AI Matters

Annotated Table of Contents

Links

SIGAI website: http://sigai.acm.org/
Newsletter: http://sigai.acm.org/aimatters/
Blog: http://sigai.acm.org/ai-matters/
Twitter: http://twitter.com/acm_sigai/
Edition DOI: 10.1145/3284751

Join SIGAI

Students $11, others $25
For details, see http://sigai.acm.org/
Benefits: regular, student

Also consider joining ACM.

Our mailing list is open to all.

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in this Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library and in any Digital Library related services
- to allow users to make a personal copy of the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will refer requests for republication directly to you.

Submit to AI Matters!

We're accepting articles and announcements now for the next issue. Details on the submission process are available at http://sigai.acm.org/aimatters.

AI Matters Editorial Board

Contact us: aimatters@sigai.acm.org

Contents Legend


Book Announcement


Ph.D. Dissertation Briefing


AI Education


Event Report


Hot Topics


Humor


AI Impact


AI News


Opinion


Paper Précis


Spotlight


Video or Image

Details at http://sigai.acm.org/aimatters

Welcome to AI Matters 4(3)

Amy McGovern, Editor (University of Oklahoma; aimatters@sigai.acm.org)
DOI: 10.1145/3284751.3284752

Issue overview

Welcome to the third issue of the fourth volume of the AI Matters Newsletter. In this issue, we have some great news: we are bringing back the AI student essay contest! Please see the first article for details on this fun contest which is due January 10, 2019. In addition to the prizes discussed in the article, the winners will have their articles features in AI Matters in the issue following their selection.

Our second big news item is an in-depth summary report of what ACM SIGAI has been working on for the last year. This report is in-depth and shows the wide extent that ACM SIGAI has been reaching out, from starting a new conference, helping with a variety of existing conferences, creating new awards, supporting members, adding new content to this newsletter, and increasing our social media presence.

In our regular columns, don't forget to check out upcoming AI related events, summarized well in Michael Rovatsos' events column. We got extra lucky this issue with two educational submissions! Michael Guerzhoy contributes an article on using deep learning and tensorflow in machine learning classes and Nathan Sturtevant discusses how to teach search algorithms to students in AI classes.

Our interview series is back! Marion Neumann has joined our contributing authors and will be working on interviews as well as other series. In this issue, Marion interviews Kristian Kersting, Professor in Computer Science and Deputy Director of the Centre for Cognitive Science at the Technical University of Darmstadt, Germany.

Larry Medsker's policy column looks into the recent report from the WEF on the future of jobs with increasing automation as well as vehicle automation. Finally, we close with two paper contributions. The first is a discussion of the Partnership on AI to Benefit People and Society. The second is a thoughtful paper on the Mechanism Design for Social Good agenda.

Submit to AI Matters!

Thanks for reading! Don't forget to send your ideas and future submissions to *AI Matters*! We're accepting articles and announcements now for the next issue. Details on the submission process are available at http://sigai.acm.org/aimatters.

Amy McGovern is chief Editor of AI Matters. She is a Professor of computer science at the University of Oklahoma and an adjunct professor of meteorology. She directs the Interaction, Discovery, Exploration and Adaptation (IDEA) lab. Her research focuses on machine learning and data mining with applications to high-impact weather.

2018 ACM SIGAI Student Essay Contest on Artificial Intelligence Technologies

Nicholas Mattei (IBM Research; n.mattei@ibm.com)

DOI: 10.1145/3284751.3284753

Abstract

Win one of several prizes including $500USD or a chat with a leading AI Researcher. For complete details, including submission requirements, formats, and judges please see: www.tinyurl.com/SIGAIEssay2018

2018 Topic

The ACM Special Interest Group on Artificial Intelligence (ACM SIGAI) supports the development and responsible application of Artificial Intelligence (AI) technologies. From intelligent assistants to self-driving cars, an increasing number of AI technologies now (or soon will) affect our lives. Examples include Google Duplex (Link) talking to humans, Drive.ai (Link) offering rides in US cities, chatbots advertising movies by impersonating people (Link), and AI systems making decisions about parole (Link) and foster care (Link). We interact with AI systems, whether we know it or not, every day.

Such interactions raise important questions. ACM SIGAI is in a unique position to shape the conversation around these and related issues and is thus interested in obtaining input from students worldwide to help shape the debate. We therefore invite all students to enter an essay in the 2018 ACM SIGAI Student Essay Contest, to be published in the ACM SIGAI newsletter *AI Matters*, addressing one or both of the following topic areas (or any other question in this space that you feel is important) while providing supporting evidence:

- What requirements, if any, should be imposed on AI systems and technology when interacting with humans who may or may not know that they are interacting with a machine? For example, should they be required to disclose their identities? If so, how? See, for example,"Turing?s Red Flag" in CACM (Link).
- What requirements, if any, should be imposed on AI systems and technology when making decisions that directly affect humans? For example, should they be required to make transparent decisions? If so, how? For example, the IEEE's discussion of Ethically Aligned Design (Link).

Each of the above topic areas raises further questions, including

- Who is responsible for the training and maintenance of AI systems? See, for example, Google?s (Link), Microsoft?s (Link), and IBM?s (Link) AI Principles.
- How do we educate ourselves and others about these issues and possible solutions? See, for example, new ways of teaching AI ethics (Link).
- How do we handle the fact that different cultures see these problems differently? See, for example, Joi Ito?s discussion in Wired (Link).
- Which steps can governments, industries, or organizations (including ACM SIGAI) take to address these issues? See, for example, the goals and outlines of the Partnership on AI (Link).

All sources must be cited. However, we are not interested in summaries of the opinions of others. Rather, we are interested in the informed opinions of the authors. Writing an essay on this topic requires some background knowledge. Possible starting points for acquiring such background knowledge are: the revised ACM Code of Ethics (Link), especially Section 3.7, and a discussion of why the revision was necessary (Link), IEEE?s Ethically Aligned Design (Link), and the One Hundred Year Study on AI and Life in 2030 (Link).

Format and Eligibility

The ACM SIGAI Student Essay Contest is open to all ACM SIGAI student members at the time of submission. (If you are a student but not an ACM SIGAI member, you can join ACM SIGAI before submission for just US $11 at https://goo.gl/6kifV9 by selecting Option 1, even if you are not an ACM member.) Essays can be authored by one or more ACM SIGAI student members but each ACM SIGAI student member can (co-)author one essay.

All authors must be SIGAI members at the time of submission. All submissions not meeting this requirement will not be reviewed.

Essays should be submitted as pdf documents of any style with at most 5,000 words via email to AI Matters Easychair.

The deadline for submissions is January 10th, 2019.

Prizes

All winning essays will be published in the ACM SIGAI newsletter *AI Matters*. ACM SIGAI provides five monetary awards of USD 500 each as well as 45-minute Skype sessions with one of the following AI researchers:

- Joanna Bryson, Reader (Assoc. Prof) in AI, University of Bath
- Murray Campbell, Senior Manager, IBM Research AI
- Eric Horvitz, Managing Director, Microsoft Research
- Peter Norvig, Director of Research, Google
- Iyad Rahwan, Associate Professor, MIT Media Lab and Head of Scalable Corp.
- Francesca Rossi, AI and Ethics Global Lead, IBM Research AI
- Toby Walsh, Scientia Professor of Artificial Intelligence, UNSW Sydney, Data61 and TU Berlin

Judges and Judging Criteria

Winning entries from last year's essay contest can be found in recent issues of the ACM SIGAI newsletter *AI Matters*, specifically Volume 3, Issue 3 and Volume 3, Issue 4.

Entries will be judged by the following panel of leading AI researchers and ACM SIGAI officers. Winning essays will be selected based on depth of insight, creativity, technical merit, and novelty of argument. All decisions by the judges are final.

- Rediet Abebe, Cornell University
- Emanuelle Burton, University of Illinois at Chicago
- Sanmay Das, Washington University in St. Louis
- John P. Dickerson, University of Maryland
- Virginia Dignum, Delft University of Technology
- Tina Eliassi-Rad, Northeastern University
- Judy Goldsmith, University of Kentucky
- Amy Greenwald, Brown University
- H. V. Jagadish, University of Michigan
- Sven Koenig, University of Southern California
- Benjamin Kuipers, University of Michigan
- Nicholas Mattei, IBM Research
- Alexandra Olteanu, Microsoft Research
- Rosemary Paradis, Leidos
- Kush Varshney, IBM Research
- Roman Yampolskiy, University of Louisville
- Yair Zick, National University of Singapore

Main organization by Nicholas Mattei (IBM Research), AI and Society Officer with involvement from Sven Koenig (University of Southern California), ACM SIGAI Chair; Sanmay Das (Washington University in St. Louis), ACM SIGAI Vice Chair; Rosemary Paradis (Leidos), ACM SIGAI Secretary/Treasurer; Benjamin Kuipers (University of Michigan), ACM SIGAI Ethics Officer; and Amy McGovern (University of Oklahoma), ACM SIGAI AI Matters Editor-in Chief.

In case of questions, please first check the ACM SIGAI blog for announcements and clarifications Link. You can also contact the ACM SIGAI Student Essay Contest Organizers at *sigai@member.acm.org*.

Nicholas Mattei is a Research Staff Member in the IBM Research AI group at the IBM TJ Watson Research Laboratory. His research focuses on the theory and practice of AI, developing systems and algorithms to support decision making.

ACM SIGAI Activity Report

Sven Koenig (elected; ACM SIGAI Chair)
Sanmay Das (elected; ACM SIGAI Vice-Chair)
Rosemary Paradis (elected; ACM SIGAI Secretary/Treasurer)
John Dickerson (appointed; ACM SIGAI Labor Market Officer)
Yolanda Gil (appointed; ACM SIGAI Past Chair)
Katherine Guo (appointed; ACM SIGAI Membership and Outreach Officer)
Benjamin Kuipers (appointed; ACM SIGAI Ethics Officer)
Hang Ma (appointed; ACM SIGAI Information Officer)
Nicholas Mattei (appointed; ACM SIGAI AI and Society Officer)
Amy McGovern (appointed; ACM SIGAI Newsletter Editor-in-Chief)
Larry Medsker (appointed; ACM SIGAI Public Policy Officer)
Todd Neller (appointed; ACM SIGAI Education Activities Officer)
Plamen Petrov (appointed; ACM SIGAI Industry Liaison Officer)
Michael Rovatsos (appointed; ACM SIGAI Conference Coordination Officer)
David G. Stork (appointed; ACM SIGAI Award Officer)
DOI: 10.1145/3284751.3284754

Abstract

We are happy to present the annual activity report of ACM SIGAI, covering the period from July 2017 to June 2018.

