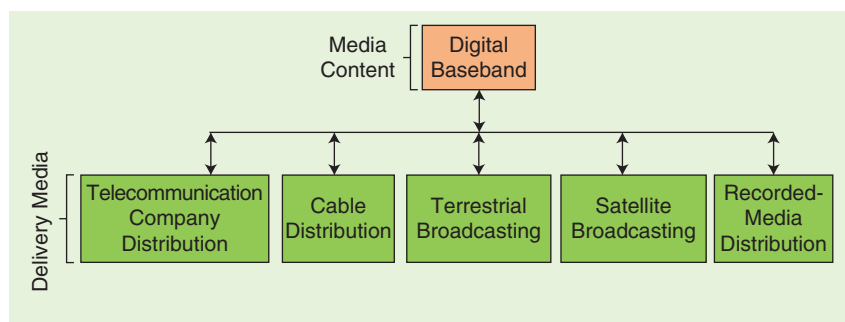


FIGURE 2. The digital-media distribution (à la MPEG).

Table 1. The media-related standards committees (1980s).

Organization	Media	Committee
ITU-T	Speech	SG XV WP 1
	Video	SG XV WP 2
ITU-R	Audio	SG 10
	Video	SG 11
IEC	Audio recording	SC 60A
	Video recording	SC 60B
	Audio-visual equipment	TC 84
	Receivers	SC 12A and G
ISO	Photography	TC 42
	Cinematography	TC 36

ITU-T: ITU Telecommunication Standardization Sector.

the mind-set that the future distribution of digital media would be similar to the one sketched in Figure 1 for analog media: The “baseband signal” of each delivery medium would be digital, thus using new technology but differing for each industry and possibly for each country/region.

MPEG (which, at the time, did not have a name yet) was conceived as an organization that would unify the digital baseband and provide generic (i.e., domain-independent) digital media compression, as depicted in Figure 2. IT distribution was not part of the original plan because the Internet was just an interesting and useful research tool. In the mid-1990s, few would think that a system designed for computer communication would carry the megabits/second streams required by

digital media. If anything, telecommunication companies were working on a solution called the *asynchronous transfer mode* to achieve that.

### A place to execute the MPEG plan

The first problem to be tackled concerned the fact that the MPEG digital baseband had to be established on international standards that had global validity. The second problem was that the standards should not be industry specific. When the MPEG idea took shape in July 1987, selecting a home to implement the idea was the primary concern. There were many choices, as shown by the list of international committees in Table 1 that were created for various reasons (regulation or simply the necessity of an independent technical reference) and cater to the needs of different industries.

Since MPEG was conceived to be industry neutral, committees that were already developing standards in the traditional media were considered to be unsuitable because they represented vested interests. The choice fell to ISO Technical Committee (TC) 97, Computers and Information Processing, which had the Subcommittee (SC) 2 character sets and information coding, which included the Working Group (WG) 8 coding of audio and picture information.

In 1987, ISO TC 97 merged with IEC TC 83, Information Technology Equipment. The resulting joint technical committee (JTC) was called *ISO/IEC JTC 1, Information Technology*. SC 2 and its WGs, including WG 8, became part of JTC 1. MPEG was established as the Experts Group on Moving Pictures within ISO/IEC JTC 1/SC 2/WG 8, in 1988. Note that the group was an

informal organizational entity within the ISO. In 1991, SC 2/WG 8 seceded from SC 2 and became SC 29. WG 8’s Moving Picture Experts Group, from which the MPEG name was derived, became WG 11, Coding of Moving Pictures and Audio (but everyone in the industry and general public called it *MPEG*).

### The rules that drive MPEG work

Many know about MPEG, but most think of it as yet another committee producing admittedly useful standards. However, it is a different beast than most other committees, for a number of reasons:

- 1) *Common standards for multiple industries:* From its early days, MPEG developed a process to achieve the common understanding that enables it to lay down the specifications of the common standards. Since MPEG does not have an industry of its own, it uses its Requirements Group to develop generic specifications that are used to interact with the trade associations and standards organizations that represent the main industries affected by its standards.
- 2) *Technology standards as toolkits:* Basic media-compression technology should be shared by all industries, but individual sectors do not necessarily need the same functions and performance. Therefore, industry-agnostic standardization must account for all of the requirements of the industries that a standard will serve and have the means to flexibly allocate technology and performance to one industry without encumbering others. MPEG has made its standards toolkit based.
- 3) *Standards for the market:* When MPEG came to the fore, a company that had developed a successful product would try to get a “standard” stamp on it, share the technology with its competitors, and enjoy the economic benefits of its technology. MPEG reshuffled the order of the steps. Instead of waiting for the market to decide which technology would win, MPEG experts developed collaboratively defined “best standards” by assessing individual

technologies and deciding which would be selected on the basis of pre-defined criteria.

- 4) *Standards that anticipate the future:* When technology moves fast or accelerates, waiting is a luxury that MPEG cannot afford. The original MPEG-1 and MPEG-2 standards were developed at a time when the industry did not have the enabling technology. Today, the components of the MPEG-I standard, which is under development, rely on interfaces that are still shaky or just hypothetical. Having standards that lead, rather than trail, technology is a tough trial-and-error game. However, it is the only one possible for digital media. MPEG has been a part of this competition for the past 30 years.
- 5) *Compete and collaborate:* MPEG favors competition to the maximum extent possible. This is achieved by
  - calling for solutions that must be comprehensively described (i.e., without black boxes) to qualify for consideration
  - having MPEG experts assess the merit of proposed technologies, followed by collaboration based on
    - Test Model, a software platform made up of components selected by MPEG experts
    - improvement of different areas of the Test Model with Core Experiments that add software to implement accepted technologies until the standard is stable.
- 6) *Industry-friendly standards:* These are achieved by
  - making the standards independent of the display format, instead of tying the display to one capture format
  - serving one need without encumbering other users through the notion of profiles and levels, the former being a subset of general interoperability and the latter a grade of performance within a profile
  - standardizing the decoder and leaving room for improved versions.
- 7) *Standards for audio and video together:* Today, we take these for granted, but it was not always so.

Audio and video used to belong to different departments within companies and standards organizations. MPEG was the first standard group to raise the industry's attention to the integrated packages that include the systems component. At the same time, MPEG standards are toolkits, so specific users can select the components they need and replace them with others from different sources.

- 8) *Technology is always on the move:* The constant change of the digital media technology landscape informs MPEG standards. Most of the developers of the technologies that are incorporated in MPEG standards invest the royalties from previous standards to develop new technologies for new standards.
- 9) *The role of research for MPEG standards:* MPEG is not in the research business, but without research there would be no MPEG. The MPEG work plan promotes corporate/academic research to push companies to improve their technologies for future participation in the MPEG standards.
- 10) *Rethinking what we are:* The industries participating in MPEG-1 included telecommunication and consumer electronics. With MPEG-2 the "club" was enlarged to terrestrial and satellite broadcasting and cable TV. With MPEG-4, IT companies joined. Later, research institutions and academia signed up in droves (today, they account for approximately one-quarter of the membership). With MPEG-I, MPEG faces new challenges because the demand for standards for immersive services and applications is there, but technology immaturity deprives MPEG of its usual "anchors."
- 11) *Standards as enablers, not disablers:* Since MPEG standards are not owned by a specific industry,

when developing a new standard, MPEG must assess and accommodate all legitimate functional requirements (for example, the value of functionality, cost of implementation, ability to aggregate the proposed functionality with others, and so on). Similarly, the decision to accept or reject a legitimate request from an industry to add a functionality to an existing standard is driven only by the value that the proposal brings, as substantiated by use cases.

- 12) *Standards need a business model:* Standardization is not a commercial business, but there is a business model that has guided MPEG and driven its standard-development work. During the late 1980s, when MPEG began its mission, industry and academia had worked on video-compression technolo-

gies for some three decades and filed many patents. A royalty-free video coding standard (that the ISO/IEC/ITU call *Option 1*) was certainly possible but probably unattractive because of its low performance.

Therefore, MPEG decided to develop standards that had the best performance, without consideration of the intellectual property rights that were involved. Most patent holders have reinvested their royalties from existing standards in technologies for future ones. The MPEG business model has created a standard-producing machine that feeds itself with new technologies.

**In the mid-1990s, few would think that a system designed for computer communication would carry the megabits/second streams required by digital media.**

## The MPEG organization

Since its early days, MPEG has created and disbanded several groups for specific needs to manage its broad range of activities. Figure 3 depicts the current MPEG organization. New ideas are typically submitted to the Requirements Group. If they are considered worthy of further explorations, they are discussed in ad-hoc groups. An ad-hoc group



will report its findings at the following MPEG meeting. After a few iterations, MPEG will produce a call for evidence (CFE) or call for proposals (CFPs), with publicity through a press release. At that time, an ad-hoc group is charged with disseminating the CFE or CFP, prepare for the logistics of the tests, and perform a first assessment of the responses. This does not mean that the Requirements Group is no longer involved with the standard, since the group typically continues the development of use cases and requirements while the technical work progresses. When necessary, new requirements are reviewed with the appropriate technology groups and may give rise to new CFPs whose outcomes are fed into the

**Many know about MPEG, but most think of it as yet another committee producing admittedly useful standards.**

work of the relevant technology groups after an assessment.

Today, the task of the Test Group is no longer confined to assessing proposal quality at CFE and CFP time. Designing and executing tests with the appropriate quality to support core experiments has

become the norm, especially in the Video Group. If a proposal is requested, the Requirements Group, in conjunction with the Test Group and possibly

one or more others, reviews the results from the ad-hoc group and makes a final assessment. If the submitted technologies are judged to be sufficient to initiate the development of a standard, the activity is transferred to the appropriate group(s) for completion.

Feeding documents into the MPEG process is of vital importance because the number of uploaded submissions can easily cross the level of 1,000–1,500 documents per meeting. This has been done electronically since 1995. For MPEG experts, a place with so many people discussing requirements and assessing, integrating, and testing media technologies is exciting. They know what happens when and where through a web application that gives a full view of the meetings that are held and topics of discussion.

A key MPEG feature is the immediate availability of the necessary technical expertise to discuss matters that cross organizational boundaries. The speed of development and quality of MPEG standards would hardly be possible if the necessary expertise to address multifaceted issues could not be deployed in a timely manner. Figure 4 breaks down the activities at an MPEG meeting. For instance, out of a total of 97 activities:

- 1) Twenty-nine relate to the processing of standards through the canonical stages of the committee draft, draft international standard, and final draft international standard and the equivalent for amendments, technical reports, and corrigenda. In other words, at every meeting, MPEG is working on roughly 10 deliverables (i.e., standards, amendments, technical reports, and corrigenda) at different stages of the approval process.
- 2) Twenty-two correspond to working drafts, i.e., new activities that have not entered the approval stages.
- 3) Eight concern technologies under consideration; in other words, new technologies that are being considered to enhance existing standards.
- 4) Eight pertain to requirements, typically for new standards.
- 5) Six support core experiments.

Figure 4 does not provide a quantitative measure of how many documents are produced for each activity and how big they were. As an example, point-cloud compression has 20 core experiments and eight exploration experiments under way, while MPEG-5 Essential Video Coding (EVC) has only one large core experiment. An average value of the

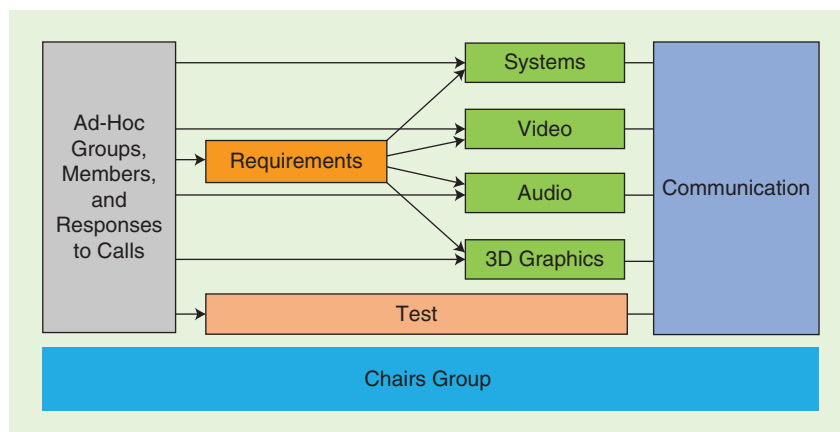


FIGURE 3. The MPEG workflow.




FIGURE 4. The activities at an MPEG meeting.

activities at the March 2019 meeting is provided by dividing the number of output documents (212) by the number of activities (97), yielding 2.2.

### The MPEG ecosystem

MPEG standards are different from most others because they consider the needs of the many industries that intend to use them and, hence, want to have a say in how they are developed: consumer electronics, IT, broadcasting, telecommunications, and so forth. Taking into account so many views is a burden for developers, but the standards that are eventually produced are abstracted from the small (or big) needs that are specific to individual

sectors. Figure 5 depicts how MPEG has succeeded in its role of abstracting the needs of its client digital-media industries.

Figure 5, however, does not describe all of the ecosystem actors. In MPEG-1, the consumer electronics industry was often able to develop the technology it needed to make products that used the standard. With MPEG-2, this was less the case, and independent companies offering encoding and decoding chips sprang up. Today, the industry that implements (as opposed to using or selling products based on) MPEG standards has grown to be a very important element of the ecosystem. It typically provides components to companies

that manufacture a complete product (sometimes this happens inside the same company, but the logic is the same).

