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: 10.1145/3465074

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

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 7(1)

Iolanda Leite, co-editor (Royal Institute of Technology (KTH); aimatters@sigai.acm.org)
Anuj Karpatne, co-editor (Virginia Tech; aimatters@sigai.acm.org)
DOI: 10.1145/3465074.3465075

Issue overview

Welcome to the first issue of this year's AI Matters Newsletter!

We start with a report on upcoming SIGAI Events by Dilini Samarasinghe and Conference reports by Louise Dennis, our conference coordination officers. In our regular Education column, Duri Long, Jonathan Moon, and Brian Magerko introduce two "unplugged" activities (i.e., no technology needed) to learn about AI focussed on K-12 AI Education. We then bring you our regular Policy column, where Larry Medsker covers several topics on AI policy, including the role of Big Tech on AI Ethics and an interview with Dr. Eric Daimler who is the CEO of the MIT-spinout Conexus.com.

Finally, we close with four article contributions. The first article discusses emerging applications of AI in analyzing source code and its implications to several industries. The second article discusses topics in the area of physical scene understanding that are necessary for machines to perceive, interact, and reason about the physical world. The third article presents novel practices and highlights from the Fourth Workshop on Mechanism Design for Social Good. The fourth article provides a report on the "Decoding AI" event that was conducted online by ViSER for high school students and adults sponsored by ACM SIGAI.

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.

Iolanda Leite is co-editor of AI Matters. She is an Associate Professor at the School of Electrical Engineering and Computer Science at the KTH Royal Institute of Technology in Sweden. Her research interests are in the areas of Human-Robot Interaction and Artificial Intelligence. She aims to develop autonomous socially intelligent robots that can assist people over long periods of time.

Anuj Karpatne is co-editor of AI Matters. He is an Assistant Professor in the Department of Computer Science at Virginia Polytechnic Institute and State University (Virginia Tech). He leads the Physics-Guided Machine Learning (PGML) Lab at Virginia Tech, where he develops novel ways of integrating scientific knowledge (or physics) with machine learning methods to accelerate scientific discovery from data.

Events

Dilini Samarasinghe (University of New South Wales; d.samarasinghe@adfa.edu.au)
DOI: 10.1145/3465074.3465076

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.For more information about conference support visit sigai.acm.org/activities/requesting_sponsorship.html.

21st ACM International Conference on Intelligent Virtual Agents (IVA 2021)

Kyoto, Japan, 14-17 September, 2021
http://www.iva2021.org/
ACM IVA is an annual interdisciplinary 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 recognised as 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.

The conference seeks to publish cutting-edge research on the design, application, and evaluation of IVAs, as well as basic research on social perception, dialog modeling, and social behavior planning. Further, submissions on central theoretical issues, uses of virtual agents in psychological research and showcases of working applications are also encouraged. Special theme topics include but are not limited to: experiments and methodological issues with Socially Interactive Agents (SIA) and HRI research when interaction is restricted (such as during Covid-19); IVAs for behavior change; IVAs for social inclusion/feelings of loneliness, mental helper IVAs; IVA's for deliberation and information understanding; cognitive aid IVA's; and IVA aided education (remote) education.
Submission deadline: May 1, 2021

18th International Conference on Informatics in Control, Automation and Robotics (ICINCO 2021)

Virtual online event, July 6-8, 2021
http://www.icinco.org/
ICINCO 2021 is targeted at emphasizing the connection between informatics applications pervasive in the areas of Control, Automation and Robotics. It expects to bring together researchers, engineers and practitioners interested in the application of informatics to these areas. The conference facilitates four simultaneous tracks: intelligent control systems and optimization; robotics and automation; signal processing, sensors, systems modeling and control; and industrial informatics.

Companies interested in presenting their products/methodologies or researchers interested in holding a tutorial, workshop or special session can find further information on the conference website. Due to the ongoing COVID-19 pandemic, the conference will be streamed online this year.
Submission deadlines: Doctoral Consortium: May 12, 2021; Special Sessions: May 6, 2021; Abstraction Track: May 12, 2021

36th IEEE/ACM International Conference on Automated Software Engineering (ASE 2021)

Melbourne, Australia, November 15-19, 2021
https://conf.researchr.org/home/ase-2021
ASE 2021 is the premier research forum for Automated Software Engineering. Each year, it brings together researchers and practitioners from academia and industry to discuss foundations, techniques, and tools for automating the analysis, design, implementation, testing, and maintenance of large software systems. In addition to the regular conference tracks (research track, doctoral symposium, journal-first papers, tutorials, workshops), the conference is also encouraging submissions to several other tracks including: artifact evaluation, industry showcase, new ideas and emerging results, tool demonstra-

tions and late breaking news.
Submission deadlines: Research Papers: April 23, 2021; Late Breaking News: May 28, 2021; Tool Demonstrations: June 11, 2021; Doctoral Symposium: July 13, 2021; Artifact Evaluation: July 14, 2021

1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'21)

Virtual online event, October 5-8, 2021
http://www.eaamo.org
The inaugural ACM conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'21) aims to highlight work where techniques from algorithms, optimization, and mechanism design, along with insights from other disciplines, can help improve equity and access to opportunity for historically disadvantaged and underserved communities. The conference is organized by the Mechanism Design for Social Good (MD4SG) initiative, and builds on the MD4SG workshop series and tutorials at conferences including ACM EC, ACM COMPASS, and WINE. EAAMO'21 will feature keynote presentations and panels and contributed presentations on research papers, surveys, problem pitches, datasets, and software demonstrations. In line with the MD4SG core values of bridging research and practice, the conference aims to provide an international forum for researchers as well as policy-makers and practitioners in various government and non-government organizations, community organizations, and industry to build interdisciplinary, multi-stakeholder research pipelines.
Submission deadline: June 3, 2021