The scope of ACM SIGAI consists of the study of intelligence and its realization in computer systems (see also our website at `sigai.acm.org`). This includes areas such as

> autonomous agents, cognitive modeling, computer vision, constraint programming, human language technologies, intelligent user interfaces, knowledge discovery, knowledge representation and reasoning, machine learning, planning and search, problem solving, and robotics.

Our members come from academia, industry, and government agencies worldwide. We are thrilled to be able to report that our membership numbers increased by about 10 percent over the past year!

The terms of three of our officers (one of the two education activities officers, one of the two newsletter co-editors in chief, and the information officer) came to an end. We thank them for their valuable service and have now started to provide certificates of appreciation for outgoing officers. We appointed a new information officer and are currently looking for a new newsletter co-editor in chief and additional column editors. We also created and filled two new officer positions to be able to serve our members even better, namely an AI and society officer and a labor market officer.

Conferences

We helped to found the AAAI/ACM AI, Ethics, and Society (AIES) conference, a high-profile, multi-disciplinary meeting that addresses the impact of AI on society (including aspects such as value alignment, data handling and bias, regulations, and workforce impact) in a scientific context. We have a 50 percent financial stake in AIES, which we expect to become the prime international conference in the field. The inaugural AIES, held on February 13, 2018 in New Orleans directly before the AAAI conference, was a great success, attracting 162 submissions and selling out (see a recent summary in our newsletter AI Matters). The AI and society officer organized the AIES doctoral consortium, which received over 60 submissions but could only accept 20 students. AIES 2019 is now in the planning

phase.

We sponsored the following conferences in addition to AIES 2018:

- The IEEE/WIC/ACM International Conference on Web Intelligence (WI 2017), Leipzig, Germany, August 23-26, 2017
- The 32nd International Conference on Automated Software Engineering, Urbana-Champaign, USA, October 30-November 3, 2017
- The 9th International Conference on Knowledge Capture (K-CAP 2017), Austin, TX, USA, December 4-6, 2017
- The CRA Summit on Technology and Jobs, Washington D.C., USA, December 12, 2017
- The 23rd International Conference on Intelligent User Interfaces (IUI 2018), Tokyo, Japan, March 7-11, 2018
- The 13th Annual ACM/IEEE International Conference on Human Robot Interaction (HRI 2018), Chicago, USA, March 5-8, 2018

We approved the following in-cooperation and sponsorship requests from events covering a wide thematic and geographical range across the international AI community:

- The 20th International Conference on Enterprise Information Systems (ICEIS 2018), Funchal, Portugal, March 21-24, 2018
- The 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10-15, 2018
- The 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE 2018), Montpellier, France, September 3-7, 2018
- The 18th International Conference on Intelligent Virtual Agents (IVA 2018), Sydney, November 5-8, 2018
- The 2nd ACM Computer Science in Cars Symposium (CSCS 2018), Munich, Germany, September 13-14, 2018
- The 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2018), Seville, Spain, September 18-20, 2018
- International Conference on the Foundations of Digital Games (FDG 2018), Malm, Sweden, August 7-10, 2018
- The 10th International Joint Conference on Computational Intelligence, Seville, Spain, September 18-20, 2018
- The 15th International Conference on Informatics in Control, Automation, and Robotics (ICINCO 2018), Porto, Portugal, July 29-31, 2018
- The First IEEE International Conference on Artificial Intelligence and Virtual Reality (IEEE-AIVR 2018), Taichung, Taiwan, December 10-12, 2018

ACM SIGAI membership benefits include reduced registration fees for many of the in-cooperation and sponsored conferences and access to many proceedings in the ACM Digital Library.

We also have an agreement with the Association for the Advancement of AI (AAAI) to jointly organize the annual joint job fair at the AAAI conference, where attendees can find out about job and internship opportunities from representatives from industry, universities, and other organizations. The job fair in 2018 was run by two of our officers and attracted over 30 organizations offering jobs and hundreds of job seekers (see the recent summary in AI Matters). Similarly, we have an agreement with AAAI to jointly sponsor the annual joint doctoral consortium at the AAAI conference, which provides an opportunity for Ph.D. students to discuss their research interests and career objectives with the other participants and a group of established AI researchers that act as their mentors.

Awards

We started a new award, the ACM SIGAI Industry Award for Excellence in AI, which will be given annually to an individual or team in industry who created a fielded AI application in recent years that demonstrates the power of AI techniques via a combination of the following features: novelty of application area, novelty, and technical excellence of the approach, importance of AI techniques for the approach, and actual and predicted societal impact of the application. The award will be accompanied by a prize of US$5,000 and be presented at

the International Joint Conference on AI (IJCAI), starting in 2019 - thanks to a collaboration with the IJCAI Board of Trustees.

The ACM SIGAI Autonomous Agents Research Award was presented at the International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2018 - thanks to a collaboration with the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) - to Craig Boutilier, Principal Research Scientist at Google, for his seminal contributions to research on decision-making under uncertainty, game theory, and computational social choice.

Member Support

We support our members in different ways. For example, we nominate them for awards or support their nominations. We also nominate publications of recent, significant, and exciting AI research results that are of general interest to the computer science research community to the Research Highlights track of the Communications of the ACM (CACM) to make important AI research results visible to many computer scientists. We concentrate our financial support on our student members and provided funding to conferences so that they can award scholarships to their student attendees. We also supported the review process for the applications of students to attend the Heidelberg laureate forum to meet the recipients of the most prestigious awards in mathematics and computer science, including the ACM A.M. Turing Award and the ACM Prize in Computing.

To understand the interests of our members better, we conducted a membership survey and learned, among other things, that the majority of our members have been SIGAI members for two or more years. They are interested in AI resources (such as digital content access and webinars) as well as networking, and consider it important that we cover (among other things) hot AI topics, AI ethics, and AI applications with an industry focus. We also reached out to other groups to learn about their outreach activities, including to regional groups (such as the Northeast Ohio ACM chapter) and larger ACM Special Interest Groups (such as SIGCHI).

Member Communication

We communicate with our members via email announcements, the SIGAI newsletter AI Matters, the AI Matters blog, and webinars.

AI Matters

We continued to expand the scope of our newsletter AI Matters this year, introduced the EasyChair system for the submission of manuscripts and added additional column editors to the editorial team. We publish four issues of AI Matters per year that are openly available on the ACM SIGAI website at `sigai.acm.org/aimatters/` and feature articles of general interest to our members. Recurring columns have included:

- AI Interviews (with interesting people from academia, industry, and government),
- AI Amusements (including AI humor, puzzles, and games),
- AI Education (led by the education activities officer),
- AI Policy Issues (led by the public policy officer),
- AI Buzzwords (which explains new AI concepts or terms),
- AI Events (which includes conference announcements and reports, led by the conference coordination officer),
- AI Dissertation Abstracts and
- News from AI Groups and Organizations.

The Symposium on Educational Advances in AI (EAAI) will feature an undergraduate research track for the Birds of a Feather faculty-mentored undergraduate research challenge based on one of the AI Education columns. AI Matters also published the 8 winning entries from the ACM SIGAI Student Essay Contest on the Responsible Use of AI Technologies. One of these essays had more than 2,000 accesses (including accesses to the pdf file of the complete issue that contained the article).

AI Matters Blog

The AI Matters Blog is openly available on the ACM SIGAI website at `sigai.acm.org/aimatters/blog/` and serves as a forum for important announcements and news. For example, we post new

information every two weeks in the blog to survey and report on current AI policy issues and raise awareness about the activities of other organizations that share interests with ACM SIGAI. We are also open to posts by regional groups and have invited the Northeast Ohio ACM chapter to contribute.

Webinars

We extended our webinars on AI topics as part of our commitment to lifelong learning. In particular, we offered monthly webinars from November 2017 to June 2018 that focused on the application of AI technology to real-world problems and were presented by speakers typically involved with both academic research and industrial implementations:

- November 10, 2017
  Dan Moldovan
  Founder of Lymba Corporation, Professor in the Computer Science Department at the University of Texas at Dallas and Co-Director of its Human Language Technology Research Institute
  Topic: "On the Evolution of NLP, QA, and IE, and Current Research and Commercial Trends"
- December 15, 2017
  Peter Elkin
  Professor and Chair of the University at Buffalo Department of Biomedical Informatics
  Topic: "HTP-NLP: A New NLP System for High Throughput Phenotyping"
- January 12, 2018
  Lionel Jouffe
  Co-founder and CEO of France-based Bayesia S.A.S.
  Topic: "Data Mining, Knowledge Modeling, and Causal Analysis with Bayesian Networks"
- February 23, 2018
  Kristian Hammond
  Professor of Computer Science at Northwestern University and co-founder of Narrative Science
  Topic: "Communicating with the New Machine: Human Insight at Machine Scale"
- March 15, 2018
  Tomek Strzalkowski
  Director of the Institute for Informatics, Logics, and Security Studies and Professor at SUNY Albany
  Topic: "Advances in Socio-Behavioral Computing"
- May 7, 2018
  Jussi Karlgren
  KTH Royal Institute of Technology and Helsinki University and founding partner of Gavagi, a text analysis company
  Topic: "Explicitly Encoded High-Dimensional Semantic Spaces"
- June 4, 2018
  Maja Matarić
  Professor and Chan Soon-Shiong Chair in the Computer Science Department, Neuroscience Program, and the Department of Pediatrics at the University of Southern California
  Topic: "Socially Assistive Robotics"

The webinars were streamed live but can still be watched on demand at `learning.acm.org/webinar/`. They turned out to be very popular, typically reaching a thousand or more viewers.