MPEG standards can be implemented using various combinations of software, hardware, and hybrid technologies. The choice for the hardware is very wide, from various integrated-circuit architectures to analog technologies. The latter choice is for devices with extremely low power consumption, although with limited compression. Devices that use neural networks are soon to come. Other technologies are likely to find use in the future, such as quantum computing and genomic technologies.

Figure 6 represents with some accuracy the way MPEG acquires the necessary

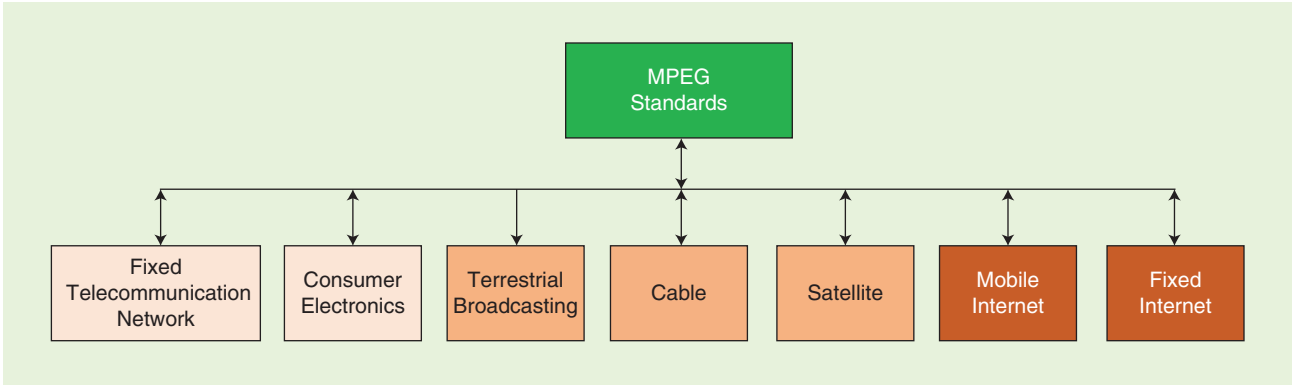


FIGURE 5. The MPEG client digital-media industries.

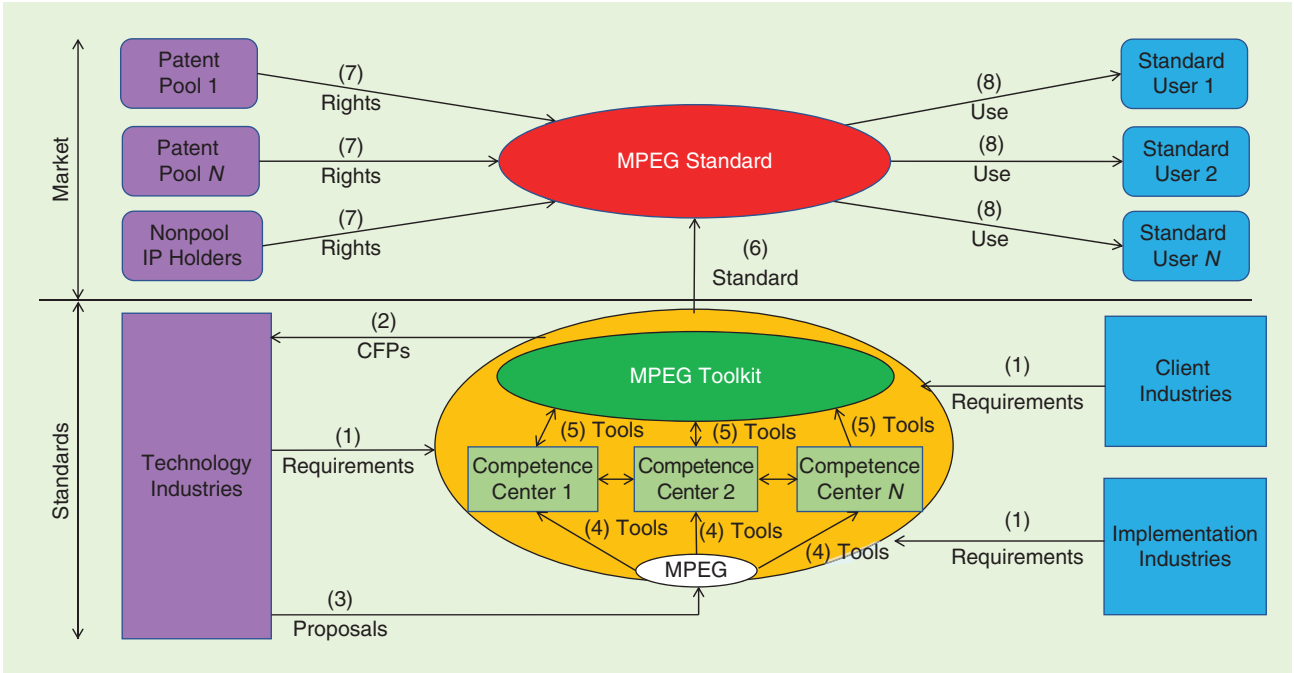


FIGURE 6. The MPEG standard workflow.

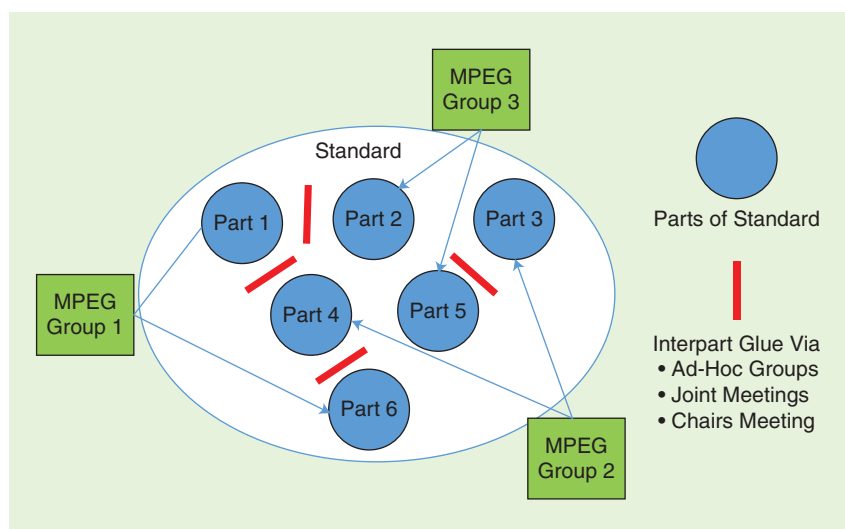


FIGURE 7. The structure of the MPEG standards.

technologies to develop a standard and how implementers of a standard can acquire the necessary rights. When MPEG intends to develop a standard, it solicits requirements from the industry. Three segments of the industry can typically provide requirements: eventual users of the standard, implementers of the standard, and providers of technology to develop the standard. Note that a specific company may belong to one or more than one industry, in the sense that one department of a company may belong to one industry and another department may belong to a different one.

MPEG receives (step 1), assesses, refines, and harmonizes the requirements received. When these have reached sufficient maturity, a CFP is issued (step 2). Companies belonging to the technology industry submit proposals in response to the CFP (step 3). MPEG assesses the coding tools contained in the proposals and assigns them to specific competence centers, for example, audio, video, 3D graphics, systems, etc. (step 4), which adapt and refine the tools (possibly interacting between them) and adds the selected ones to those in the MPEG Toolkit (step 5), where all the technologies adopted in past standards are contained.

When the development of the standard is completed (step 6), companies claiming rights to the tools included in the standard may decide to join one or more than one patent pool or not to

join any patent pool at all. Users of the standard (e.g., implementers and service providers) get use rights from patent holders (step 7) for use in their products, services, or applications (step 8).

Although MPEG has no role in the last step, the perceived performance of MPEG as a provider of industry standards is highly dependent on how this step unfolds.

Conformance testing typically plays a role in assessing an implementation's compliance with the standard.

MPEG's role cannot be described by the simple standards-provider/client-industry relationship. MPEG is a complex ecosystem that works because all of its entities play the role that is proper to them. MPEG has devised an organization that enables it to deploy the necessary level of expertise in specific technology areas, such as video and audio coding and file formats. At the same time, the organization enables it to identify instances where interfaces between different media subsystem are needed.

Figure 7 describes how most MPEG standards are developed. Different groups with diverging competences develop parts of a standard, making sure that the distinct components talk to one another. Some elements are designed to work with others in systems that are identified during the context-objectives-use-cases phase. However, many

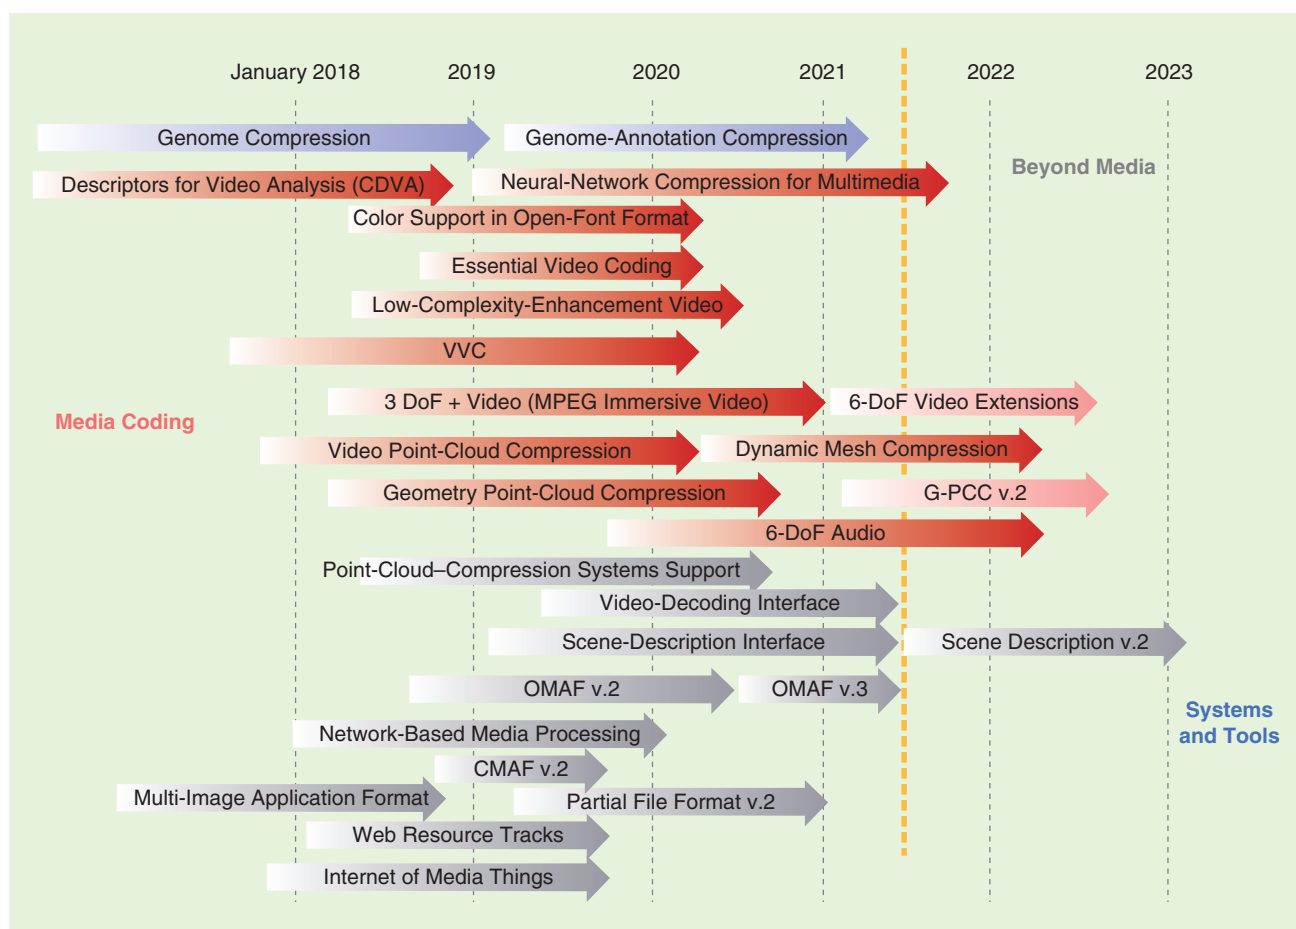
parts are not tightly bound because, in general, it is possible to use them separately. In other cases, there are parts from other sources that must work tightly together, and that is where MPEG provides the “glue” by using ad-hoc groups, joint meetings, chair meetings, and so on.

## What is next for MPEG

Through its 31-year history, MPEG has produced some 200 specifications, tens of editions and hundreds of amendments. Its mission is not over, as the workplan in Figure 8 shows. The following is a brief summary of the main standards planned for the next few years.

### Media coding

- 1) *Versatile Video Coding (VVC)*: This is the flagship video-compression activity that will deliver another round of improved performance. It is expected to be the platform on which MPEG will build new technologies for immersive visual experiences.
- 2) *EVC*: This shorter-term project has less-ambitious goals than VVC. EVC is designed to satisfy the urgent need for a standard with a simplified IP landscape.
- 3) *Immersive visual technologies*: These investigate technologies applicable to the visual information captured by different camera arrangements.
- 4) *Point-cloud compression*: This refers to two standards capable of compressing 3D point clouds captured with multiple cameras and depth sensors: video and geometry point-cloud compression. The algorithms in both standards are scalable and progressive, and they support random access to point-cloud subsets. Video point-cloud compression is lossy, and geometry point-cloud compression is also lossless.
- 5) *Immersive audio*: MPEG-H 3D audio supports a 3-degree of freedom (DoF) (yaw, pitch, and roll) experience at the movie “sweet spot.” More complete user experiences, however, are needed, such as 6 DoF (adding x, y, and z). These can be achieved with additional metadata and rendering technology.