Public Policy and AI Ethics

Within ACM, we work with the ACM US Public Policy Council (USACM) through the membership of our public policy officer in USACM and the participation of our members in US public policy issues related to computing and information technology. For example, our public policy officer works with the USACM leadership and the Electronic Privacy Information Center to petition the Office of Science and Technology Policy of the White House to construct and publicize a formal process by which the public might have input into the work of the recently-named Select Committee on Artificial Intelligence. He also studies how organizations collect and analyze data and whether these practices are consistent with recommendations by the USACM working groups on algorithmic accountability, transparency, and bias. Finally, he works on recommendations for possible changes to the ACM policy regarding data privacy of ACM SIGAI and ACM members who use EasyChair to submit articles for publication, including to AI Matters.

Outside of ACM, we helped to found the AAAI/ACM AI, Ethics, and Society (AIES) conference (as detailed above). We also partici-

pate in the executive committee of the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems to ensure that every technologist is educated, trained, and empowered to prioritize ethical considerations in the design and development of autonomous and intelligent systems.

Planning for the Future

In our last progress report, we promised to increase our understanding of our current membership. To further this objective, we performed a membership survey this year. We promised to intensify our activities to reach industry professionals. To further this objective, we substantially extended our webinars that focus on the application of AI technology to real-world problems and started the ACM SIGAI Industry Award for Excellence in AI. We promised to reach out to more AI groups worldwide that could benefit from ACM support, such as providing financial support, making the proceedings widely accessible in the ACM Digital Library, and providing speakers via the ACM Distinguished Speakers program. To further this objective, our membership officer contacted a number of such groups. We promised to expand the number of co-sponsored and in-cooperation conferences and to continue our efforts to further the discussion on the responsible use of AI technologies. To further these objectives, we helped to found the AAAI/ACM AI, Ethics, and Society conference. In the next year, we intend to continue the reorganization of AI Matters to be able to provide more content while spreading the production effort among more editors. We also intend to intensify our activities to support students and regional chapters and become even more international in our news coverage. Furthermore, the election of the next SIGAI leadership team is coming up and we are looking for strong candidates who are interested in shaping the future of ACM SIGAI!

Events

Michael Rovatsos (University of Edinburgh; mrovatso@inf.ed.ac.uk)
DOI: 10.1145/3284751.3284755

This section features information about upcoming events relevant to the readers of AI Matters, including those supported by SIGAI. We would love to hear from you if you are are organizing an event and would be interested in cooperating with SIGAI, or if you have announcements relevant to SIGAI. For more information about conference support visit sigai.acm.org/activities/requesting_sponsorship.html.

International Summit of Artificial Intelligence

Sophia Antipolis, France, November 7-9, 2018
sophia-summit.com/sophia2018/en

Europe's first science and technology park and digital cradle, Sophia Antipolis, will host, from November 7th to 9th, 2018, the top Artificial Intelligence specialists for a summit dedicated to researchers, businesses, economic partners, students and the general public. This summit will be the culmination of other events that, throughout 2018, have promoted and challenged Artificial Intelligence and the numerous applications of AI. The summit is jointly organized by the Sophia Antipolis Urban Community and University Côte d'Azur.

18th ACM International Conference on Intelligent Virtual Agents (IVA 2018)

Sydney, Australia, November 5-8, 2018
iva2018.westernsydney.edu.au

IVA 2018 is the 18th meeting of an interdisciplinary annual conference and the main leading scientific forum for presenting research on modeling, developing and evaluating intelligent virtual agents (IVAs) with a focus on communicative abilities and social behavior. IVAs are interactive digital characters that exhibit human-like qualities and can communicate with humans and each other using natural human modalities like facial expressions, speech and gesture. They are capable of real-time perception, cognition, emotion and action that allow them to participate in dynamic social environments. In addition to presentations on theoretical issues, the conference encourages the showcasing of working applications. Students, academics and industry professionals with an interest in learning about and presenting the most cutting-edge research conducted today in the multi-disciplinary field of intelligent virtual agents. Advances in IVA research are enabling increasingly autonomous agents that are now being utilized across a growing number of fields, including counseling, entertainment, medicine, the military, psychology and teaching. Researchers from the fields of human-human and human-robot interaction are also encouraged to share work with a relevance to IVAs.

1st IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR 2018)

Taichung, Taiwan, December 10-12, 2018
aivr.asia.edu.tw/2018/

The First IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR 2018) is a unique event, addressing researchers and industries from all areas of AI as well as Virtual, Augmented, and Mixed Reality. It provides an international forum for the exchange between those fields, to present advances in the state of the art, identify emerging research topics, and together define the future of these exciting research domains. We invite researchers from Virtual, as well as Augmented Reality (AR) and Mixed Reality (MR) to participate and submit their work to the program. Likewise, any work on AI that has a relation to any of these fields or potential for the usage in any of them is welcome.

24th International Conference on Intelligent User Interfaces (IUI 2019)

Los Angeles, CA, USA, March 17-20, 2019
iui.acm.org/2019

ACM IUI 2019 is the 24th annual meeting of the intelligent interfaces community and serves as a premier international forum for

reporting outstanding research and development on intelligent user interfaces. ACM IUI is where the Human-Computer Interaction community meets the Artificial Intelligence community. We are also very interested in contributions from related fields, such as psychology, behavioral science, cognitive science, computer graphics, design, the arts, etc.

32nd International Conference on Industrial, Engineering & Other Applications (IEA/AIE 2019)

Graz, Austria, July 9-11, 2019
ieaaie2019.ist.tugraz.at
The 32nd International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems will continue the tradition of emphasizing on applications of applied intelligent systems to solve real-life problems in all areas. These areas include: engineering, science, industry, automation & robotics, business & finance, medicine and biomedicine, bioinformatics, cyberspace, and human-machine interactions. IEA/AIE 2019 will have a special focus on automated driving and autonomous systems such that contributions dealing with such systems or their verification and validation are especially welcome.

Michael Rovatsos is the Conference Coordination Officer for ACM SIGAI, and a faculty member of the School of Informatics at the University of Edinburgh, UK. His research in in multiagent systems and human-friendly AI. Contact him at mrovatso@inf.ed.ac.uk.

AI Education Matters: Teaching with Deep Learning Frameworks in Introductory Machine Learning Courses

Michael Guerzhoy (Princeton University, University of Toronto, and the Li Ka Shing Knowledge Institute, St. Michael's Hospital; guerzhoy@princeton.edu)

DOI: 10.1145/3284751.3284756

Introduction

In this article, we demonstrate an assignment[1] in which students use TensorFlow to build a face recognition system. Students build shallow and deep neural networks for face recognition in TensorFlow and use transfer learning to obtain near-perfect performance on a simple face recognition task. Visualizing neural networks in order to explain how they work is central to the assignment.

The emergence of frameworks such as TensorFlow (Abadi et al., 2016) and PyTorch[2] and the explosive growth of ready-made open source machine learning systems on the Web changed the practice of machine learning. The pedagogy of machine learning faces challenges that are analogous to the challenges introductory computer science pedagogy started facing with the advent of high-level programming languages and GUI libraries (Astrachan, Bruce, Koffman, Kölling, & Reges, 2005): we would like for students to use the tools practitioners use to build real-world systems, but at the same time we want students to have a thorough understanding of how systems can be built "from scratch."

We resolve this dilemma by advocating for an assignment style that approximates building a new machine learning system prototype as closely as possible. We avoid giving out incomplete starter code and asking students to "fill in the blanks." Instead, we give students working code that is relevant to the task they are working on and that is similar to code that they might find on the Web. Students use as much or as little of the code as they want to complete the assignment.

[1] http://modelai.gettysburg.edu/2018/nnfaces/

[2] Excellent introductory tutorials are available at http://www.tensorflow.org and http://www.pytorch.org

Figure 1: Demonstrating overfitting by visualizing the weights of a one-hidden-layer neural network

We ask students to visualize the models that they are working with, building on recent research work on visualizing and understanding neural networks (Zeiler & Fergus, 2014) (Springenberg, Dosovitskiy, Brox, & Riedmiller, 2014). Visualizing ML models is a nontrivial technical task which allows students to practice developing machine learning code, it is helpful when debugging ML models, and, finally, it allows students to gain a better intuition for why the models work and when they can fail. Pedagogically, this task is both nontrivial and feasible for students to accomplish.

Model Assignment: Neural Networks for Face Recognition with TensorFlow

In the assignment, students train both shallow and deep networks to classify faces of six famous actors. Students adapt handout TensorFlow code in order to build and train the networks.

Students train a one-hidden-layer neural network for face classification. Students then improve the performance of their system by train-

ing a convolutional neural network to perform face classification; as a first step, students use the activations of a pre-trained AlexNet network (Krizhevsky, Sutskever, & Hinton, 2012) as the features for their face classifier.

Students visualize the shallow and the deep networks that they train.

Introducing TensorFlow

TensorFlow is introduced by giving students working code that trains a one-hidden-layer neural network on the MNIST dataset. We also provide a beginner-friendly implementation of AlexNet in TensorFlow, along with weights pretrained on ImageNet[3]. The handout is designed to be very easily readable and only uses the most elementary TensorFlow features. This makes it possible for students to modify the code to suit their needs without having to immediately understand the highly-engineered neural network implementations that are included with TensorFlow. Students can (and do) reuse the generic handout code later on in their career.

Visualizing Feedforward Networks

Visualizing the weights of shallow neural networks allows students to understand one-hidden-layer neural networks in terms of template matching, and allows students to explore overfitting. Students can demonstrate that, with enough hidden units, a one-hidden-layer neural network will "memorize" faces (see Fig. 1).

Students start exploring the visualization of deep neural networks by computing the gradient of a network's output with respect to its input using TensorFlow and visualizing it. For the bonus component of the assignment, students are asked to explore a visualization strategy of their choice for understanding the convolutional network they train. Guided Backpropagation (Springenberg et al., 2014)) is a common technique for students to attempt to implement.

[3] http://github.com/guerzh/tf_weights

Conclusion

Deep Learning frameworks should be introduced early in machine learning courses. Handout code that trains simple networks can be given out. Performing transfer learning and visualizing the neural networks that students train are two interesting tasks that can be assigned to students.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... Zheng, X. (2016). TensorFlow: A System for Large-Scale Machine Learning. In *USENIX Symposium on Operating Systems Design and Implementation* (Vol. 16, pp. 265–283).

Astrachan, O., Bruce, K., Koffman, E., Kölling, M., & Reges, S. (2005). Resolved: objects early has failed. In *ACM SIGCSE Bulletin* (Vol. 37, pp. 451–452).

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (pp. 1097–1105).

Springenberg, J. T., Dosovitskiy, A., Brox, T., & Riedmiller, M. (2014). Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.

Zeiler, M. D., & Fergus, R. (2014). Visualizing and Understanding Convolutional Networks. In *European Conference on Computer Vision* (pp. 818–833).

Michael Guerzhoy is a Lecturer in the Center for Statistics and Machine Learning at Princeton University, an Assistant Professor (Status Only) in the Dept. of Statistical Sciences at the University of Toronto, and a Scientist at the Li Ka Shing Knowledge Institute, St. Michael's Hospital. His professional interests are in computer science and data science education and in applications of machine learning to healthcare.

AI Education Matters: Teaching Search Algorithms

Nathan R. Sturtevant (University of Alberta; nathanst@ualberta.ca)
DOI: 10.1145/3284751.3284757

Introduction

In this column, we share advice and resources for teaching and learning about heuristic search algorithms. These are algorithms like A* (Hart, Nilsson, & Raphael, 1968), IDA* (Korf, 1985) and others that use admissible heuristics to find optimal and suboptimal solutions to shortest path problems.

While A* and IDA* are well-known and well-covered in many courses, they can still be confusing to students studying them for the first time. But, there is a significant body of work that goes beyond these basic algorithms that can be used as additional projects and to encourage advanced students. These includes topics such as:

- Improving search with better heuristics (Culberson & Schaeffer, 1996; Sturtevant, Felner, Barer, Schaeffer, & Burch, 2009) and constraints (Goldberg, Kaplan, & Werneck, 2006; Rabin & Sturtevant, 2016)
- Exploiting path symmetries on grids (Harabor & Grastien, 2011; Sturtevant & Rabin, 2016)
- Any-angle path planning (Nash & Koenig, 2013)
- 3D Path Planning (Brewer & Sturtevant, 2018)
- Suboptimal path planning (Hatem & Ruml, 2014)
- Bidirectional path planning (Sturtevant & Felner, 2018)
- Multi-agent path planning (Felner et al., 2017)

Textbooks

There are two common textbooks that provide an introduction to heuristic search.

Russell and Norvig's *Artificial Intelligence: A Modern Approach* (Russell & Norvig, 2009) covers the basics of search with heuristics and gives several examples of different heuristics for different domains. This book is relatively high-level with mostly introductory material.

On the opposite spectrum is Stefan Edelkamp's Textbook, *Heuristic Search: Theory and Applications* (Edelkamp & Schrödl, 2012). This excellent book goes into great detail about most aspects of heuristic search, however it may be too advanced for most undergraduates.

Video Introductions

There are many videos available on YouTube which demonstrate the A* algorithm and go through examples of how it works. Our Moving AI web pages[1] feature videos on different aspects of search algorithms that are both available on YouTube or can be download for offline use.

Interactive Web Demos

More useful than static videos are interactive web pages that allow users to interact with the search process by changing problem instances, search algorithms, and other parameters.

Amit Patel has developed a broad range of tutorials for subjects of interest to game developers. Among these, he has many different pages that explore and explain search algorithms, some statically[2] and some interactively.[3]

Two other notable pages have been developed by Xueqiao (Joe) Xu[4] and Dave Churchill[5].

[1] https://www.movingai.com/
[2] http://theory.stanford.edu/~amitp/GameProgramming/
[3] https://www.redblobgames.com/pathfinding/a-star/introduction.html
[4] https://qiao.github.io/PathFinding.js/visual/
[5] http://www.cs.mun.ca/~dchurchill/search/

Each of these pages allow you to interact with a variety of search algorithms and search parameters, visualizing how the algorithms run.

Finally, we have recently released a broad set of demos[6] that we have been using for our cross-listed undergraduate/graduate course on heuristic search. In addition to the regular A*-like algorithms, this page also has demos of IDA*, bidirectional search, heuristics, constraints, and other search techniques. These demos are still under active development, particularly when we teach our heuristic search course. In addition to providing these demos, we give them in class with challenge questions, such as 'Find a problem where A* updates the $g-$cost of a state on the open list'. This gives student more purpose when exploring them.

Model AI Assignments

Model AI Assignments are free, peer-reviewed assignment materials made available in order to advance AI education. There are several model assignments that explore aspects of heuristic search.

Sven Koenig's has two Model AI assignments available. "Any-Angle Path Planning"[7] is for undergraduate or graduate artificial intelligence classes and covers algorithms that use a grid representation of the obstacles in the world when agents are able to move through free space.

"Fast Trajectory Replanning"[8] is also for undergraduate and graduate students and covers algorithms for re-planning when the world representation changes or the agents discover obstacles they were not previously aware of.

"The Pac-Man Projects" by John DeNero and Dan Klein also has a heuristic search project, "Project #1: Search Project"[9] in which "students implement depth-first, breadth-first, uniform cost, and A* search algorithms. These algorithms are used to solve navigation and traveling salesman problems in the Pac-Man world. "

[6] https://movingai.com/SAS/

[7] http://idm-lab.org/project-m/project2.html

[8] http://idm-lab.org/project-m/project1.html

[9] http://modelai.gettysburg.edu/2010/pacman/projects/search/search.html

Conferences for Heuristic Search

For those looking for recent research in the area of heuristic search, there are several venues that regularly publish work in heuristic search. These include the Symposium on Combinatorial Search (SoCS)[10], which specializes particularly in search. The International Conference on Planning and Scheduling (ICAPS)[11] has a broader scope, but regularly contains significant research in heuristic search. Finally, major conferences like the International Joint Conference on Artificial Intelligence (IJCAI)[12] and the AAAI Conference on Artificial Intelligence[13] also publish research in this subfield of Artificial Intelligence.

References

Brewer, D., & Sturtevant, N. R. (2018). Benchmarks for pathfinding in 3d voxel space. *Symposium on Combinatorial Search (SoCS)*, 143 – 147.

Culberson, J. C., & Schaeffer, J. (1996). Searching with pattern databases. *Advances in Artificial Intelligence*, 402-416.

Edelkamp, S., & Schrödl, S. (2012). *Heuristic search - theory and applications*. Academic Press. Retrieved from http://www.elsevierdirect.com/product.jsp?isbn=9780123725127

Felner, A., Stern, R., Shimony, S. E., Boyarski, E., Goldenberg, M., Sharon, G., ... Surynek, P. (2017). Search-based optimal solvers for the multi-agent pathfinding problem: Summary and challenges..

Goldberg, A. V., Kaplan, H., & Werneck, R. F. (2006). Reach for A*: Efficient point-to-point shortest path algorithms. In *ALENEX* (pp. 129–143).

Harabor, D. D., & Grastien, A. (2011). Online graph pruning for pathfinding on grid maps. In *AAAI* (pp. 1114–1119).

[10] http://search-conference.org

[11] http://www.icaps-conference.org

[12] http://www.ijcai.org/

[13] http://www.aaai.org/Conferences/conferences.php

Hart, P., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, *4*, 100–107.

Hatem, M., & Ruml, W. (2014). Simpler bounded suboptimal search. In *AAAI* (pp. 856–862).

Korf, R. E. (1985). Iterative-Deepening-A*: An optimal admissible tree search. In *IJCAI* (p. 1034-1036).

Nash, A., & Koenig, S. (2013). Any-angle path planning. *AI Magazine*, *34*(4), 9.

Rabin, S., & Sturtevant, N. R. (2016). Combining bounding boxes and jps to prune grid pathfinding. In *AAAI.*

Russell, S., & Norvig, P. (2009). *Artificial intelligence: A modern approach* (3rd ed.). Upper Saddle River, NJ, USA: Prentice Hall.

Sturtevant, N. R., & Felner, A. (2018). A brief history and recent achievements in bidirectional search. In *Aaai conference on artificial intelligence.*

Sturtevant, N. R., Felner, A., Barer, M., Schaeffer, J., & Burch, N. (2009). Memory-based heuristics for explicit state spaces. *IJCAI*, 609–614.

Sturtevant, N. R., & Rabin, S. (2016). Canonical orderings on grids. *IJCAI*, 683–689.

Nathan R. Sturtevant is a Professor of Computer Science at the University of Alberta doing research in heuristic search, artificial intelligence, and games. He has implemented his research in commercial games and is a senior member of the Association for the Advancement of Artificial Intelligence.

AI Profiles: An Interview with Kristian Kersting

Marion Neumann (Washington University in St. Louis; m.neumann@wustl.edu)
DOI: 10.1145/3284751.3284758

Introduction

This column is the sixth in our series profiling senior AI researchers. This month we interview Kristian Kersting, Professor in Computer Science and Deputy Director of the Centre for Cognitive Science at the Technical University of Darmstadt, Germany.

Figure 1: Kristian Kersting

Biography

After receiving his Ph.D. from the University of Freiburg in 2006, he was with the MIT, Fraunhofer IAIS, the University of Bonn, and the TU Dortmund University. His main research interests are statistical relational artificial intelligence (AI), probabilistic deep programming, and machine learning. Kristian has published over 170 peer-reviewed technical papers and co-authored a book on statistical relational AI. He received the European Association for Artificial Intelligence (EurAI, formerly ECCAI) Dissertation Award 2006 for the best AI dissertation in Europe and two best-paper awards (ECML 2006, AIIDE 2015). He gave several tutorials at top AI conferences, co-chaired several international workshops, and cofounded the international workshop series on Statistical Relational AI (StarAI). He regularly serves on the PC (often at senior level) for several top conference and co-chaired the PC of ECML PKDD 2013 and UAI 2017. He is the Speciality Editor in Chief for Machine Learning and AI of Frontiers in Big Data, and is/was an action editor of TPAMI, JAIR, AIJ, DAMI, and MLJ.