**FIGURE 8.** The MPEG work plan (October 2019). CDVA: compact descriptors for video analysis; CMAF: common media-application format.

## Systems and tools

- 1) *Omnidirectional media format (OMAF)*: This format supports the interoperable exchange of omnidirectional (360°) content. In version 1, users can only yaw, pitch, and roll their head, but version 2 will support more functionalities.
- 2) *Storage of point-cloud conversion data in MP4 File Format*: MPEG is developing systems to enable the storage and movement of compressed point clouds through dynamic adaptive streaming over HTTP, MPEG media transport, and so forth.
- 3) *Scene-description interface*: MPEG is extending an existing scene-description technology to enable rich immersive experiences.
- 4) *Service interface for immersive media*: Network-based media processing will enable a user to access potentially sophisticated process-

ing functionality made available by a network service via a standard application-programming interface.

- 5) *When things are media things*: The Internet of Media Things will enable the creation of networks of intelligent devices, such as sensors and actuators.

## Beyond media

- 1) *Standards for biotechnology applications*: MPEG is finalizing the five parts of the MPEG-G standard and developing a new standard for the compression of annotations to DNA.
- 2) *Coping with neural networks everywhere*: MPEG is developing compression of neural networks for multimedia content description and analysis (part 17 of MPEG-7), after receiving responses to its CFP for neural-network compression.

## A future for MPEG

MPEG has worked without interruption for 31 years, adding industry members and witnessing technologies morph as new ones come to the fore. Today, MPEG faces another major change and has to “shed its skin” one more time.

## MPEG is about compression

MPEG has produced five generations of video-compression standards, each offering additional compression and features. More of the same is expected for the sixth (VVC). MPEG has produced an equivalent number of audio coding standards. Will industry keep asking for more compression? I would like to answer with a resolute yes, as I believe that there will always be a need for it but not necessarily in the “old way.” If we take the example of the light field, we see that more compression is needed



if the distribution of that kind of information is going to become real within the current (human) generation. The answer to the compression question is, at least in the current time frame, however, that the compression we have is enough, although we need new standards for other noncompression features (6 DoF, in the case of audio). The trend is less noticeable in video, but 3-DoF+ moves in the same direction as audio, and the exact nature of 6-DoF-video is still evolving. Point-cloud compression is probably a different story because we are still low on the learning curve.

### *MPEG is also about systems aspects*

Systems aspects have been the enabling factor of the success of many MPEG standards, and the future will not change that role. The trend toward immersive media will require an even deeper integration between compressed media and the systems that permeate it. This can be seen from the requirements that are being identified in an activity called *immersive media access and delivery*, where four dimensions are identified:

- time (as usual)
- space (the ability to retrieve only the relevant parts of the media)
- quality (the ability to access portions of the media with the desired quality)
- object (the ability to access specific parts of specific objects of interest).

### *Is MPEG only about media?*

My answer is no: It is not and should not be only about media. During the past 30 years, MPEG has shown that it has been able to add domains of expertise and learn to speak their language. The fact that all of the media domains today

speak the same (technical) language is due in no small part to the efforts made by MPEG to understand the needs of different industries, convert them to requirements, develop the technologies, and quantize the standards into

profiles and levels.

This workflow has been in operation for 27 years, namely, since MPEG invented profiles and levels and consistently applied them to talk to different communities using the same

language. This does not mean that there are no challenges in talking to a new industry. MPEG spent more than three years speaking to the genomic community and identifying and validating requirements before starting the development of the MPEG-G standard.

### *The business context is changing*

MPEG does not have much to say on business matters, but its work is highly influenced by companies; for example, in terms of requirements, expert support, timelines, and so on. During the first 10 years of its existence, MPEG ruled the field of audio and video coding. Through the following 10 years, some proprietary solutions popped up, but the arena was still largely dominated by MPEG standards. During the past 10 years, MPEG has seen proprietary solutions gain more strength and occupy areas that its standards used to cover exclusively. The market spoke, and it was right. MPEG should not complain about competition. Operating in a competitive environment is always beneficial because it pushes those affected to do a better job. However, this is the easy part of the story because the question is, “Should MPEG continue undeterred on its way, or should it rethink its role?”

## Conclusions

For decades, MPEG has managed a large WG, designed multithreaded work programs, sought and received support from industry, developed standards, and driven the evolution and relentless expansion of the digital-media business. MPEG did all that as a WG; i.e., the lowest-level ISO organizational unit, which is recommended to be “reasonably limited in size.” The reality is that, in 1989, MPEG had 100 members; in 1999, it had 200; and, in 2019, it had 1,500, with 600 experts attending its quarterly meetings. MPEG has played a role much above its status, and I do not think there should be complaints about the results. MPEG can continue as a WG for another 30 years, but, from now on, it should have a status equivalent to the position it holds, considering that the ISO has tolerated the “MPEG exception.” Now, there should be no exception.

Should MPEG be split in smaller pieces? Who would replace a successful organization with an untested one? Who would take responsibility if the new one failed? Something more effective is needed. The Italian National Standards Institute has proposed leveraging the strong MPEG identity, organization, and brand in a new setting that offers solid governance while preserving MPEG’s ability to develop and execute strategic plans that serve the needs of the industry. By the time this article is published, we will hopefully know if MPEG has a future.

## Author

**Leonardo Chiariglione** (leonardo@chiariglione.org) conceived, founded with Hiroshi Yasuda, and chairs the Moving Picture Experts Group, the International Organization for Standardization/International Electrotechnical Commission working group that has brought digital audio and video standards to the media industries. He is the chief executive officer of CEDEO.net.

# The Challenge: From MPEG Intellectual Property Rights Ontologies to Smart Contracts and Blockchains

**T**he Moving Picture Experts Group (MPEG) is an International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) working group that develops media coding standards. These standards include a set of ontologies for the codification of intellectual property rights (IPR) information related to media. The Media Value Chain Ontology (MVCO) facilitates rights tracking for fair, timely, and transparent payment of royalties by capturing user roles and their permissible actions on a particular IP entity. The Audio Value Chain Ontology (AVCO) extends MVCO functionality related to the description of IP entities in the audio domain, e.g., multitrack audio and time segments. The Media Contract Ontology (MCO) facilitates the conversion of narrative contracts to digital ones. Furthermore, the axioms in these ontologies can drive the execution of rights-related workflows in controlled environments, e.g., blockchains, where transparency and interoperability is favored toward fair trade of music and media. Thus, the aim of this article is to create awareness of the MPEG IPR ontologies developed in the last few years and the work currently taking place addressing the challenge identified toward the execution of such ontologies as smart contracts on blockchain environments.

## Background

### Motivation

Copyright legislation has continuously evolved with the aim of reviving the music industry in terms of fair and increased revenues returned to artists and rights holders,

improved multiterritory licensing, timely payments, and overall, more transparency, e.g., the United States Music Modernization Act [1] and the European Union's Copyright Directive Reform [2]. Meanwhile, several key artists and musicians have turned their hopes for resolving these issues to technology and, in particular, blockchain [3], [4].

Blockchain emerged in 2008 as the technology that underpins bitcoin. It operates as a shared ledger that continuously records transactions or information. Its database structure, where there is a timestamp on each entry and information linking it to previous blocks, makes it not only transparent but exceptionally difficult to tamper with.

Initiatives investigating blockchain have been launched around the world. In the United States, the Open Music Initiative (OMI) [3] has been launched by the Berklee Institute for Creative Entrepreneurship, harnessing the expertise of the Massachusetts Institute of Technology Media Lab, in decentralized platforms, whose mission is to promote and advance the development of open source standards and innovation related to music and to help ensure proper compensation for all creators, performers, and rights holders of music. OMI's focus is 1) on new works, rather than the vast legacy music catalog, with the aim that the same principles can be applied to legacy music retrospectively; and 2) on achieving interoperability among infrastructures, databases, and systems so they can be accessed, shared, and exchanged by all stakeholders.

In Europe, one of blockchain's evangelists is the Grammy-award-winning U.K. singer, songwriter and producer Imogen Heap. She has launched a blockchain project, Mycelia [4]. Although still in its early stages, she intends Mycelia to

be an entire ecosystem that uses blockchain as a way to shake up the music industry. Mycelia's mission is to

- 1) empower a fair, sustainable, and vibrant music industry ecosystem involving all online music interaction services
- 2) unlock the huge potential for creators and their music-related metadata so an entirely new commercial marketplace may flourish
- 3) ensure that all involved are paid and acknowledged fully
- 4) set commercial, ethical, and technical standards to exponentially increase innovation for the music services of the future
- 5) connect the dots with all those involved in this shift from our current outdated music industry models while exploring new technological solutions to enliven and improve the music ecosystem.

Such missions can be accomplished thanks to MPEG IPR ontologies, which can be used by music and media value chain stakeholders to share and exchange all metadata and contractual information connected to creative works, in a standardized and therefore interoperable way, leading to transparent payment of royalties and reduced time spent searching for the right data. The latter is due to inference and reasoning capabilities inherently associated with ontologies. That is, knowledge and data can be derived by evidence (facts) and logic based on rich semantic copyright models expressed by MPEG IPR ontologies. In this way, the data derived are unambiguously interpretable, facilitating efficient processing in business-to-consumer and business-to-business (B2B) music and media value chains.

However, while enthusiasm is growing for blockchain, it is likely to be

several years before we see it rolled out in a wide-scale, mainstream capacity. Blockchain enables value to be transferred over the Internet. For contractual music and media asset trading, smart contracts can be used to encode the terms and conditions of a contract. They validate contractual agreements between stakeholders before a blockchain value transfer is enabled [5]. In other words, smart contracts, implemented via software, could allow music and media royalties to be administered almost instantaneously and manage usage allowances and restrictions.

Rather than passing through intermediaries, revenue from a stream or download could be distributed automatically to rights holders, according to agreed terms and conditions (e.g., splits), as soon as an asset is downloaded or streamed [6], [7].

That is, while various smart-contract solutions abound, it is likely that the technology will really only take off once there is a clear consensus in business about which standards will prevail [8]. So the challenge that naturally arises is as follows. How can MPEG IPR standardized ontologies be converted to smart contracts that can be executed on existing blockchain environments, thus enriching blockchain environments with inference and reasoning capabilities inherently associated with ontologies? Note that this process will increase trust among music and media value chain stakeholders for sharing data in the ecosystem since the data will be cryptographically secured and verified by a blockchain.

From the other side, while plenty of research literature deals with semantic-level interoperability of ontologies (linking different ontologies) and protocol-level interoperability of blockchains (transferring verified data from one to another), the interoperability gap between them has not yet been sufficiently bridged [9]. Toward this direction, MPEG is not going to develop any

blockchain-based technology or any new language for smart contracts. However, in the last few years MPEG has developed MPEG IPR ontologies, which facilitate the conversion of narrative contracts to digital ones. Thus, MPEG's aim is to further develop the means

(e.g., protocols and application programming interfaces) for converting MPEG IPR ontologies to smart contracts executable on existing blockchain environments. In that way, MPEG is going to close the interoperability gap between MPEG IPR ontologies (and consequent-

ly the Semantic Web) and blockchains.

Last but not least, a standards-based fair and sustainable trade of music and media ecosystem is envisaged [10] based on widely deployed MPEG technologies (e.g., audiovisual codecs, file formats, and streaming protocols) [11], including emerging MPEG IPR ontologies executed as smart contracts on blockchain environments.

### Issuing body

MPEG, officially known as *ISO/IEC JTC1/SC29/WG11*, is a working group of Standardization Subcommittee 29 of the Joint Technical Committee 1 of the ISO and the IEC, that develops and facilitates international standards within the field of audio, picture, multimedia, and hypermedia information coding.

In the last few years, MPEG has developed a number of standardized ontologies catering to the needs of the music and media industry with respect to codification of IPR information toward the fair trade of music and media. These MPEG IPR ontologies have been developed using the World Wide Web Consortium's (W3C's) Resource Description Framework (RDF), under the MPEG-21 Multimedia Framework (ISO/IEC 21000) family of standards, and include *MVCO* (ISO/IEC 21000, Part 19), its extension with

respect to multitrack audio and time segments, known as *AVCO* (ISO/IEC 21000, Part 19/Amendment 1) and *MCO* (ISO/IEC 21000, Part 21). With respect to the latter, an equivalent standard has also been developed using W3C's XML, known as *Contract Expression Language* (ISO/IEC 21000, Part 20). Next, the aforementioned MPEG IPR ontologies are described. Terms in italics are further defined in the standards given in "Resources."

## Technology

### The MVCO

#### Main entities

The MVCO [12] is an ontology that formalizes the media value chain. The MVCO was designed to satisfy a number of requirements, which in turn led to defining three entities of top importance: *IP entities*, as they are transformed along their life cycle, relevant *actions* that can be performed on such entities, and types of *users* whose actions are rights, obligations, or something else foreseen by IP law.

IP entities are objects (e.g., work, manifestation, instance, product) in the media value chain, subject to protection by copyright law. The very first entity in the chain is the abstract creation, the *work*, which is the result of any intellectual endeavor with enough creativity. Works are pure, abstract entities with no material incarnation whatsoever. Derivative works are special types of works derived from an existing work. Works are fixated into physical *manifestations*, which are the very first incarnation of works. Manifestations can be *instanced* and *copied*, or they can be transformed into commercial *products*. Whereas the logical schema of IP entities resembles the Functional Requirements for Bibliographic Records (FRBR) chain [13], the source is somewhat different: the MVCO, catering to the needs of music and media stakeholders, codifies the IP entities mentioned by copyright legislation (as defined by worldwide treaties, such as the Berne Convention), whereas the FRBR is inspired by the needs of librarians.