Getting to Know Kristian Kersting

When and how did you become interested in AI?

As a student, I was attending an AI course of Bernhard Nebel at the University of Freiburg. This was the first time I dived deep into AI. However, my interest in AI was probably triggered earlier. Around the age of 16, I think, I was reading about AI in some popular science magazines. I did not get all the details, but I was fascinated.

What professional achievement are you most proud of?

We were collaborating with biologists on understanding better how plants react to (a)biotic stress using machine learning to analyze hyperspectral images. We got quite encouraging results. The first submission to a journal, however, got rejected. As you can imagine, I was disappointed. One of the biologists from our team looked at me and said ”Kristian, do not worry, your research helped us a lot.” This made me proud. But also the joint work with Martin Mladenov on compressing linear and quadratic programs using fractional automorpishms. This provides optimization flags for ML and AI compilers. Turning them on makes the compilers attempt to reduce the solver costs, making ML and AI automatically faster.

What would you have chosen as your career if you hadn’t gone into CS?

Physics, I guess, but back then I did not see any other option than Computer Science.

What do you wish you had known as a Ph.D. student or early researcher?

That "sleep is for post-docs," as Michael Littman once said.

Artificial Intelligence = Machine Learning. What's wrong with this equation?

Machine Learning (ML) and Artificial Intelligence (AI) are indeed similar, but not quite the same. AI is about problem solving, reasoning, and learning in general. To keep it simple, if you can write a very clever program that shows intelligent behavior, it can be AI. But unless the program is automatically learned from data, it is not ML. The easiest way to think of their relationship is to visualize them as concentric circles with AI first and ML sitting inside (with deep learning fitting inside both), since ML also requires to write programs, namely, of the learning process. The crucial point is that they share the idea of using computation as the language for intelligent behavior.

As you experienced AI research and education in the US and in Europe, what are the biggest differences between the two systems and what can we learn from each other?

If you present a new idea, US people will usually respond with "Sounds great, let's do it!", while the typical German reply is "This won't work because ...". Here, AI is no exception. It is much more critically received in Germany than in the US. However, this also provides research opportunities such as transparent, fair and explainable AI. Generally, over the past three decades, academia and industry have been converging philosophically and physically much more in the US than in Germany. This facilitate the transfer of AI knowledge via well-trained, constantly learning AI experts, who can then continuously create new ideas within the company/university and keep pace with the AI development. To foster AI research and education, the department structure and tenure-track system common in the US is beneficial. On the other hand, Germany is offering access to free higher education to all students, regardless of their origin. AI has no borders. We have to take it from the ivory towers and make it accessible for all.

What is the most interesting project you are currently involved with?

Deep learning has made striking advances in enabling computers to perform tasks like recognizing faces or objects, but it does not show the general, flexible intelligence that lets people solve problems without being specially trained to do so. Thus, it is time to boost its IQ. Currently, we are working on deep learning approaches based on sum-product networks and other arithmetic circuits that explicitly quantify uncertainty. Together with colleagues—also from the Centre of Cognitive Science—we combining the resulting probabilistic deep learning with probabilistic (logical) programming languages. If successful, this would be a big step forward in programming languages, machine learning and AI.

AI is grown up - it's time to make use of it for good. Which real-world problem would you like to see solved by AI in the future?

Due to climate change, population growth and food security concerns the world has to seek more innovative approaches to protecting and improving crop yield. AI should play a major role here. Next to feeding a hungry world, AI should aim to help eradicate disease and poverty.

We currently observe many promising and exciting advances in using AI in education, going beyond automating Piazza answering, how should we *make use of AI to teach AI*?

AI can be seen as an expanding and evolving network of ideas, scholars, papers, codes and showcases. Can machines read this data? We should establish the "AI Genome", a dataset, a knowledge base, an ongoing effort to learn and reason about AI problems, concepts, algorithms, and experiments. This would not only help to curate and personalize the learning experience but also to meet the challenges of reproducible AI research. It would make AI truly accessible for all.

What is your favorite AI-related movie or book and why?

"Ex Machina" because the Turing test is shaping its plot. It makes me think about current

real-life systems that give the impression that they pass the test. However, I think AI is hard than many people think.

Help us determine who should be in the AI Matters spotlight!

If you have suggestions for who we should profile next, please feel free to contact us via email at aimatters@sigai.acm.org.

AI Policy Matters

Larry Medsker (George Washington University; lrm@gwu.edu)
DOI: 10.1145/3284751.3284759

Abstract

AI Policy is a regular column in AI Matters featuring summaries and commentary based on postings that appear twice a month in the AI Matters blog (https://sigai.acm.org/aimatters/blog/). Selected posts are summarized in issues of *AI Matters*.

Introduction

The SIGAI Public Policy goals are to

- promote discussion of policies related to AI through posts in the AI Matters blog on the 1st and 15th of each month,
- help identify external groups with common interests in AI Public Policy,
- encourage SIGAI members to partner in policy initiatives with these organizations, and
- disseminate public policy ideas to the SIGAI membership through articles in the newsletter.

I welcome everyone to make blog comments so we can develop a rich knowledge base of information and ideas representing the SIGAI members.

WEF Report on the Future of Jobs

The World Economic Forum recently released a report (link) on the future of jobs . Their analyses refer to the Fourth Industrial Revolution (link) and their Centre for the Fourth Industrial Revolution. The report states that

> The Fourth Industrial Revolution is interacting with other socio-economic and demographic factors to create a perfect storm of business model change in all industries, resulting in major disruptions to labour markets. New categories of jobs will emerge, partly or wholly displacing others. The skill sets required in both old and new occupations will change in most industries and transform how and where people work. It may also affect female and male workers differently and transform the dynamics of the industry gender gap. The Future of Jobs Report aims to unpack and provide specific information on the relative magnitude of these trends by industry and geography, and on the expected time horizon for their impact to be felt on job functions, employment levels and skills.

The report concludes that by 2022 more jobs can be created than the number lost but that various stakeholders, including those making education policy, must make wise decisions.

Vehicle Automation: Safe Design, Scientific Advances, and Smart Policy

Updating previous columns on terminology and popular discourse about AI, a current example is the impact on policy of the way we talk about automation. “Unmanned Autonomous Vehicle (UAV)” is a term that justifiably creates fear in the general public, but talk about a UAV usually misses the necessary roles of humans and human decision making. Likewise, discussions about an “automated decision maker (ADM)” ignores the social and legal responsibilities of those who design, manufacture, implement, and operate “autonomous” systems. The AI community has an important role to influence correct and realistic use of concepts and issues in discussions of science and technology systems that increase automation. The concept “hybrid system” might be more helpful here for understanding the potential and limitations of combining technologies and humans into AI and Autonomous Systems (AI/AS) requiring less, but not zero, need for humans over time.

Safe Design

In addition to avoiding confusion and managing expectations, design approaches and analyses of the performance of existing systems with automation are crucial to developing

safe systems with which the public and policymakers can feel comfortable. In this regard, stakeholders should read information on design of systems with automation components, such as the IEEE report "Ethically Aligned Design" (link), which is subtitled "A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems". The report says about AI and Autonomous Systems (AI/AS):

> We need to make sure that these technologies are aligned to humans in terms of our moral values and ethical principles. AI/AS have to behave in a way that is beneficial to people beyond reaching functional goals and addressing technical problems. This will allow for an elevated level of trust between humans and our technology that is needed for a fruitful pervasive use of AI/AS in our daily lives.

See also Ben Shneiderman's excellent summary and comments (link) on the report as well as the YouTube video of his Turing Institute Lecture on "Algorithmic Accountability: Design for Safety". See also his proposal for a National Algorithms Safety Board (link).

Advances in AI/AS Science and Technology

Another perspective on the automation issue is the need to increase safety of systems through advances in science and technology. In a future blog, we will present the transcript of an interview with Dr. Harold Szu, about the need for a next generation of AI that moves closer to brain-style computing and that incorporates human behaviors into AI/AS systems. Dr. Szu was the founder and former president, and former governor, of the International Neural Network Society. He is acknowledged for outstanding contributions to ANN applications and scientific innovations.

Policy and Ethics

Over the summer 2018, increased activity in congress and state legislatures focused on understandings, accurate and not, of "unmanned autonomous vehicles" and what policies should be in place. The following examples are interesting for possible interventions, but also for the use of AI/AS terminology:

- House Energy & Commerce Committee's press release (link): the SELF DRIVE Act.
- CNBC Commentary (link) by Reps. Bob Latta (R-OH) and Jan Schakowsky (D-IL).
- Politico, 08/03/2018.: "Trial lawyers speak out on Senate self-driving car bill", by Brianna Gurciullo with help from Lauren Gardner.

  > AV NON-STARTER: After being mum for months, the American Association for Justice said publicly Thursday that it has been pressing for the Senate's self-driving car bill, S. 1885 (115) (definitions on p.42 link), to stipulate that companies can't force arbitration, our Tanya Snyder reports for Pros. The trial lawyers group is calling for a provision to make sure when a person, whether a passenger or pedestrian, is injured or killed by a driverless car, that person or their family is not forced into a secret arbitration proceeding, according to a statement. Senate Commerce Chairman John Thune (R-S.D.) has said that arbitration has been "a thorny spot" in bill negotiations.

ACM Code of Ethics

On Tuesday, July 17, ACM announced the updated Code of Ethics and Professional Conduct. Please note the message from ACM Headquarters and check the link below: "We would like your support in helping to reach as broad an audience of computing professionals as possible with this news." The updated Code is at link.

We encourage you to share the updated Code with your friends and colleagues at that time. If you use social media, please take part in the conversation around computing ethics using the hashtags #ACMCodeOfEthics and #IReadTheCode. And if you are not doing so already, please follow the @TheOfficialACM and @ACM_Ethics Twitter handles to share and engage with posts about the Code.

News From the ACM US Technology Policy Committee

The USACM has a new name Please note the change and remember that SIGAI will continue to have a close relationship with the ACM

US Technology Policy Committee (link). Here is a reminder of the purpose and goals: “The ACM US Technology Policy Committee is a leading independent and nonpartisan voice in addressing US public policy issues related to computing and information technology. The Committee regularly educates and informs Congress, the Administration, and the courts about significant developments in the computing field and how those developments affect public policy in the United States. The Committee provides guidance and expertise in varied areas, including algorithmic accountability, artificial intelligence, big data and analytics, privacy, security, accessibility, digital governance, intellectual property, voting systems, and tech law. As the internet is global, the ACM US Technology Policy Committee works with the other ACM policy entities on publications and projects related to cross-border issues, such as cybersecurity, encryption, cloud computing, the Internet of Things, and internet governance.”

The ACM US Technology Policy Committee’s New Leadership ACM named Prof. Jim Hendler as the new Chair of the ACM U.S. Technology Policy Committee (formerly USACM) under the new ACM Technology Policy Council. In addition to being a distinguished computer science professor at RPI, Jim has long been an active USACM member and has served as both a committee chair and as an at-large representative. This is a great choice to guide USACM into the future within ACM’s new technology policy structure. Please add individually to the SIGAI Public Policy congratulations to Jim. Our congratulations and appreciation go to outgoing Chair Stuart Shapiro for his outstanding leadership of USACM.

Larry Medsker is Research Professor of Physics and founding director of the Data Science graduate program at the Data Science graduate program at The George Washington University. Dr. Medsker is a former Dean of the Siena College School of Science, and Professor in Computer Science and in Physics, where he was a co-founder of the Siena Institute for Artificial Intelligence. His research and teaching continues at GW on the nature of humans and machines and the impacts of AI on society and policy[a]

[a] http://humac-web.org/. Professor Medskers research in AI includes work on artificial neural networks and hybrid intelligent systems. He is the Public Policy Officer for the ACM SIGAI.

The Partnership on AI

Jeffrey Heer (University of Washington; jheer@uw.edu)
DOI: 10.1145/3284751.3284760

The Partnership on AI to Benefit People and Society (or PAI) is a consortium of industrial, non-profit, and academic partners for promoting public understanding and beneficial use of Artificial Intelligence (AI) technologies. Primary goals of the PAI are "to study and formulate best practices on AI technologies, to advance the public's understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society."

PAI membership includes representatives from over 70 partner organizations (https://www.partnershiponai.org/partners/), spanning nine countries (primarily in North America, Europe, and Asia) and including technology companies of various sizes, academic institutions, and a number of non-profits (including ACM). The Board of Directors represents a number of large tech companies (Amazon, Apple, Facebook, IBM, Google/DeepMind, Microsoft) as well as non-profits (ACLU, MacArthur Foundation, OpenAI) and academics (Arizona State University, Harvard, UC Berkeley). The founding co-chairs are Eric Horvitz (Microsoft) and Mustafa Suleyman (DeepMind), with day-to-day operations of the PAI run by a dedicated staff in San Francisco, CA.

The PAI's research and outreach efforts center around six thematic pillars:

- Safety-Critical AI
- Fair, Transparent, and Accountable AI
- AI, Labor, and the Economy
- Collaborations Between People and AI Systems
- Social and Societal Influences of AI
- AI and Social Good

For more details regarding each thematic pillar, see https://www.partnershiponai.org/about/#our-work.

The PAI kick-off meeting was held in Berlin in October 2017, with wide-ranging discussions organized around the above pillars. A primary of outcome of the meeting and subsequent membership surveys was the formation of three initial working groups focused on "Safety-Critical AI", "Fair, Transparent, and Accountable AI", and "AI, Labor, and the Economy". Each of these groups are now conducting a series of meetings to develop best practices and resources for their respective topics.

I am a member of the Fair, Transparent, and Accountable AI (or FTA) working group, co-chaired by Ed Felten (Princeton) and Verity Harding (DeepMind). The charter of the FTA working group is to study issues and opportunities for promoting fair, accountable, transparent, and explainable AI, including the development of standards and best practices as a way to avoid and minimize the risk that AI systems will undermine fairness, equality, due process, civil liberties, or human rights. The FTA group has held two in-person meetings so far: the first in London in May 2018 and the second at Princeton in August 2018. The initial meeting resulted in the formation of four specific projects.

The "Papers" project aims to produce a set of three educational primers, one each for fairness, transparency, and accountability. The goal is to provide a starting point for people to learn about fairness in AI, how it is conceived, and what people are currently doing about it. The purpose is not just to come up with definitions, but provide structure for how to think about problems of fairness, transparency and accountability in the context of the work that Partner organizations are conducting.

The "Case Studies" project focuses of identifying multiple case studies that illustrate how FTA concerns can be addressed, as well as develop a framework for the analysis of such case studies. Meanwhile, the "Grand Challenges" project is working to create challenges or endorse existing contests developed around FTA issues in practice.

The “Diverse Voices” project seeks to understand different communities, how their voices can be involved in tech policy making, and the impact of AI deployment. The project intends to engage with these groups in a meaningful manner and ensure under-represented groups’ voices are heard. One goal is to provide guidance for outreach, as many companies are interested in deploying ethical AI but are unsure how to address the problem.

Each of these projects is a work-in-progress, with additional milestones and meetings planned over the coming year.

Jeffrey Heer is a Professor of Computer Science & Engineering at the University of Washington and Co-Founder of Trifacta Inc., a provider of interactive tools for scalable data transformation. His research spans human-computer interaction, visualization, data science, and interactive machine learning. He also serves as an ACM representative to the Partnership on AI.

Mechanism Design for Social Good

Rediet Abebe (Cornell University; red@cs.cornell.edu)
Kira Goldner (University of Washington; kgoldner@cs.washington.edu)
DOI: 10.1145/3284751.3284761

Introduction

Across various domains—such as health, education, and housing—improving societal welfare involves allocating resources, setting policies, targeting interventions, and regulating activities. These solutions have an immense impact on the day-to-day lives of individuals, whether in the form of access to quality healthcare, labor market outcomes, or how votes are accounted for in a democratic society. Problems that can have an outsized impact on individuals whose opportunities have historically been limited often pose conceptual and technical challenges, requiring insights from many disciplines. Conversely, the lack of interdisciplinary approach can leave these urgent needs unaddressed and can even exacerbate underlying socioeconomic inequalities.

To realize the opportunities in these domains, we need to correctly set objectives and reason about human behavior and actions. Doing so requires a deep grounding in the field of interest and collaboration with domain experts who understand the societal implications and feasibility of proposed solutions. These insights can play an instrumental role in proposing algorithmically-informed policies. In many cases, the input data for our algorithms may be generated by strategic and self-interested individuals who have a stake in the outcome of the algorithm. To get around this issue, we can deploy techniques from *mechanism design*, which uses game theory to align incentives or analyze the strategic behavior of individuals who interact with the algorithms.

The *Mechanism Design for Social Good* (MD4SG) research agenda is to address problems for which insights from algorithms, optimization, and mechanism design have the potential to improve access to opportunity. These include allocating affordable housing services, designing efficient health insurance markets, setting subsidies to alleviate economic inequality, and several other issues affecting many individuals' livelihoods. This research area falls at the interface of artificial intelligence, theoretical computer science, and the social sciences. Since the fall of 2016, the authors of this piece have been co-organizing the Mechanism Design for Social Good research group, workshop series, and colloquium series Abebe and Goldner [2016, 2018]. The group comprises a large network of researchers from various disciplines, including computer science, economics, sociology, operations research, and public policy. Members of the group partner with domain experts in non-government organizations, think tanks, companies, and other entities with a shared mission. The mission is to explore new frontiers, garner interest in directions in which algorithmic and mechanism design insights have been under-utilized but have the potential to inform innovative interventions, and highlight exemplary work.

In this piece, we discuss three exciting research avenues within MD4SG. For each of these, we showcase ongoing work, underline new directions, and discuss potential for implementing existing work in practice.

Access to Opportunity in the Developing World

New technologies and data sources are frequently leveraged to understand, evaluate, and address societal concerns across the world. In many developing nations, however, there is a lack of information regarding underlying matters—whether that be the prevalence of diseases or accurate measurements of economic welfare and poverty—due to the unavailability of high-quality, comprehensive, and reliable data UN [2014]. This limits the implementation of effective policies and interventions. An emerging solution, which has been successfully demonstrated by the Information Communication Technology for Development (ICT4D) research community, has been to take advantage of high phone and Internet

penetration rates across developing nations to design new technologies which enable collection and sharing of high-quality data. There has also been recent work from within the AI community to use new data sources to close this information gap Abebe et al. [2018], Jean et al. [2016]. Such AI-driven approaches surface new algorithmic, modeling, and mechanism design questions to improve the lives of many under-served individuals.

A prominent example is in agriculture, which accounts for a large portion of the economy in many developing nations. Here, viral disease attacks on crops is a leading cause of food insecurity and poverty. Traditional disease surveillance methods fail to provide adequate information to curtail the impact of diseases Mwebaze and Biehl [2016], Mwebaze and Owomugisha [2016], Quinn et al. [2011]. The Cassava Adhoc Surveillance Project from Makerere University implements crowd-sourcing surveillance using pictures taken by mobile phones in order to address this gap Mutembesa et al. [2018]. The tool is set up as a game between farmers and other collaborators, and aims to collect truthful, high-value data (e.g., data from hard-to-reach locations). This approach underlines interesting challenges, such as how to optimally incentivize individuals to collect high-quality information and how to augment this information with existing methods. Similar issues arise in other domains—e.g., in citizen science and in computational sustainability Xue et al. [2016a,b]. Finding solutions in the context of the developing world may therefore have a broader global impact.