A *user* is defined as an individual or organization acting in the media value chain. The types of roles a user could undertake revolve around the IP entities, e.g., a *creator* is defined as the user who creates a work; an *adaptor* is the user who adapts a work to produce an adaptation. These roles or very similar ones are also acknowledged by copyright legislation. Other roles include producer, distributor, and, finally, the end user.

The types of actions that can be performed also revolve around the IP entities. *Create work* is the action whose result is a new work, *produce* is the action whose result is a product, and so forth. In addition, some other actions do not produce any new IP entity. Such actions include a public communication or an end-user action (e.g., play and print), but they are legal concepts with explicit mentions and provisions in copyright legislation.

The relationship between a user and a particular IP entity type (e.g., work, adaptation, product, copy) is specified through the concept of *role*. The actions that a user performs on a given IP entity determine the role of that user with respect to the IP entity in question. Users get roles (e.g., creator, adaptor, producer, end user) that attribute them rights over actions (e.g., create work, make adaptation, produce, distribute, synchronize) that can be exercised on specific IP entities. Any given user may undertake any number of roles within a given value chain. Figure 1 illustrates these relationships among actions, users, and IP entities.

#### Authorization model

The MVCO, by defining the relationships between users, actions, and IP entities, serves well to depict a static picture of the IP information. However, in real life, rights are transferable and the MVCO needed to support this dynamic nature of rights.

The transfer of rights are authorized by signatures on agreements or contracts that grant *permissions*. A permission relates an IP entity to a right in transit between the original rights owner and the new rights owner. Permissions have an intrinsic dynamic nature: they are

granted, invoked, and revoked. Instances of a user class will probably be actual companies or persons; instances of works will be actual works. However, instances of permissions are far more interesting because they could refer either to the past or the future.

That is, an instance permission (e.g., Alice's permission to play a song) would be related to both an end-user instance (e.g., Alice) and an action instance (e.g., play a song). However, what is the interpretation of an action instance? It might be an action effectively executed in the past (e.g., Alice played a song), but it might also be an action to be performed in the future, as a mere possibility (e.g., Alice can play a song). This is commonly referred in the literature as *event factuality* and suggests that action instances can be marked as executed acts or as possible acts.

Permissions can also be granted conditionally, that is, subject to certain conditions (*facts*). Facts can be seen as propositions with an alethic (e.g., true or false) value. These propositions can be combined with logical operators (e.g., conjunction and disjunction) to create more complex conditions. The evaluation of conditions against a certain context would determine whether a permission would actually be granted or not. In such a context, permissions can also be expressed as prohibitions (negation of a permission) and obligations (the prohibition of not doing something).

Finally, the MVCO supports to some extent the so-called *copyright exceptions*, a notion present in IP law to enable the reasonable use of copyrighted assets in certain cases. For example, complete quotes are allowed for scientific purposes, and parody is also permitted. The MVCO provides mechanisms for specifying such copyright exceptions,

although the exceptions themselves are not specified.

#### The AVCO

The AVCO facilitates transparent IPR management even when content reuse is involved. This relates in particular to widespread adoption of interactive music services (remixing, karaoke, and collaborative music creation) enabled by MPEG-A: Interactive Music Application Format [14], also known as *Stem* [15], which raises the issue of rights monitoring when reuse of audio IP entities is involved, such as tracks or even segments of tracks in new derivative works.

AVCO addresses this issue by extending MVCO functionality related to the description of composite IP entities in the audio domain, whereby the components of a given IP entity can be located in time and, in the case of multi-track audio, in association with specific tracks. To do so, AVCO introduces, as shown in Figure 2, the concepts of

- 1) *timeline* [16]: a linear and coherent piece of time in relation to time-based IP entities, e.g., a vocal track can be associated with such a timeline
- 2) *interval*: a temporal entity defined by a start and end points on a given timeline, e.g., the chorus interval of a vocal track
- 3) *segment*: a slice of an IP entity with boundaries defined by the interval's start and end points, e.g., the chorus interval's IP entity
- 4) *track*: a single track of a multitrack audio IP entity, e.g., the vocal track's IP entity.

The introduction of an additional *reuse* action enables querying and granting permissions for the reuse of existing IP entities to create new derivative composite IP entities.

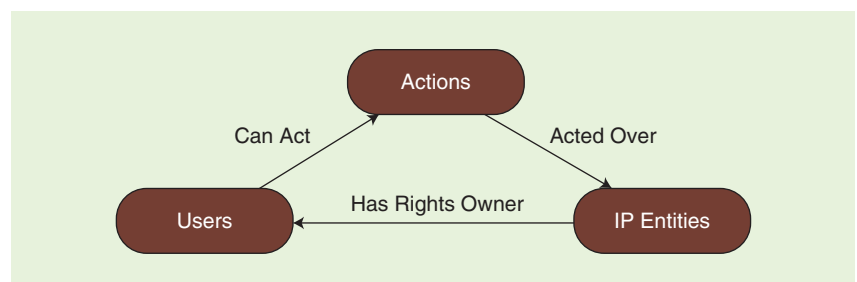


FIGURE 1. MVCO-defined relationships among actions, users, and IP entities.



## Relationships for IP entity segments and tracks

The AVCO-defined classes and relationships are illustrated in Figure 3. Since IP entities in the audio domain constitute timed media, a timeline can be associated with them. That is, an IP entity through the property *interval* is linked to an interval (Interval class instance), which in turn, through the property *onTimeLine* is

associated with a timeline. The property interval is also handy to be expressed that a segment exists within a specified interval on a timeline.

A segment is usually in a part-of relationship with an existing IP entity linked to it through the *hasSegment* property. However, a segment may also contain an IP entity different than the existing (reused) one. In either case, since a

segment is subsumed by the *IPEntity* class, it is an IP entity with its own value chain resolving to its rights holders.

In the case of multitrack audio resources, an IP entity is related to a specific track with the *hasTrack* property. To be expressed that a segment exists on a certain track, it is linked to the respective track using the *onTrack* property.


In that way, reused IP entities may exist in specified segments of existing IP entities and, in the case of multitrack audio IP entities, on specified tracks.

## The MCO

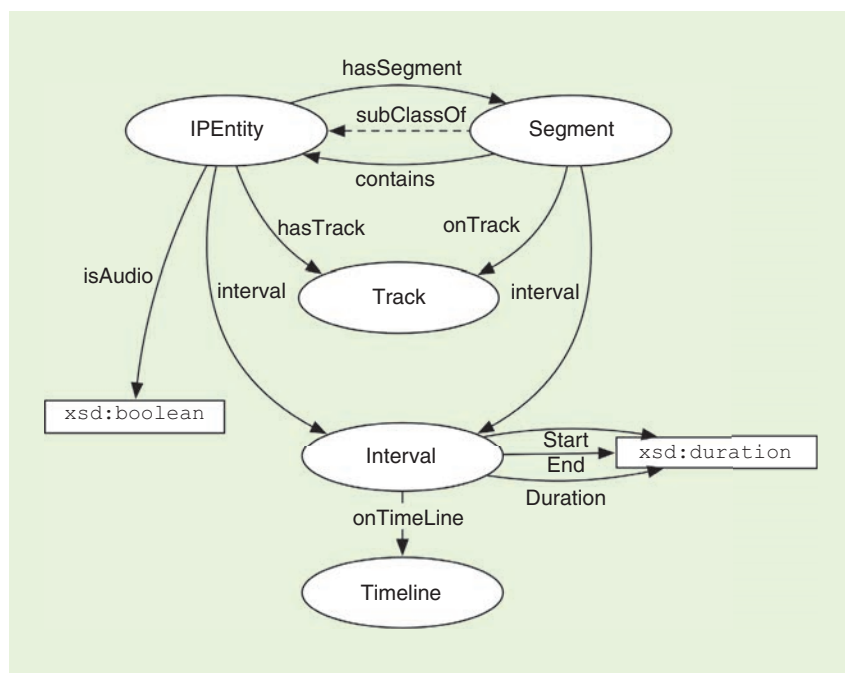
The MCO [17] facilitates the conversion of narrative contracts to digital ones and permits the creation of new contracts in machine-readable electronic formats. It consists of a core model (mco-core) and two extensions. The core model, as shown in Figure 4, builds on top of MVCO generic deontic statements (encompassing the concepts of permission, prohibition, and obligation) by providing the elements for modeling the basic structure of media contracts (e.g., contract and parties identification and relationships with other contracts). The two extensions are 1) Exploitation of Intellectual Property Rights and 2) Payments and Notifications.

## Exploitation of IPR

The extension for the exploitation of IPR (mco-ipre) provides the means to express the rights for exploiting media content, as is typical among audiovisual production companies and broadcasters. In such a context, the most commonly used rights for media exploitation are those for *public performance* (e.g., where the public is present), *fixation* (e.g., when a performance is recorded on a tangible medium), and *communication to the public* (e.g., where the public is reached by means of a communication technology). As in narrative contracts, these exploitation rights might be associated with a wide set of conditions (*facts*; e.g., number of broadcast transmissions, time periods, territories, languages, exclusivity, royalty percentages), *modalities* (e.g., linear/broadcast, nonlinear/broadband),



**FIGURE 2.** Recordings representing visualized multitrack audio. A segment exists within an interval on a timeline.



**FIGURE 3.** ACVO-defined classes and relationships for the representation of IP entities that contain other existing IP entities. Segments can also be associated with individual tracks of a multitrack audio IP entity.

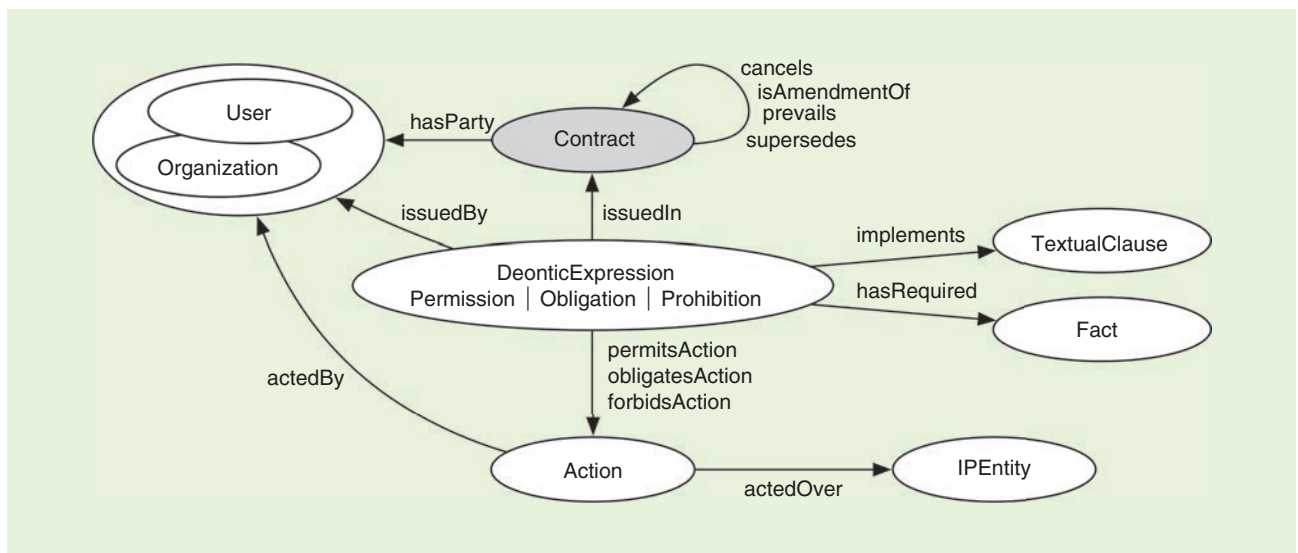


FIGURE 4. The main elements of the MCO model.

and *access policies* (e.g., free of charge, subscription, pay per view).

In the main model (mco-core), actions are permitted when the required conditions are met (e.g., the required facts are true). However, with this extension on exploitation of IPR (mco-ipre), dependencies between different actions can also be specified. That is, the occurrence of an action, such as the exploitation of a right, can trigger a condition for another action. This mechanism allows the specification of complex rights' dependencies, such as, for instance, in the so-called catch-up TV service (a combination of both linear/broadcast and nonlinear/broadband communication to the public) offered by a number of broadcasters. As an example, consider a broadcasting operator who has acquired the right from a production company to *broadcast* a TV episode. The broadcasting operator has also acquired the right to make the TV episode available on demand from its website to its subscribers via *broadband* access but only after the TV episode has been broadcast. In this case, the latter right (communication to the public via broadband) is dependent upon the use of the former communication to the public via broadcast.

#### Payments and notifications

The extension for payments and notifications (mco-pane) provides means to define specific obligations for

completing a media contract scenario. Both payments and notifications are typically obligated actions that can either be triggered by (as a consequence of) rights exploitation actions or required as a precondition to rights exploitation actions.

Eventually, the MCO can be used for the conversion of narrative media contracts to digital ones and vice versa. Such an MCO-based rights management system has been built and used by Radiotelevisione Italiana to store, access, and modify information on media rights purchased and used across its departments involved in activities ranging from media production to broadcast scheduling, improving the efficiency of media operations. Furthermore, interorganizational (B2B) rights management interoperability could be achieved by the deployment of MCO open standards by other media production companies and broadcasting operators.