Lack of information also leads to inefficiencies in existing systems, presenting a possibility to introduce solutions that abide by existing cultural and technological constraints. For instance, large price discrepancies and major arbitrage opportunities present in markets for agricultural products in Uganda suggest large market inefficiencies Ssekibuule et al. [2013]. To alleviate this, Newman et al. [2018] introduce Kudu—a mobile technology that functions over feature phones via SMS service. Kudu facilitates transactions between farmers in rural areas and buyers at markets in cities by allowing sellers and buyers to post their asks and offers. Kudu has been adapted by users across Uganda and many trades have been realized through this system.

Availability of new technologies also presents opportunities to tackle fundamental problems related to poverty. Advances in last-mile payment technologies, for example, enable large-scale, secure cash transfers. GiveDirectly leverages this and the popularity of mobile money across the world to create a system where donors can directly transfer cash to recipients GD, Blattman and Niehaus [2014]. GiveDirectly moves the decision about how to use aid from policy-makers to recipients, giving recipients maximum flexibility. Such aid generates heterogeneity in outcomes—e.g., families may use aid to start a business, pay rent, cover health costs, and so on. Policy-makers used to prioritizing specific outcomes may be uncomfortable by such a model. A research question then is: can we predict how a given population will use aid? Likewise, how can we target people for whom the interventions will make the largest difference? Aid has historically been targeted on the basis of finding the most deprived people. The ability to model heterogeneous treatment effects opens the door for designing more nuanced mechanisms that fairly and efficiently allocate subsidies in order to maximize a desired outcome.

Problems in the developing world surface unique challenges at the intersection of AI, ICT4D, and development economics. Solutions often have to be implemented in resource-constrained environments (e.g., over feature phones or with low network connectivity) Brunette et al. [2013], Patel et al. [2010]. Key populations of interest (e.g., women, people living in rural parts, individuals with disabilities) may not be easily accessible Sultana et al. [2018], Vashistha et al. [2015b,a]. Individuals may have low-literacy Sambasivan et al. [2010]. Lack of understanding of socio-cultural norms and politics, furthermore, may inhibit proposed interventions Vashistha et al. [2018]. All of these highlight the need for a multi-stakeholder approach that leverages technological advances, innovative technical solutions, and partnerships with individuals and organizations that will be impacted by the solutions. MD4SG fosters one such environment in which insights from across these disciplines inform the design of algorithms and mechanisms to improve the lives of individuals across the world.

Labor, Platforms, and Discrimination

Online platforms are ubiquitous, providing a vast playground for algorithm design and artificial intelligence. Every policy decision, however, impacts and interacts with the platform's strategic users. In this section, we will focus on online labor markets and how discrimination effects stem from a platform's decisions. Past work begins to investigate some aspects of platforms, of strategic agents, and of discrimination in labor markets, but there are still major opportunities for work at the intersection, and insights from mechanism design and AI are ripe for the job.

One central issue surrounding labor markets is that of *hiring*, in which a firm takes information about a potential candidate and makes an employment decision. Firms act as classifiers, labeling each applicant as "hire" or "not hire" based on an applicant's "features," such as educational investment or a worker's productivity reputation. In the process of making hiring decisions, however, the firm may potentially make discriminatory decisions—perhaps by using protected attributes, or by not correcting for differences in applications that stem from systemic discrimination Bertrand and Mullainathan [2004], Marlowe et al. [1996]. Bias in hiring decisions may arise due to implicit human bias or algorithmic bias, in which algorithms replicate human and/or historic discrimination that is reflected in the data on which they are trained Broussard [2018], Eubanks [2018], Noble [2018], O'Neil [2016].

One recent line of work investigates hiring policies that achieve diversity or statistical parity (with respect to certain groups) among the hired workers, and how workers make their investment decisions (e.g. whether to attend college) based on the hiring policies in place. Coate and Loury [1993], Fryer Jr and Loury [2013], Hu and Chen [2017] study settings where there is some known underlying bias or historical discrimination against certain groups; the aim is to characterize hiring policies that are optimal-subject-to-fair-hiring, and to quantify any loss in efficiency compared to optimal-but-discriminatory policies. These works explore two settings: first, when hiring decisions must be "group-blind," that is, they cannot take group membership into account, and second, when they are "group-aware". The aim is to choose hiring policies that will mitigate discrimination against protected categories. Hu and Chen [2017] highlight additional complexity that arises in dynamic settings where workers are hired based on investment decisions (e.g. college GPA) in an initial temporary labor market (e.g. internships) and this job creates a worker's initial productivity reputation that is then used in the permanent labor market. Many of these findings also discuss "trade-offs" between group-blind and group-aware policies.

Another aspect of labor markets is that a worker may have the ability to pay to change a feature of her application in some illegitimate or unfair way in order to improve her outcome in the labor market. Hardt et al. [2016] examine this problem from a robust machine learning perspective. Under certain assumptions of the cost required to change an applicant's reputation, they characterize classifiers that optimally compare to the original reputation (before the applicant modified it).

These are only two aspects at the interplay between hiring and strategic agents; hiring, furthermore, is only one aspect of the labor market. Consider today's popular online labor markets, such as Mechanical Turk, Upwork, Task Rabbit, and Lyft, in which the platform's goal is to match workers to employers or jobs. In these labor markets, the platform's decisions, even at a granular level, have a large impact on the workers and firms. Consider the following platform decisions. Visibility: How many firms can workers see at a time? What capacity do they have to search job offers? Can workers see jobs and jobs see workers? Initiation: Which side (or both) can submit applications? Initiate messaging? Set contract terms? Information: What information is displayed about parties on the opposite side? Name? Photo? Ethnicity? Wage history? Reputation?

Each of these decisions impacts the outcome—not only the quality of the match, but also whether (and how much) discrimination occurs. In a recent paper, Levy and Barocas [2017] outline categories of platform decisions which may mitigate or perpetuate discrimination in labor markets, including the high-level categories of setting platform discrimination policies or norms,

structuring information and interactions, and monitoring/evaluating discriminatory conduct.

In offline labor markets, it may be challenging or infeasible to collect data to understand the nature and extent of discrimination. Online labor markets, on the other hand, yield rich data about employer-employee interactions and present the possibility of conducting experiments aimed at reducing bias and discrimination or other desired societal objectives. For instance, Barach and Horton [2017] look at the impact on hiring of hiding workers' wage history. Horton and Johari [2015], Horton [2018] look at the impact of trying to elicit additional information (features) from workers or firms, and the impact of this strategically-reported information on hiring. Horton [2017] examines who the hired worker population is when a minimum wage is imposed on one platform. Each of these provide insights into labor dynamics that may inform platform design and interventions.

Online labor markets provide a rich playground for techniques from algorithms, AI, and mechanism design to study how each aspect of platform design impacts discriminatory effects, workers' actions, and the desired objective for the platform.

Allocating Housing and Homelessness Resources

Allocation of resources—such as public housing, housing vouchers, and homelessness services—has a long history in the economics and computation literature. Even simple-to-state problems here have given rise to challenging research questions, many of which are still open. Increased scarcity of housing resources, growing need for services, and the use of algorithmic decision-making tools all open up several avenues with major opportunities for reforming policies and regulations. Here, we discuss some foundational work, new challenges, and opportunities that emerge at the nexus of algorithm and mechanism design, AI, and the social sciences.

Millions of individuals across the US have been evicted or are at risk of experiencing eviction every year. In groundbreaking work, Desmond [2012, 2016] shows that eviction is much more common than was previously documented. By compiling the first ever evictions database, Desmond reveals that there is an estimate of 2.3 million evictions in 2016 alone and argues that eviction is a cause of poverty EL [2018]. Using this database, and other similar datasets, we may be able to employ a combination of machine learning and statistical techniques to gain a better understanding of what causes housing instability and homelessness. We can then build on this work to design algorithms and mechanisms that can improve on allocation policies. For instance, Kube et al. [2018] use counter-factual predictions to improve homelessness service provisions. By doing so, they realize some gains on reducing the number of families experiencing repeated episodes of homelessness. At the same time, Eubanks [2018] emphasizes that caution must be taken when using automated decision-making tools for allocating limited resources in such high-stakes scenarios. Such tools may be used to reduce failure rates by caseworkers, but, if not approached with care, can deepen already existing inequalities. Furthermore, the use of such tools alone is limited; it does not address the lack of housing and homelessness resources or eliminate human biases or discrimination. It is crucial to take advantage of the confluence of insights from cross many disciplines in order to serve the needs of such vulnerable populations.

An issue that is growing in prominence in housing contexts is that of information. Little is documented about how landlords or housing authorities screen applications and make decisions. One exception is the work of Ambrose and Diop [2016], which shows that there is increased restriction in access to rental housing since landlords mitigate information asymmetry by investing in screening tenants. With the increased use and availability of data about individuals, it is of paramount importance to understand the role of information in the decision-making process of entities, such as landlords or housing agencies, who have enormous discretion in how and whether families are housed.

Although the introduction of automated tools introduces acute challenges related to housing, the use of algorithmic techniques dates back several decades and there are many fundamental problems that remain unsolved. An early work here is that of Hylland and Zeck-

hauser [1979], which considers the "house allocation problem" of assigning each individual to one item, such as a house. They introduce a mechanism which satisfies natural efficiency and fairness notions but is not incentive-compatible. That is, individuals may be able to improve their outcome by misreporting their true preferences. Since then, several mechanisms have been proposed, including the popular Randomized Serial Dictatorship (RSD) mechanism, which uses a random lottery Abdulkadiroğlu and Sönmez [1998]. This mechanism is used as a standard mechanism in many domains, including housing. While it is incentive-compatible, it fails to satisfy the fairness criteria of Hylland and Zeckhauser [1979]. An important question is then the design of incentive-compatible, fair, and efficient mechanisms for the house allocation problem.