#### Usage example

The MPEG IPR ontologies can be used as data models, e.g., knowledge graphs, for representing media rights. That is, actual users, media assets, and rights can be represented in RDF, instantiating MCO/MVCO/AVCO classes. The next RDF statements declare a work identified by an International Standard Musical Work Code with exploitation rights assigned to a certain PartyA:

```
:myWork a mvco:Work;  
mvco:hasRightsOwner  
  "PartyA";  
:myWork owl:sameAs  
  "T-034.524.680-C";
```

The exploitation rights on this work may be described in a contract represented using the MCO. A basic MCO contract follows. This allows the communication to the public right of the aforementioned work to be transferred from PartyA to PartyB:

```
:a Contract a mco-core:  
Contract;  
mco-core:hasParty  
  "PartyA", "PartyB";  
[  
  a mvco:Permission;  
  mvco:permitsAction mco-  
ipre:CommunicationTo  
ThePublic;  
  mco-core:issuedIn:  
aContract;  
  mco-core:actedBy "PartyB";
```

In this usage example, the joint use of terms defined in the mvco ontology (such as mvco:Permission), in mco-core (such as mco-core:Contract), and in mco-ipre (such as the communication to the public) has been shown. In practice, contracts will contain a number of restrictions and obligations (such as payments).

#### Further technical developments

MPEG IPR ontologies can be used by music and media value chain stakeholders to share and exchange in an

## Resources

### Standards

- *Information Technology—Multimedia Framework (MPEG-21), Part 19: Media Value Chain Ontology*, Standard ISO/IEC 21000-19, June 2010.
- *Information Technology—Multimedia Framework (MPEG-21), Part 8: Reference software/AMD2 Reference Software for Media Value Chain Ontology*, Standard ISO/IEC 21000-8/AMD2, Nov. 2011.
- *Information Technology—Multimedia Framework (MPEG-21), Part 19: Media Value Chain Ontology/AMD1 Extensions on Time-Segments and Multi-Track Audio*, Standard ISO/IEC 21000-19:2010/AMD1, June 2018.
- *Information Technology—Multimedia Framework (MPEG-21), Part 8: Reference Software/AMD4 Media Value Chain Ontology Extensions on Time-Segments and Multi-*

*Track Audio*, Standard ISO/IEC 21000-8:2008/AMD4, Oct. 2018.

- *Information Technology—Multimedia Framework (MPEG-21), Part 21: Media Contract Ontology*, Standard ISO/IEC 21000-21 (2nd ed.), May 2017.
- *Information Technology—Multimedia Framework (MPEG-21), Part 20: Contract Expression Language*, Standard ISO/IEC 21000-20 (2nd ed.), Dec. 2016.

### Software

- Media Value Chain Ontology: <https://tinyurl.com/y6tsr9as>.
- Audio Value Chain Ontology: <https://standards.iso.org/iso-iec/21000/-8/ed-2/en/amd/4>. (Note: Source code files provided replace the corresponding MVCO ones.)
- Media Contract Ontology: <https://standards.iso.org/iso-iec/21000/-21/ed-2>.

interoperable way all metadata and contractual information connected to creative works, leading to transparent payment of royalties and reduced time spent searching for the right data.

Such MPEG IPR ontology standards should convince music and media industry stakeholders to accept technology developments catering to the needs of music and media rights transparency built upon open standards. Related information can be found in “Resources.”

Furthermore, an MPEG ad hoc group, known as *MPEG-21 Contracts to Smart Contracts*, has recently been established to investigate and develop hooks (e.g., protocols and application programming interfaces) for converting MPEG IPR ontologies to smart contracts executable on existing blockchain environments, thus further increasing trust among music and media value chain stakeholders for sharing data in the ecosystem. In that way, the group is going to also bridge the interoperability gap between MPEG IPR ontologies (and consequently the Semantic Web) and blockchains. Though the MPEG ad hoc group is in its infancy, it has attracted a significant number of industrial and academic experts from both the semantic and the blockchain communities committed to work on the identified challenge. The work is in an exploratory phase and a publishable

working draft is expected soon. Such developments toward a *semantic music and media blockchain* have the potential to unlock both the Semantic Web and the creative economy.

### Authors

**Panos Kudumakis** (p.kudumakis@qmul.ac.uk) received his B.Sc. degree in electrical and computer engineering from the National Technical University of Athens, Greece, in 1992 and his Ph.D. degree in electronic and electrical engineering from King's College London in 1997. He is a research manager at qMedia, Queen Mary University of London. He serves as leader of the United Kingdom delegation of the International Organization for Standardization/International Electrotechnical Commission Joint Technical Committee 1/Standardization Subcommittee 29 international committee and as chair of the British Standards Institution IST/37 national mirror committee, which he joined in 1998. He is also a chair of the Moving Picture Experts Group (MPEG) ad hoc group on MPEG-21 contracts to smart contracts. His research interests include digital audio signal processing, music information retrieval, interactive media formats, intellectual property rights, middleware architectures, and standardization. He is a Senior Member of the IEEE.

**Thomas Wilmering** (t.wilmering@qmul.ac.uk) received his B.Sc. degree in music technology from London Metropolitan University in 2005, his M.Sc. degree in music technology from the University of York, United Kingdom, in 2006, and his Ph.D. degree in electronic engineering from the Queen Mary University of London in 2014, where he is currently a postdoctoral researcher in the Centre for Digital Music. His research interests include semantic audio and linked data technology. He is a coauthor of the Audio Value Chain Ontology.

**Mark Sandler** (m.sandler@qmul.ac.uk) received both his B.Sc. and Ph.D. degrees in electronic engineering from the University of Essex, Colchester, United Kingdom, in 1978 and 1983, respectively. He is a professor of signal processing and director at the Centre for Digital Music, Queen Mary University of London. He has supervised more than 40 Ph.D. students, and he is a principal investigator of the large Engineering and Physical Sciences Research Council grant (EP/L019981/1) on fusing semantic and audio technologies for intelligent music production and consumption. His research interests include music signal processing, music informatics, audio processing, semantic audio, and Semantic Web for audio. He is a Fellow of the IEEE and the Royal Academy of Engineering.



**Víctor Rodríguez-Doncel** (vrodri-guez@fi.upm.es) received his B.Sc. degree in telecommunication engineering from the University of Valladolid, Spain, in 2001 and his Ph.D. degree in computer science from the Technical University of Catalonia, Barcelona, Spain, in 2010. He is an assistant professor in the Artificial Intelligence Department at the Technical University of Madrid, Spain. He is conducting research within the Ontology Engineering Group in different areas of the Semantic Web, language technologies, and law. He is a coauthor of the Media Value Chain Ontology.

**Laurent Boch** (laurent.boch@rai.it) received his M.Sc. degree in electronic engineering from the Polytechnic University of Turin, Italy, in 1990. He is a manager of research projects in the Centre for Research and Innovation at Radiotelevisione Italiana, Turin, Italy. His research interests include the evolution of audiovisual archives. He is a coauthor of Media Contracts Ontology and Contracts Expression Language.

**Jaime Delgado** (jaime.delgado@ac.upc.edu) received both his M.Sc. and Ph.D. degrees in telecommunications engineering from the Technical University of Catalonia, Barcelona, Spain, in 1983 and 1987, respectively. He is a professor at the Technical University of Catalonia, Barcelona, Spain. Since 1989, he has been a frequent contributor to standards, including Moving Picture Experts Group (MPEG) intellectual property rights ontologies, and he has taken on various leading roles in standardization activities, including with the MPEG, the European Telecommunications Standards Institute, and the European Federation for Medical Informatics. His research interests include privacy and security, eHealth and genomic information, metadata, digital rights management, multimedia information, distributed applications, and standardization.

## References

- [1] U.S. House of Representatives, 115th Congress. (2018, Apr. 10). *H.R.5447: Music Modernization Act*. [Online]. Available: <https://www.congress.gov/bill/115th-congress/house-bill/5447/all-info>
- [2] European Commission, "Modernization of the EU copyright rules," 2018. [Online]. Available: <https://ec.europa.eu/digital-single-market/en/modernisation-eu-copyright-rules>
- [3] Rethink Music. "Fair music: Transparency and payment flows in the music industry," BerkleeCE,

2015. [Online]. Available: [www.rethink-music.com/research/fair-music-transparency-and-payment-flows-in-the-music-industry](http://www.rethink-music.com/research/fair-music-transparency-and-payment-flows-in-the-music-industry)

- [4] I. Heap, "Mycelia for music." Accessed on: Dec. 9, 2019. [Online]. Available: <http://myceliaformusic.org/>
- [5] M. E. Peck, "Blockchains: How they work and why they'll change the world," *IEEE Spectr.*, vol. 54, no. 10, pp. 26–35, Oct. 2017. doi: 10.1109/MSPEC.2017.8048836.
- [6] M. O'Dair, Z. Beaven, D. Neilson, R. Osborne, and P. Pacifico, "Music on the blockchain," Middlesex Univ., London, July 2016. [Online]. Available: [https://www.mdx.ac.uk/\\_data/assets/pdf\\_file/0026/230696/Music-On-The-Blockchain.pdf](https://www.mdx.ac.uk/_data/assets/pdf_file/0026/230696/Music-On-The-Blockchain.pdf)
- [7] B. Rosenblatt, "The future of blockchain technology in the music industry," *Entertainment Sports Layer*, vol. 35, no. 1, pp. 12–20, 2019.
- [8] K. Cagle, "Rise of the smart contract," *Forbes*, Mar. 2019. [Online]. Available: <https://www.forbes.com/sites/cognitiveworld/2019/03/10/rise-of-the-smart-contract/>
- [9] H. M. Kim, M. Laskowski, and N. Nan, "A first step in the co-evolution of blockchain and ontologies: Towards engineering an ontology of governance at the blockchain protocol level," *SSRN Electron. J.*, Jan. 6, 2018. doi: 10.2139/ssrn.3097443.
- [10] P. Kudumakis, M. Sandler, A.-C. G. Anadiotis, I. S. Venieris, A. Difino, X. Wang, G. Tropea, M. Grafl et al., "MPEG-M: A digital media ecosystem for interoperable applications," *Signal Process. Image Commun.*, vol. 29, no. 1, pp. 150–166, Jan. 2014. doi: 10.1016/j.image.2013.10.006.

- [11] P. Kudumakis, T. Wilmering, M. Sandler, and J. Foss, "MPEG intellectual property rights ontologies for media trading and personalization," in *Proc. 1st Int. Workshop Data-Driven Personalization of Television (DataTV'19)*, Manchester, U.K. 2019. [Online]. Available: [http://ceur-ws.org/Vol-2423/DataTV2019\\_paper\\_2.pdf](http://ceur-ws.org/Vol-2423/DataTV2019_paper_2.pdf)
- [12] V. Rodríguez-Doncel and J. Delgado, "A media value chain ontology for MPEG-21," *IEEE Multimedia Mag.*, vol. 16, no. 4, pp. 44–51, Oct.–Dec. 2009.
- [13] B. Tillett, "What is FRBR? A conceptual model for the bibliographic universe," *Australian Library J.*, vol. 54, no. 1, pp. 24–30, 2005. doi: 10.1080/00049670.2005.10721710.
- [14] I. Jang, P. Kudumakis, M. Sandler, and K. Kang, "The MPEG Interactive Music Application Format standard," *IEEE Signal Process. Mag.*, vol. 28, no. 1, pp. 150–154, Jan. 2011. doi: 10.1109/MSP.2010.939073.
- [15] M. Le Goff, C. Carrier, and S. Walker, "Introducing Stem: A new multichannel audio format," in *Proc. Int. Society for Music Information Retrieval Conf. (ISMIR)*, 2015. [Online]. Available: <http://ismir2015.uma.es/LBD/LBD39.pdf>
- [16] Y. Raimond, S. A. Abdallah, M. Sandler, and F. Giasson, "The music ontology," in *Proc. Int. Society for Music Information Retrieval Conf. (ISMIR)*, 2007, pp. 417–422. [Online]. Available: <http://cite.seerx.ist.psu.edu/viewdoc/download?doi=10.1.1.173.5403&rep=rep1&type=pdf>
- [17] V. Rodríguez-Doncel, J. Delgado, S. Llorente, E. Rodríguez, and L. Boch, "Overview of the MPEG-21 media contract ontology," *Semantic Web*, vol. 7, no. 3, pp. 311–332, Mar. 2016.

SP



The IEEE Foundation is raising awareness, creating partnerships, and generating financial resources needed to combat global challenges.

**Our goal is to raise \$30 million by 2020.**

**DONATE NOW**  
**[ieeefoundation.org](http://ieeefoundation.org)**





## In Memory of A.H. “Steen” Gray Jr.

**M**y brother, Augustine Heard “Steen” Gray Jr., died at home in Florence, Oregon, on 28 October 2019. From the late 1960s through the early 1980s, we shared several research interests, which gave me a front seat to and a peripheral role in what seemed to me a golden age of digital speech research. Our professional and personal lives were entangled, so I hope I may be forgiven for my natural bias in these reflections on my admiration and affection for Steen as a brother, teacher, colleague, writer, friend, and remarkable human being.

Steen was born in Long Beach, California, in 1936 to Commander Augustine H. Gray, United States Navy, and Elizabeth Jordan Gray. As a Navy brat, he moved several times, living in Honolulu, Hawaii; Buenos Aires, Argentina; and Newport, Rhode Island before settling in Coronado, California, in 1942. Steen graduated from Coronado High School in 1953, when he was given a Bank of America Science and Mathematics Award, a hint of his future interests. In 1959, Steen received his S.B. and S.M. degrees in electrical engineering at the Massachusetts Institute of Technology (MIT), where he was part of the electrical engineering cooperative program (Course VI-A), working for General Electric several terms. His final project was logic design for ERMA, the first major banking computer.