Due to increased scarcity of resources, current allocation protocols often involves waiting lists and priority groups. Policy constraints make wait-list design a dynamic rationing problem rather than the static assignment problem discussed above. Dynamic mechanisms present several technical and practical challenges; e.g., incentive-compatibility may be infeasible in dynamic settings due to waiting time trade-offs for applicants. There are consequential design decisions related to how to manage wait-lists and different metropolitan areas have different policies (e.g., setting priority groups, conditions under which individuals are removed from the waiting list, set of choices, and many others). Each of these policies impacts the allocation dynamics, waiting time, and quality of matches. Recent work has studied how to design mechanisms satisfying various desiderata and quantify differences in quality of matches across various mechanisms Arnosti and Shi [2018], Leshno [2017], Thakral [2016], Waldinger [2017].

Conclusion

As the use of algorithmic and AI techniques becomes more pervasive, there is a growing appreciation of the fact that the most impactful solutions often fall at the interface of various disciplines. The Mechanism Design for Social Good research agenda is to foster an environment in which insights from algorithms and mechanism design can, in conjunction with the social sciences, be used to improve access to opportunity, especially for communities of individuals for whom opportunities have historically been limited. In this piece, we have highlighted MD4SG research avenues related to issues in developing nations, labor markets, and housing. For each of these, we have discussed the need to work in close partnership with a wide range of stakeholders to set objectives that best address the needs of individuals and propose feasible solutions with desired societal outcomes. There are numerous other domains in which this kind of interdisciplinary approach for designing algorithms and mechanisms can improve the lives of many individuals; we invite readers to learn more through our colloquium and workshop series.

Acknowledgments

We are indebted to members of the Mechanism Design for Social Good research group–Ellora Derenoncourt, Alon Eden, Lily Hu, Manish Raghavan, Sam Taggart, Daniel Waldinger, and Matt Weinberg—our co-organizer Irene Lo, and our advisors Anna Karlin and Jon Kleinberg for their generosity in sharing their knowledge and support throughout the past two years. We additionally thank Mutembesa Daniel, Paul Niehaus, Fabian Okeke, and Aditya Vashistha for helpful discussions and pointers. We are grateful, as ever, for the numerous researchers in the economics and computation research communities for their enthusiastic support and guidance throughout the development of the group and workshop series.

References

Givedirectly: Send money directly to the extreme poor. https://givedirectly.org/.

A world that counts: Mobilising the data revolution for sustainable development. Technical report, November 2014. Report prepared at the request of the United Nations Secretary-General by the Independent Expert Advisory Group on Data Revolution for Sustainable Development.

The Eviction Lab. https://evictionlab.org/, 2018.

Atila Abdulkadiroğlu and Tayfun Sönmez.

Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

Rediet Abebe and Kira Goldner. Mechanism Design for Social Good. http://md4sg.com, 2016.

Rediet Abebe and Kira Goldner. A report on the workshop on mechanism design for social good. *ACM SIGecom Exchanges*, 16 (2):2–11, 2018.

Rediet Abebe, Shawndra Hill, Jennifer Wortman Vaughan, Peter M Small, and H Andrew Schwartz. Using search queries to understand health information needs in africa. *arXiv preprint arXiv:1806.05740*, 2018.

Brent W Ambrose and Moussa Diop. Information asymmetry, regulations, and equilibrium outcomes: Theory and evidence from the housing rental market. 2016.

Nick Arnosti and Peng Shi. Design of lotteries and waitlists for affordable housing allocation. 2018.

Moshe Barach and John J Horton. How do employers use compensation history?: Evidence from a field experiment. 2017.

Marianne Bertrand and Sendhil Mullainathan. Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013, 2004.

Christopher Blattman and Paul Niehaus. Show them the money: Why giving cash helps alleviate poverty. *Foreign Affairs*, 93 (3):117–126, 2014.

Meredith Broussard. *Artificial Unintelligence: How Computers Misunderstand the World*. MIT Press, 2018.

Waylon Brunette, Mitchell Sundt, Nicola Dell, Rohit Chaudhri, Nathan Breit, and Gaetano Borriello. Open data kit 2.0: expanding and refining information services for developing regions. In *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications*, page 10. ACM, 2013.

Stephen Coate and Glenn C Loury. Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, pages 1220–1240, 1993.

Matthew Desmond. Eviction and the reproduction of urban poverty. *American Journal of Sociology*, 118(1):88–133, 2012.

Matthew Desmond. *Evicted: Poverty and profit in the American city*. Broadway Books, 2016.

Virginia Eubanks. *Automating Inequality: How High-tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press, 2018.

Roland G Fryer Jr and Glenn C Loury. Valuing diversity. *Journal of Political Economy*, 121 (4):747–774, 2013.

Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122. ACM, 2016.

John J Horton. Price floors and employer preferences: Evidence from a minimum wage experiment. 2017.

John J Horton. Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. 2018.

John J Horton and Ramesh Johari. At what quality and what price?: Eliciting buyer preferences as a market design problem. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 507–507. ACM, 2015.

Lily Hu and Yiling Chen. Fairness at equilibrium in the labor market. *arXiv preprint arXiv:1707.01590*, 2017.

Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political economy*, 87(2): 293–314, 1979.

Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.

Amanda R Kube, Sanmay Das, and Patrick J. Fowler. Allocating interventions based on counterfactual predictions: A case study on homelessness services. 2018.

Jacob Leshno. Dynamic matching in overloaded waiting lists. 2017.

Karen Levy and Solon Barocas. Designing against discrimination in online markets. *Berkeley Tech. LJ*, 32:1183, 2017.

Cynthia M Marlowe, Sandra L Schneider, and Carnot E Nelson. Gender and attractiveness biases in hiring decisions: Are more experienced managers less biased? *Journal of applied psychology*, 81(1):11, 1996.

Daniel Mutembesa, Christopher Omongo, and Ernest Mwebaze. Crowdsourcing real-time viral disease and pest information: A case of nation-wide cassava disease surveillance in a developing country. In *HCOMP*, pages 117–125, 2018.

Ernest Mwebaze and Michael Biehl. Prototype-based classification for image analysis and its application to crop disease diagnosis. In *Advances in Self-Organizing Maps and Learning Vector Quantization*, pages 329–339. Springer, 2016.

Ernest Mwebaze and Godliver Owomugisha. Machine learning for plant disease incidence and severity measurements from leaf images. In *Machine Learning and Applications (ICMLA), 2016 15th IEEE International Conference on*, pages 158–163. IEEE, 2016.

Neil Newman, Lauren Falcao Bergquist, Nicole Immorlica, Kevin Leyton-Brown, Brendan Lucier, Craig McIntosh, John Quinn, and Richard Ssekibuule. Designing and evolving an electronic agricultural marketplace in Uganda. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, page 14. ACM, 2018.

Safiya Umoja Noble. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, 2018.

Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books, 2016.

Neil Patel, Deepti Chittamuru, Anupam Jain, Paresh Dave, and Tapan S Parikh. Avaaj otalo: a field study of an interactive voice forum for small farmers in rural india. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 733–742. ACM, 2010.

John Alexander Quinn, Kevin Leyton-Brown, and Ernest Mwebaze. Modeling and monitoring crop disease in developing countries. In *AAAI*, 2011.

Nithya Sambasivan, Ed Cutrell, Kentaro Toyama, and Bonnie Nardi. Intermediated technology use in developing communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 2583–2592, New York, NY, USA, 2010. ACM.

Richard Ssekibuule, John A Quinn, and Kevin Leyton-Brown. A mobile market for agricultural trade in uganda. In *Proceedings of the 4th Annual Symposium on Computing for Development*, page 9. ACM, 2013.

Sharifa Sultana, François Guimbretière, Phoebe Sengers, and Nicola Dell. Design within a patriarchal society: Opportunities and challenges in designing for rural women in bangladesh. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 536:1–536:13. ACM, 2018.

Neil Thakral. The public-housing allocation problem. Technical report, Technical report, Harvard University, 2016.

Aditya Vashistha, Edward Cutrell, Gaetano Borriello, and William Thies. Sangeet swara: A community-moderated voice forum in rural india. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 417–426. ACM, 2015a.

Aditya Vashistha, Edward Cutrell, Nicola Dell, and Richard Anderson. Social media platforms for low-income blind people in india. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '15, pages 259–272. ACM, 2015b.

Aditya Vashistha, Fabian Okeke, Richard Anderson, and Nicola Dell. 'you can always do better!: The impact of social proof on participant response bias. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 552. ACM, 2018.

Daniel Waldinger. Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. 2017.

Yexiang Xue, Ian Davies, Daniel Fink, Christopher Wood, and Carla P. Gomes. Avicaching: A two stage game for bias reduction in citizen science. In *Proceedings of*

the 2016 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '16, pages 776–785, 2016a.

Yexiang Xue, Ian Davies, Daniel Fink, Christopher Wood, and Carla P. Gomes. Behavior identification in two-stage games for incentivizing citizen science exploration. In Michel Rueher, editor, *Principles and Practice of Constraint Programming*, pages 701–717, Cham, 2016b. Springer International Publishing.

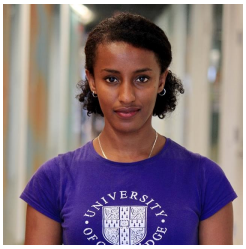

Rediet Abebe is a PhD candidate in computer science, advised by Jon Kleinberg. Her research focuses on algorithms, artificial intelligence, and their applications to social good. In particular, she uses algorithmic and computational insights to better understand socioeconomic inequality and inform interventions for improving access to opportunity. Her work is generously supported by fellowships and scholarships through Facebook, Google, and the Cornell Graduate School. Prior to Cornell, she completed an M.S. and a B.A. in Applied Mathematics and Mathematics from Harvard University and an M.A. in Mathematics from the University of Cambridge. She was born and raised in Addis Ababa, Ethiopia.

Kira Goldner is a fifth-year PhD student at the University of Washington in the Department of Computer Science and Engineering, advised by Anna Karlin. Her research focuses on problems in mechanism design, particularly on (1) maximizing revenue in settings that are motivated by practice and (2) on mechanism design within health insurance and online labor markets. She is a 2017-19 recipient of the Microsoft Research PhD Fellowship and was a 2016 recipient of a Google Anita Borg Scholarship. Kira received her B.A. in Mathematics from Oberlin College and also studied at Budapest Semesters in Mathematics.