A.H. “Steen” Gray Jr.

Following his graduation from MIT, Steen spent a year as a physics instructor at San Diego State College, teaching electronics and digital computer courses. I was a senior in high school that year and Steen let me audit his digital computer class, where I learned Boolean algebra, logic gates, and simple logical design. He was a natural teacher, and I learned that technical topics did not have to be boring. Steen also infected me with some of his musical tastes, most notably the Weavers and Pete Seeger, along with other folk music and blues gaining popularity in late 1950s California. Steen took me to a Pete Seeger concert, where I first saw and heard a 12-string guitar. I soon owned a Stella 12 string (sold by Sears).

Steen met, courted, and married Mary Averill Forneret of Montréal, Canada, a graduate of McGill University. Steen decided to continue his graduate work, but his enjoyment of his mathematics courses at MIT led him to apply to the mathematics department instead

of electrical engineering at California Institute of Technology (Caltech). He was turned down due to his lack of an undergraduate degree in math, but he was offered a compromise: admission to the Ph.D. program in the Department of Engineering Science, essentially an Applied Mathematics Department in the Division of Engineering. Steen accepted the offer, and he and Averill moved to Pasadena. In 1963, Steen and Averill’s son Augustine “Tino” Heard Gray III was born.

In 1964, Steen received his Ph.D. degree from Caltech, and his dissertation “Stability and Related Problems in Randomly Excited Systems” provided more hints of his future work. He was appointed assistant professor of electrical engineering at the University of California, Santa Barbara (UCSB), in July, the same month and year that the department was created. Steen and Averill moved to Santa Barbara and quickly immersed themselves in local activities, including the Get Oil Out (GOO) movement and the Santa Barbara Zoo. They made almost annual three-week voyages in a large (often unreliable) recreational vehicle to visit Averill’s sister on Hornby Island off Vancouver Island, passing through and falling in love with Florence, Oregon, on their way.

As one of the founding faculty members of the Electrical Engineering Department, Steen helped to build its academic program. He taught courses in circuits, numerical analysis, linear systems, transform techniques, digital signal processing, and complex variables.

Steen and I were in fairly frequent contact from 1966 to 1969 while I was at USC getting my Ph.D. degree. I recall him talking about his interests in discrete-time linear filters and Fourier analysis, and we chatted about my initial efforts to find a dissertation project in Shannon rate-distortion theory that focused on simple examples of autoregressive sources. Steen suggested using linear difference equations to solve for the eigenvalues of an almost Toeplitz inverse autocorrelation matrix, which solved my special case and got me started.

We visited Acres of Books in Long Beach, which was then the largest and oldest family-owned bookstore in California. As with music, brotherly osmosis resulted in my acquiring many of his literary tastes of the time, including late 19th-century adventure fiction and pulp fiction collected and published by Arkham House. On occasion, Steen and Averill would drive to Manhattan Beach for parties my bandmates and I threw, including our electrified versions of familiar folk songs and blues.

In 1968, Steen was promoted to associate professor, and John D. Markel arrived at the Speech Communications Research Laboratory (SCRL). John began applying the digital filtering and  $z$  transform expertise he had acquired at Arizona State and Motorola to speech processing. John's salary and an eventual doctoral study grant allowed him to pursue a Ph.D. degree at nearby UCSB, so he began looking for a dissertation topic. He first approached Glen Culler, who was famous for his computer networking and his "online" signal processing system networking graphic displays, large computers, and keyboards to provide a real-time signal processing laboratory in the classroom. Culler was working on a variation of the analysis/synthesis approach to voice coding (vocoding), which encoded speech by fitting a parametric model to observed speech segments (analysis), quantizing the parameters for transmission, and reconstructing the approximate model at the receiver to synthesize the speech. Culler used sinusoids modulated by Gaussian wave functions, later called *Gabor wavelets*, as synthesis mod-

els. His analysis involved algorithms (known only to himself), which took advantage of the extraordinary computing power he commanded. Culler turned John down flat, but John decided to persevere on the topic and began to look for technical help.

John met Steen during the math portion of his Ph.D. entrance exam and afterward approached Steen for the help that Culler had denied. In John's words, "This is where I went to Steen and realized that he was a true genius with specialized skill in closed-form solutions. From there it's a great story of Steen's generosity and ability and how we ended up complementing each other for 30 some IEEE articles."

In a 2003 email, Steen related his view of his early interactions with John:

He had moved into more practical & pragmatic speech research but was lacking in math. I liked applying math to anything and we had a lot of fun. The digital filtering stuff (lattice, ladder, normalized) dropped mostly out of playing with the math and limitations of fixed point arithmetic. The elliptic filter (unrelated to speech) was because somebody was selling programs to do it and I thought it would be nice to create one and put one in the public domain—besides which it got us a paper.

Steen joined John's dissertation committee, and John brought up a new project at SCRL. Hisashi Wakita had been very impressed by a December 1966 report from Nippon Telegraph and Telephone in Japan that had been written by a young Japanese doctoral student at Nagoya University named Fumitada Itakura and his Ph.D. supervisor Shuzo Saito [1]. Wakita had worked with Saito and was one of the few speech researchers in the United States who could read the Japanese report. Like Culler's approach, it was an analysis/synthesis vocoder, but Itakura used stationary Gaussian autoregressive processes as the synthesis model and fit a model to the observed data by approximating a maximum likelihood (ML) estimate of a parameter vector, specifying the model based on an observed speech sample vector.

Since the models were Gaussian, the ML estimator corresponded to optimal linear prediction. They were also stationary, so the method was equivalent to fitting the Toeplitz autocorrelation or the power spectral density to sample-based estimates. There was also explicit pitch detection for determining a driving function for the autoregressive or all-pole filter for modeling voiced sounds. This approach to vocoding is an example of what would be named *linear predictive coding (LPC)* by Bishnu Atal of Bell Labs in 1969; Itakura and Atal share credit for inventing LPC. John's dissertation research suggested that Culler's approach was not promising, but the Itakura approach captivated Hisashi, John, and Steen.

Hisashi saw the connections with the physics of speech production and modeling the vocal tract with acoustic tubes, John saw the approach as designing all-pole filters with correlation matching that of observed signals, and Steen saw the problem as fitting stationary autoregressive random processes to observed signals so that the spectral densities were the best possible match. Itakura's approach was all of those things and more.

In 1969, Steen joined SCRL as a consultant (SCRL did not pay consulting rates, only ordinary employee hourly salary, and there were no benefits), and in 1971, Steen became both a consultant and summer employee at SCRL.

As the SCRL work proceeded, a series of reports and articles reported advances in filter design, public domain heavily commented Fortran algorithms for speech processing and coding, relations of LPC to new and classic formulations of speech as well as to other fields, and an analysis of quantization effects. These pieces combined to provide a detailed description and prescription for vocoder design consistent with the Itakura formulation. Standing out among the numerous manuscripts were two highly referenced 1973 SCRL reports [2], [3] and a (*the*, in my opinion) classic book on the topic, *Linear Prediction of Speech* by Markel and Gray [4], which included the distilled wisdom of the reports and many published articles.

The book was a masterpiece, describing the many approaches to the problem, the underlying theory and algorithms, linear filter design, spectral analysis, pitch extraction and fundamental frequency estimation, and Fortran implementation. The U.S. military and the National Security Agency (NSA) eventually adopted and standardized the Markel/Gray implementation of the Itakura approach to LPC for secure 2,400 bit/s digital vocoders. A decade ago, Dr. Joseph P. Campbell of Lincoln Laboratory (at NSA from 1979 to 1990) wrote of the book [11]:

In true engineering form, Drs. Markel & Gray reduced complicated mathematics to practice in the form of practical and efficient FORTRAN code while presenting a rigorous mathematical development. This rare combination of rigorous theory and code, which might have been the first in speech, made this book a remarkable contribution for its day. Markel & Gray was the first book that everyone joining the Speech Coding branch of the Communications group at NSA was expected to read. As a measure of the longevity of Markel & Gray's contributions, their code from 1976 is still in operational use today, over 30-years later, in the US Government's and NATO's standard secure voice terminals. ... Linear prediction also forms the foundation for the speech coding used in today's cellular telephones. Markel & Gray reduced linear prediction to practice and are largely responsible for its widespread dissemination and successful use to this day.

Meanwhile, in 1968, Culler had joined with Elmer Shapiro, Len Kleinrock, and Larry Roberts to complete the specification of the interface message processor (IMP), the basic node of the ARPAnet, which is the predecessor to the Internet. In 1969, UCSB became the third node of the ARPAnet, and Culler founded Culler-Harrison Inc. (CHI), an early developer of array processors. Both CHI and SCRL had dedicated phone lines connecting to the UCSB IMP and

SCRL was involved with real-time signal processing and operating systems on networks from the beginning.

By 1972, it was apparent to Robert Kahn at ARPA that the network speed constraints were insufficient to provide real-time communication of decent speech quality. Kahn discussed his ideas with Danny Cohen, who was then a Ph.D. student at Harvard University, Cambridge, Massachusetts, working on real-time flight simulations across a network. Danny knew about real-time signal processing and networks and knew he could get up to speed quickly on speech coding. At Kahn's suggestion, Danny moved to USC's Information Sciences Institute (ISI) in Marina del Rey, California, and in 1973, he became chair of the new Network Secure Communications (NSC), charged with exploring the possibilities of packet speech on the ARPAnet. Joining ISI on the NSC were SCRL, SRI, BBN, CHI, the University of Utah, and MIT Lincoln Lab. ARPA provided funding and active participation. Amazingly, this disparate group successfully demonstrated the first real-time understandable quality packet speech in an experiment between CHI and MIT Lincoln Lab over the ARPAnet at 3.5 kbit/s in December 1974. The speech-coding component was the Markel and Gray implementation of the Itakura LPC algorithms.

The demonstration influenced the separation of the Internet Protocol (IP)



John and Steen at SCRL, c. 1974, with Larry Pfeiffer and PDP-11 in the background. This photo was shown in *IEEE Acoustics, Speech, and Signal Processing Newsletter*, number 35, November 1975.

from the Transmission Control Protocol and laid the groundwork for voice over IP. The story is told in detail in [9] and [10] along with citations of primary sources. The underlying foundational work of John and Steen won them an IEEE Acoustics, Speech, and Signal Processing (ASSP) Senior Award, election to IEEE Fellow, and an ASSP Technical Achievement Award.

I was only dimly aware of this story as it was unfolding. I had occasional conversations with Steen about quantization and speech compression, but I only got seriously involved with Steen and John as a latecomer in 1975, when John noted that many algorithms had been published that claimed optimal quantization of LPC parameters and complained that it was difficult to interpret what *optimal* meant when so many approaches claimed the attribute.

Steen and I took this as a great excuse to work together and began to actively collaborate on research for the first time, looking primarily to the classic theoretical approach to quantization of William R. Bennett and Stuart Lloyd at Bell Labs. The project was great fun and led to a thorough analysis and comparison of many of the popular techniques for scalar quantization of various parametric representations of LPC models [5]. As a side benefit, Steen and I discovered simple proofs of the classic asymptotically optimal scalar quantization approximations, which were published in *IEEE Transactions on Information Theory* in the same month. Publishing simultaneous papers with my brother was a high point in my career.

Scalar quantization of parameter vectors is suboptimal in the Shannon sense and practically the overall quality can be significantly affected by which parameterization is chosen. So Steen and I thought that it would be interesting to try the other theory of quantization and compression, Shannon rate-distortion theory, on speech coding (an idea first suggested by Chaffee and Omura in 1974). Two of my students (Andrés Buzo and Yasuo Matsuyama) became involved with Steen, John, and me in developing algorithms for vector quantization design for LPC speech models using Lloyd's

Method I with Saito and Itakura's error matching measure, which we called the *Itakura-Saito distortion*. About this time, Steen suggested using amateur (ham) radio to provide more opportunities to chat about our joint speech research. He had an Amateur Extra class license (AA6H) and I had an Advanced class license (KB6XQ), and for several years most of our technical discussions were by radio, which surprisingly worked pretty well (and drew many curious hams).

The work was disrupted in 1977 when John and Steen left SCRL to found Signal Technology, Inc. (STI) along with Larry Pfeiffer. John became president of the corporation and Steen became secretary, as well as continuing as a part-time and summer employee. John brought several of his speech-coding and related contracts with him so the speech research could continue, but STI's main product was the Interactive Laboratory System (ILS), a Fortran-based system for real-time signal processing in education and research, which combined John and Steen's speech algorithms with a variation of Glen Culler's online signal processing system. David Wong, Steen's Ph.D. student, joined STI from SCRL as a research engineer as he was finishing his dissertation.

The initial work on vector quantization of LPC speech came to fruition in 1978 with three presentations of joint work by Steen, me, John, Buzo, and Matsuyama on a 792 bit/s LPC vocoder using vector quantization of the model filter parameters along with standard pitch and voicing algorithms from SCRL and STI. In late June, Steen presented a talk, "Speech Compression and Speech Distortion Measures," at an IEEE information theory workshop at Lake Tahoe. He played tapes for the system using codebooks designed with the Lloyd clustering algorithm with the Itakura-Saito distortion on a large training set and on a test speaker not in the training set. Unfortunately, the codebook had been trained and the tapes made at the last minute; Steen and I received the tapes only the night before the talk, and they sounded awful. Steen called John, who discovered that the original training and test sequences indeed sounded awful, but our numeri-

cal results showed a good match between input and output, and the two sounded alike (both awful). Steen handled the talk brilliantly and turned a potential disaster into good entertainment.

I presented "Source Coding and Speech Compression" at the 1978 International Telemetering Conference in Los Angeles, California (the proceedings article in [6] has the technical details of the experiment), and Steen presented "Optimal Quantizations of Coefficient Vectors in LPC Speech" at the 1978 Joint Meeting of the Acoustical Society of America and the Acoustical Society of Japan in Honolulu. In October 1980, Buzo, Steen, John, and I published "Speech Coding Based Upon Vector Quantization" [7], which provided a thorough experimental performance comparison at a range of bit rates for both fully searched and binary tree-structured vector quantized LPC [7]. The paper won Steen and John a second ASSP Senior Award.

Over the following few years, the results were elaborated on and extended in many directions, but after 1982, Steen vanished from research and papers. Steen's and my final coauthored paper (along with John Shore and my student Guillermo Rebollo) appeared in late 1981 in *IEEE Transactions on Information Theory*. In October 1982, Steen's paper with his former students David Wong and Biing-Huang (Fred) Juang provided a detailed development of an operational 800 bit/s vector quantized LPC vocoder [8] with many improvements, based on an STI project funded by the Naval Research Laboratory. This was Steen's final research publication.

Steen's abrupt disappearance from academia and research at the young age of 46 is attributable to the tragic early death in late 1979 of his and Averill's only child, Tino. Steen retired from UCSB in 1980 and moved to STI full time, where he held successive positions as senior scientist, vice president, and executive vice president through 1988. From 1988 to 1993, he held various positions with derivative companies of STI.

In 1991, Steen and Averill retired to Florence, Oregon, where they became active volunteers for the Oregon Coast

Humane Society, for whom Steen created and managed their computer network and later their webpages. Steen and Averill adopted a succession of rescue cats and dogs over the years.

Until Averill's death in 2016, Steen and I were in regular email contact; neither of us much liked telephone conversations, and we both typed quickly. Most of the conversations were about family, gadgets (computers, scanners, recorders, audio software and players, abacuses), music, and BBC comedy. We did not discuss technical issues, but we did discuss technical history, as we had discussed the history of the NSC ARPAnet packet speech project when I was preparing the talks that led to [9] and [10].

Steen's main occupation for several years was as caregiver for Averill. I last visited him in Florence in summer 2018, and he was mobile, still driving, and living independently. But he was nearly blind, no longer read, could not walk much, missed Averill terribly, and had little interest in activity. In the last email I received from him in March 2019, he said he rarely did email anymore and that "I've aged rather a bit in the past few years." His personal computing network at home as well as at the shelter office was ancient and obsolete; it was a personal challenge for him to keep it all working. I suspect that toward the end he had the same view of himself.

Steen is remembered by his students as an amazingly good teacher who could make complex technical ideas clear and seemingly simple. Over the holidays in December 1960, he managed to teach me enough freshman calculus and physics to survive my first MIT semester. I'll close these reflections with a quote from David Wong. I cannot put it any better.

Steen lived a very meaningful and honorable life. He did not seek fame or fortune or glory, but has made his marks everywhere he lived and worked. He certainly left his mark in my life. We were teacher/student, colleagues, friends through the decades (1972 to 1990s) until he moved away. We worked together and enjoyed each other's company as we both made the transition from academia into business. I still think



of him often as the one teacher who had the biggest impact on my growth, development, and moral compass. There is no finer scholar, teacher, human, citizen than Steen!! I can only hope to emulate his character, because his intellect is simply beyond my reach. I will always remember him and I hope to leave the same type of memory in those who have been in my life as he has in mine.

## Author

**Robert M. Gray** (rmgray@stanford.edu) received his B.S. and M.S. degrees from the Massachusetts Institute of Technology, Cambridge, in 1966 and received his Ph.D. degree from the University of Southern California, Los Angeles, in 1969. He is the Alcatel-Lucent Technologies Professor of Communications and Networking in the

School of Engineering, Emeritus, Stanford University, California. He is a research professor in the Electrical and Computer Engineering Department, Boston University. He received the IEEE ASSP Society Senior award in 1983 with Andrés Buzo, Augustine H. Gray Jr., and John Markel for [7]. He also received the Jack S. Kilby Medal in Signal Processing and the Claude E. Shannon Award in 2008. He is a Life Fellow of the IEEE.

## References

- [1] S. Saito and F. Itakura, "The theoretical consideration of statistically optimum methods for speech spectral density," *Elect. Commun. Lab., NTT Nippon Telegraph and Telephone, Tokyo, Rep.* 3107, Dec. 1966.
- [2] J. D. Markel, A. H. Gray, and H. Wakita, "Linear prediction of speech: Theory and practice," *Speech Commun. Res. Lab., Santa Barbara, CA, Monograph No. 10*, Sept. 1973.
- [3] J. D. Markel, A. H. Gray, and H. Wakita, "Documentation for SCRL linear prediction analysis/synthesis programs," *Speech Commun. Res. Lab., Santa Barbara, CA, Nov.* 1973.
- [4] J. D. Markel and A. H. Gray Jr., *Linear Prediction of Speech*. Berlin: Springer-Verlag, 1976.

- [5] A. H. Gray Jr., R. M. Gray, and J. D. Markel, "Comparison of optimal quantizations of speech reflection coefficients," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 25, no. 1, pp. 9–23, Feb. 1977. doi: 10.1109/TASSP.1977.1162907.
- [6] R. M. Gray, A. Buzo, Y. Matsuyama, A. H. Gray, Jr., and J. D. Markel, "Source coding and speech compression," in *Proc. Int. Telemetering Conf.*, vol. 14. San Diego, CA: International Foundation for Telemetering, Nov. 1978, pp. 871–878.
- [7] A. Buzo, A. H. Gray Jr., R. M. Gray, and J. D. Markel, "Speech coding based upon vector quantization," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 28, no. 5, pp. 562–574, Oct. 1980. doi: 10.1109/TASSP.1980.1163445.
- [8] D. Y. Wong, B.-H. Juang, and A. H. Gray, Jr., "An 800 bit/s vector quantization LPC vocoder," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 30, no. 5, pp. 770–780, Oct. 1982. doi: 10.1109/TASSP.1982.1163960.
- [9] R. M. Gray, "The 1974 origins of VoIP," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 87–90, July 2005. doi: 10.1109/MSP.2005.1458295.
- [10] R. M. Gray, *Linear Predictive Coding and the Internet Protocol*. Hanover, NH: Now Publishers Inc., 2010. [Online]. Available: <https://ee.stanford.edu/~gray/lpcip/>
- [11] J. P. Campbell, private communication, Feb. 2009.



## LECTURE NOTES (continued from page 80)

### What we have learned

Continuity is a very important property of LTI systems, without which an LTI system may not be completely specified by its impulse response. For example, not every LTI system can be described by convolution, although every convolution defines an LTI system. For linear systems, the continuity implies the BIBO stability, whereas the converse is not true in general. In other words, BIBO stability does not always guarantee the stability of a linear system.

### Acknowledgments

We would like to thank the reviewers for their helpful comments and suggestions. This work was supported in part by the National Natural Science Foundation of China under grant 61801366 and in part by the Fundamental Research Funds for the Central Universities under grant xjh012019041.

### Authors

**Ming Zhang** (itjerry@163.com) received his B.S. and M.S. degrees in information and communications engineering and his

Ph.D. degree in electronic science and technology from Xi'an Jiaotong University, China, in 2008, 2011, and 2017, respectively. From 2011 to 2014, he worked at the Huawei Technologies Co., Ltd. He is currently an associate professor with the School of Information and Communications Engineering at Xi'an Jiaotong University. His research interests include array signal processing and linear system analysis. He is a Member of the IEEE.

**Xiaoming Chen** (xiaoming.chen@xjtu.edu.cn) received his B.S. degree in electrical engineering from Northwestern Polytechnical University, Xi'an, China, in 2006 and his M.S. and Ph.D. degrees in electrical engineering from Chalmers University of Technology, Gothenburg, Sweden, in 2007 and 2012, respectively. He is currently a professor with the School of Information and Communications Engineering at Xi'an Jiaotong University, China. His research interests include multiple-input, multiple-output antennas; over-the-air testing; and reverberation chambers. He is a Senior Member of the IEEE.

### References

- [1] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing*, 4th ed. Englewood Cliffs, NJ: Prentice Hall, 2006.
- [2] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*, 3rd ed. Englewood Cliffs, NJ: Prentice Hall, 2010.
- [3] M. Vetterli, J. Kovačević, and V. K. Goyal, *Foundations of Signal Processing*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [4] L. Hörmander, "Estimates for translation invariant operators in  $L^p$  spaces," *Acta Mathematica*, vol. 104, no. 1–2, pp. 93–140, 1960.
- [5] E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*. Princeton, NJ: Princeton Univ. Press, 1971.
- [6] L. Grafakos, *Classical Fourier Analysis*, 3rd ed. New York: Springer, 2014.
- [7] P. Chakravorty, "What is a signal?" *IEEE Signal Process. Mag.*, vol. 35, no. 5, pp. 175–177, Sept. 2018.
- [8] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. Philadelphia: Society for Industrial & Applied Mathematics, 2000.
- [9] E. Kreyszig, *Introductory Functional Analysis with Applications*. New York: Wiley, 1978.
- [10] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. New York: McGraw-Hill, 1976.
- [11] P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms*. Philadelphia: Society for Industrial & Applied Mathematics, 2010.
- [12] M. Zhang and A. Zhang, "The superposition principle of linear time-invariant systems," *IEEE Signal Process. Mag.*, vol. 36, no. 6, pp. 153–156, Nov. 2019.



Please send calendar submissions to:  
Dates Ahead, Attn: Samantha Walter, E-mail: [walter.samantha@ieee.org](mailto:walter.samantha@ieee.org)

## 2020

### MARCH

#### Data Compression Conference (DCC)

24–27 March, Snowbird, Utah, United States.  
General Cochairs: Michael W. Marcellin and  
James A. Storer  
URL: <https://www.cs.brandeis.edu/~dcc/Call.html>

### APRIL

#### IEEE International Symposium on Biomedical Imaging (ISBI)

3–7 April, Iowa City, Iowa, United States.  
General Chairs: Mathews Jacob and Jong Chul Ye  
URL: <http://2020.biomedicalimaging.org/>

#### International Conference on Information Processing in Sensor Networks (IPSN)

21–24 April, Sydney, Australia.  
General Chair: Kusy Brano  
URL: <http://ipsn.acm.org/2020/index.html#>

#### IEEE International Radar Conference (RADAR)

27 April–1 May, Washington, D.C., United States.  
General Chairs: Michael Zatman and Ravi Adve  
URL: <https://radar2020.org>

### MAY

#### IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)

4–9 May, Barcelona, Spain.  
General Chairs: Ana I. Pérez Neira and  
Xavier Mestre  
URL: <https://2020.ieeeicassp.org>

#### 11th International Conference on Quality of Multimedia Experience (QoMEX)

26–28 May, Athlone, Ireland.  
General Chairs: Niall Murray and Andrew Hines  
URL: <http://qomex2020.ie/>

Digital Object Identifier 10.1109/MSP.2019.2960991  
Date of current version: 26 February 2020



The 54th Annual Asilomar Conference on Signals, Systems, and Computers will be held 1–4 November in Pacific Grove, California.

#### IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)

26–29 May, Atlanta, Georgia, United States.  
General Cochairs: Geoffrey Le  
and Xiaoli Ma  
URL: <https://spawc2020.netlify.com>

### JUNE

#### IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)

8–11 June, Hangzhou, China.  
General Chairs: Martin Haardt  
and Zhiguo Shi  
URL: <http://www.sam2020.cn>

### JULY

#### IEEE International Conference on Multimedia and Expo (ICME)

6–10 July, London, United Kingdom.  
General Chairs: Marta Mark, Ebroul Izquierdo,  
and Vlado Velisavljevic  
URL: <http://www.2020.ieeeicme.org/>

#### IEEE Statistical Signal Processing Workshop (SSP)

12–15 July, Rio de Janeiro, Brazil.  
General Chair: C. de Lamare  
URL: <http://ssp2020.cetuc.puc-rio.br/>

### SEPTEMBER

#### 2020 IEEE 22nd International Workshop on Multimedia Signal Processing

21–23 September, Tampere, Finland.  
General Chairs: Atanas Gotchev and  
Dong Tian  
URL: <https://events.tuni.fi/mm2020/>

### OCTOBER

#### IEEE International Conference on Image Processing (ICIP)

25–28 October, Abu Dhabi,  
United Arab Emirates.  
General Chairs: Mohammed Al-Mualla and  
Moncef Gabbouj  
URL: <http://2020.ieeeicip.org/>

### NOVEMBER

#### 54th Annual Asilomar Conference on Signals, Systems, and Computers

1–4 November, Pacific Grove, California,  
United States.  
General Chair: Joseph R. Cavallaro  
URL: [www.asilomarssc.org](http://www.asilomarssc.org)



**FIGURE 1.** Peter Schreier speaking at the 2018 SSP Workshop.

including the Challenges and Data Collections Subcommittee, recently established by the Technical Directions Board to oversee the SPS Challenges Program, and the IEEE SPS Autonomous Systems Initiative, which promotes technical and educational activities in autonomous systems. Other activities of the SPTM TC include the generation of instructional and motivational material in the form of text and video on the practical applications of signal processing theory and the promotion of webinars, online courses, and other educational resources covering topics within the areas of the TC.

The SPTM TC actively participates in the nomination process for major IEEE and SPS awards. Some of the paper awards for which the SPTM TC systematically nominates candidates include those for Signal Processing Best Paper, Best Overview Paper, Signal Processing Magazine Best Paper and Best Column, SPS Best Paper, and SPS Young Author Best Paper. The SPTM TC has established a very careful nomination process, coordinated by the Awards Subcommittee, which involves detailed reviews of several papers nominated by TC members. For the major awards, the TC nominates candidates for the Education Award, Early Career Technical Achievement Award, SPS Award,

and Technical Achievement Award. The SPTM TC also nominates candidates for Distinguished Lecturers and Distinguished Industry Speakers.

Another important objective of the SPTM TC is to identify emerging theoretical areas that are promising for signal processing applications and promote these areas by facilitating the disclosure of relevant research in signal processing conferences and journals. One area that has been recently included in the technical activities of the SPTM TC is graph signal processing, which focuses on the development of tools for processing data defined on irregular graph domains. New theories and methods in this field include a linear discrete signal processing framework to process data sets arising from several types of networks, including social, biological, and physical. Applications, such as health monitoring, banking, traffic control, and marketing, have benefited from the new methods developed under the umbrella of graph signal processing.

More recently, the area of quantum signal processing (QSP) has been incorporated into the SPTM TC technical activities. QSP borrows from the formalism and principles of quantum mechanics to establish a new theory for signal processing with applications in areas such as quantization and sampling methods, detection,

and parametric estimation. QSP establishes a new framework for developing new signal processing algorithms by exploiting the probabilistic nature of quantum mechanics without the need to satisfy the practical requirements dictated by physical implementations based on quantum mechanics.

The SPTM TC maintains a webpage dedicated to TC activities (<https://signalprocessingsociety.org/get-involved/signal-processing-theory-and-methods/sptm-home>), where one can find ways to participate in the TC activities, become an affiliate member, and be nominated for possible election as a TC member.

## Authors

**José C.M. Bermudez** ([j.bermudez@ieee.org](mailto:j.bermudez@ieee.org)) received his B.Sc. degree in electronic engineering from the Federal University of Rio de Janeiro (UFRJ), Brazil, in 1978, his M.Sc. degree in electrical engineering from COPPE/UFRJ in 1981, and his Ph.D. degree in electrical engineering from Concordia University, Montréal, Canada, in 1985. He is a professor in the Department of Electrical Engineering, Federal University of Santa Catarina, Florianópolis, Brazil, and the Graduate Program on Electronic Engineering and Computing, Catholic



University of Pelotas, Brazil, as well as a consultant in signal processing. His current research interests include statistical signal processing, including image processing, adaptive filters, hyperspectral image processing, and machine learning. He is senior area editor of *IEEE Transactions on Signal Processing* and chair of the Signal Processing Theory and Methods Technical Committee for the IEEE Signal Processing Society.

**Mónica F. Bugallo** (monica.bugallo@stonybrook.edu) received her B.S., M.S., and Ph.D. degrees in computer science and engineering from the University of A Coruña, Spain. She is a professor of electrical and computer engineering and associate dean for diversity and outreach of the College of Engineering and Applied Sciences at Stony Brook University, New York. Her current research interests

are in the field of statistical signal processing, with an emphasis on the theory of Monte Carlo methods and their appli-

cation to different disciplines, including biomedicine, ecology, sensor networks, and finance. In addition, she has focused on science, technology, engineering, and mathematics

education, and she has initiated several successful programs to engage students at all academic stages in the excitement of engineering and research, with a focus on underrepresented groups. She is vice chair of the IEEE Signal Processing Theory and Methods Technical Committee and has served on several technical committees of IEEE conferences and workshops. She is a Senior Member of the IEEE.

**Alle-Jan van der Veen** (a.j.vanderveen@tudelft.nl) received his M.Sc. degree

(cum laude) in 1989 and his Ph.D. degree (cum laude) in 1993, both from TU Delft, The Netherlands, where he is currently a full professor in signal processing. He is the recipient of the 1994 and 1997 IEEE Signal Processing Society (SPS) Young Author Paper Award. He was chair of the IEEE SPS Signal Processing for Communications Technical Committee and the IEEE SPS Signal Processing Theory and Methods Technical Committee, editor-in-chief of *IEEE Signal Processing Letters* and *IEEE Transactions on Signal Processing*, technical cochair of ICASSP 2011, and chair of the IEEE SPS Fellow Reference Committee and the IEEE Kilby Award Medal Selection Committee. Currently, he is a member of the IEEE SPS Nominations and Appointments Committee and director of publications of the European Association for Signal Processing. He is a Fellow of the IEEE.

SP

**As its name suggests, the SPTM TC is a cross-cutting group that focuses on theory and methods for signal processing applications.**

## Technology insight on demand on IEEE.tv

Internet television gets a mobile makeover

A mobile version of IEEE.tv is now available, plus a new app can also be found in your app store. Bring an entire network of technology insight with you:

- Generations of industry leaders.
- The newest innovations.
- Trends shaping our future.

Access award-winning programs about the what, who, and how of technology today.

**Go mobile or get the app.**

[www.ieee.tv](http://www.ieee.tv)

Available on the  
App Store

Available in  
Android  
Market





## Highlights From the Signal Processing Theory and Methods Technical Committee

The IEEE Signal Processing Theory and Methods Technical Committee (SPTM TC) is one of 13 TCs in the IEEE Signal Processing Society (SPS). Its mission is to support, nourish, and lead scientific and technological development in all theoretical aspects of signal processing. The TC has 42 volunteer members who are elected among leaders from academia and industry.

The SPTM TC promotes activities within the technical areas of digital signal processing and statistical signal processing. As such, its interests span a broad range of topics, from classical fields (signal sampling and reconstruction, digital filtering, multirate signal processing, adaptive signal processing, statistical signal analysis, estimation, and detection) to more recent and currently emerging areas (compressed sensing, signal processing over graphs and networks, and optimization). Due to its breadth and fundamental nature, this TC is often the first home for new directions within the scope of the SPS.

As its name suggests, the SPTM TC is a cross-cutting group that focuses on theory and methods for signal processing applications. The TC is interested in the theoretical aspects of the methods, which are spread across and employed in numerous domains, such as sensor arrays, communication, audio and image processing, and biomedical applications. Theoretical studies that are fundamental to the advancement of these important areas include time–fre-

quency analysis, multirate processing, optimization methods, adaptive filtering, filter design, detection and estimation theory and methods, sampling theory, transforms, algorithm performance analysis, signal and information processing over graphs, learning theory, quantization, tensor theory, signals and systems modeling, distributed processing, sparsity-aware processing, quantum signal processing, and distributed processing over networks.

The SPTM TC sponsors the IEEE Workshop on Statistical Signal Processing (SSP) (Figure 1), which is the most prestigious event for presenting and discussing new research results on statistical methods applied to signal processing (Table 1). The 2018 workshop was held in Freiburg im Breisgau, Germany. It was chaired by Peter Schreier, with Javier Vía and Arie Yeredor as technical program cochair, and attracted 250 attendees. The 2020 workshop will be held in Rio de Janeiro, Brazil, and chaired by Rodrigo de Lamare, with Vítor Nascimento and José Bermudez as technical program cochair.

The largest meeting point for the SPTM community is ICASSP. The

SPTM track had 375 paper submissions for ICASSP 2016 in Shanghai, China; 410 submissions for ICASSP 2017 in New Orleans, Louisiana; 331 submissions for ICASSP 2018 in Calgary, Alberta, Canada; and 333 submissions for ICASSP 2019 in Brighton, United Kingdom. The upcoming ICASSP 2020 (Barcelona, Spain) has 411 submissions in the SPTM track, breaking the four-year record established in 2017 and showing the continued importance of the theoretical advances in the SPTM field for the development of new technologies.

The scope of the SPTM TC aligns with that of the European Association for Signal Processing (EURASIP) Theoretical and Methodological Trends in Signal Processing Technical Area Committee, so it is not surprising that there is overlap in topics, activities, and membership. A strong example is the joint cosponsorship of seasonal schools, such as the 2019 IEEE-SPS/EURASIP Summer School on Network- and Data-Driven Learning.

The SPTM TC participates in other technical activities within the IEEE SPS,

*(continued on page 102)*

**Table 1. The recent SSP Workshop editions.**

Year	Location	Chairs	Papers	Attendees
2014	Gold Coast, Australia	Abd-Krim Seghouane and Rob Evans	143	141
2016	Palma de Mallorca, Spain	Antonio Artés-Rodríguez and Joaquín Míguez Arenas	154	193
2018	Freiburg im Breisgau, Germany	Peter Schreier	173	250

## Signal Processing for Neurorehabilitation and Assistive Technologies

### GUEST EDITORS

Dario Farina, Imperial College London, London, UK, ([d.farina@imperial.ac.uk](mailto:d.farina@imperial.ac.uk))  
 Arash Mohammadi, Concordia University, Montreal, Canada, ([arash.mohammadi@concordia.ca](mailto:arash.mohammadi@concordia.ca))  
 Nitish V. Thakor, National University of Singapore, Singapore, ([eletnv@nus.edu.sg](mailto:eletnv@nus.edu.sg))  
 Konstantinos N. Plataniotis, University of Toronto, Toronto, Canada, ([kostas@ece.utoronto.ca](mailto:kostas@ece.utoronto.ca))  
 Tülay Adalı, University of Maryland Baltimore County, Baltimore, USA, ([adali@umbc.edu](mailto:adali@umbc.edu))

### SCOPE

During the twentieth century, the number of seniors over the age of sixty has increased significantly, which has raised the alarming concern of society aging and the balance between the available resources and the need for serving an aged society. A recent study from the United Nations (UN) has shown that the number of people over the age of sixty will increase from 962 million in 2017 to 2.1 billion by 2050, and to 3.1 billion by 2100. Similarly, the average life expectancy in the US is expected to reach 75.0 and 83.1 in 2040 for men and women, respectively. In parallel with the ageing of the world population, there has been an increase in age-related health issues, such as stroke, sensorimotor disorders, Parkinson's disease, and essential tremor, which significantly impact the health care systems. With a system that is under-resourced, patients are transferred from the hospitals to home while still suffering from major functional deficits. In this ageing crisis, a potential solution is to develop technologies and techniques that can provide (a) efficient, effective, widely-accessible, and affordable means of neurorehabilitation; and (b) intuitive and agile assistance to maximize the patients' independence during activities of daily living. Biological Signal Processing (BSP) is playing an imperative role in this domain for realizing advanced, intelligent, and dynamic rehabilitation and assistive solutions. This area of research includes processing, decomposing, and decoding of bioelectrical, biomechanical, and biochemical signals. The nonstationary and nonlinear nature of biological signals and the corresponding technical challenges in the area of man-machine interfacing, call for novel and innovative techniques beyond conventional approaches. The ultimate goal is to implement practical and effective augmentation techniques for the sensorimotor capabilities of the affected individuals and to allow for either (i) Instantaneous replacement of the lost motor functions (i.e., assistive solution), or; (ii) Gradual enhancement of the residual functions (i.e., rehabilitative solution). Through the use of advanced signal processing techniques, the motor intent of the individual can be decoded and decomposed into the underlying neuromechanical activities. This would ultimately enable the fusion of human neuromechanics with robotic, bionic, and neurorehabilitative technologies. Through advanced real-time processing of multichannel and multimodal biological signals (e.g., EEG, ECG, eye gaze, video, speech, MMG, FMG, and EMG), effective and alternative treatments and diagnosis are possible. Leveraging recent advances in the field of biological signal processing and machine learning, a revolution can be foreseen in the field of neurorehabilitation and neuroprosthetics. Motivated by the above-mentioned note, this special issue seeks to provide a platform for summarizing, educating, and sharing the state-of-the-art techniques and technologies related to signal processing applied to the domain of neurorehabilitation, neuroprosthetics, and assistive systems. Topics of interest include, but are not limited to:

- Biological signal processing for efficient and high-resolution man-machine interfacing
- High-fidelity brain-computer interfacing for diagnosis and therapy
- Signal Processing for neurorehabilitation robotic technologies
- Signal processing for wearable health technologies
- Signal and video processing for Tele-rehabilitation
- Machine intelligence and biosignal processing for neuroprosthetics
- Blind source separation for decomposing high-density biosignals
- IoT applications in neurorehabilitation, and movement assistance
- Multi-modal and distributed sensing for mobile and telehealth
- Signal processing for multilateral rehabilitation
- Biofeedback technologies for regeneration of the lost perceptual capabilities

### SUBMISSION GUIDELINES

Authors are invited to submit their contributions by following the detailed instructions given in the "Information for Authors" at <https://signalprocessingsociety.org/publications-resources/ieee-signal-processing-magazine>. Manuscripts should be submitted via Manuscript Central at <https://mc.manuscriptcentral.com/spmag-ieee>. Questions about the special issue should be directed to the Guest Editors.

### TENTATIVE SCHEDULE

White papers due:	April 1, 2020	Revision due:	November 1, 2020
Invitation notification:	May 1, 2020	Final decision:	January 1, 2021
Full length Manuscripts due:	July 1, 2020	Final package due:	February 5, 2021
First review to authors:	September 1, 2020	Publish manuscript:	May 1, 2021

# MATLAB SPEAKS DEEP LEARNING

With MATLAB®, you can build and deploy deep learning models for signal processing, reinforcement learning, automated driving, and other applications. Preprocess data, train models, generate code for GPUs, and deploy to production systems.

[mathworks.com/deeplearning](https://mathworks.com/deeplearning)