23rd International Conference on Enterprise Information Systems (ICEIS 2021)

Virtual online event, April 26-28, 2021
http://www.iceis.org
ICEIS aims to bring together researchers, engineers and practitioners interested in the advances and business applications of information systems. Six simultaneous tracks will be held, covering different aspects of enterprise information systems applications, including enterprise database technology, systems integration, artificial intelligence, decision support systems, information systems analysis and specification, internet computing, electronic commerce, human factors and enterprise architecture. A special session on soft computing for smart cities services is also planned to be held at the conference as a set of oral and poster presentations with the intention of providing a focused discussion on the specialized theme.
Submission deadlines: Doctoral Consortium: March 2, 2021; Abstracts Track: March 2, 2021; Special Session: March 2, 2021

34th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE 2021)

Kuala Lumpur, Malaysia, July 26-29, 2021
https://ieeecomputer.my/ieaaie2021
IEA/AIE 2021 is the 34th event continuing the tradition of emphasizing on applications of applied intelligent systems to solve real-life problems in all areas including engineering, science, industry, automation & robotics, business & finance, medicine and biomedicine, bioinformatics, cyberspace, and human-machine interactions. The topics considered for publication include but are not limited to: adaptive control, autonomous agents, computer vision, data mining, evolutionary computation, games, information retrieval, machine learning and robotics with eight special sessions that run in parallel with the main conference.

Due to the ongoing Covid-19 pandemic both physical and virtual presentations are facilitated at the conference. Please refer to the conference website for information on registration and the program.

16th International Conference on Foundations of Digital Games (FDG 2021)

Virtual online event, August 3-6, 2021
http://fdg2021.org/
FDG 2021 is focused on presenting contributions from within and across disciplines committed to advancing knowledge on the foundations of games: computer science and engineering, humanities and social sciences,

arts and design, mathematics and natural sciences. The conference theme for 2021 is Diversity & Inclusion through Games. As game playing is increasingly being enjoyed by people of all genders, ages, ethnicities, and social economic backgrounds, it is expected to broaden the conversation to also include how games themselves can be used to increase diversity and inclusion in industry, academia, and society through FDG 2021. Papers are solicited across eight tracks discussing topics related to game development methods and technologies, analytics, artificial intelligence, game design and player experience.

The event will be hosted as a single-track virtual conference due to the ongoing Covid-19 pandemic. Please refer to the conference website for details on registration and the program.

4th AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES 2021)

Virtual online event, May 19-21, 2021
https://www.aies-conference.com/2021/
AIES 2021 is aimed at encouraging talented scholars in Computer Science, Law and Policy, the Social Sciences, Philosophy, and related fields to focus on morality, law, and political economy of data and AI. The conference is tailored for a multi-disciplinary audience and also welcomes disciplinary experts who are newer to this topic, and see ways to break new ground in their own fields by thinking about data and AI. The topics of interest to the conference include: empirical and evaluative research into impacts of AI systems; goals at which to be aimed when redesigning AI systems; representation, acquisition, and use of ethical knowledge by AI systems; and sociotechnical, legal and regulatory approaches for realising evaluative goals.

Please refer to the conference website for details on registration and the program.

Dilini Samarasinghe is the Assistant Conference Coordination Officer for ACM SIGAI, and a post-doctoral research associate at the University of New South Wales. Her research is in Artificial Intelligence, Multi-agent Systems and Evolutionary Computation. Contact her at d.samarasinghe@adfa.edu.au.

Conference Reports

Louise A. Dennis (University of Manchester; louise.dennis@manchester.ac.uk)
DOI: 10.1145/3465074.3465077

This section features brief reports from recent events sponsored or run in cooperation with ACM SIGAI.

IEEE/ACM International Conference on Automated Software Engineering (ASE2020)

Online, September 21-25, 2020
https://conf.researchr.org/home/ase-2020

The ASE conference is the premier research forum for Automated Software Engineering. Each year, it brings together researchers and practitioners from academia and industry to discuss foundations, techniques, and tools for automating the analysis, design, implementation, testing, and maintenance of large software systems.

Key objectives for ASE 2020 were to:

- Host a very high quality scientific programme
- Attract industry practitioners to the participate in the event
- Build industry/academic collaboration
- Develop students and junior faculty in the ASE community
- Support various networking and exchange of ideas for ASE community members

ASE 2020 was *supposed* to be held in Melbourne Australia at the Arts Centre, 21-25 September 2020. However, the COVID-19 pandemic meant the organizing committee, ACM and IEEE agreed it must be held in fully virtual mode. This was due to closure of the Australian borders to international arrivals and a stage 4 lockdown in the state of Victoria during the conference week.

ASE 2020 used the Whova platform to run the conference virtually. It ended up with 400 attendees in total – significantly higher than previous years in-person – and run 4 workshops, 4 tutorials, 3 keynotes, had almost 150 main conference paper presentations, and a variety of on-line networking events. These included early career researcher networking, doctoral symposium, PhD student development, industry/academic collaboration, 3 ask-me-anything sessions with senior ASE community members, a Women in ASE workshop, and interactive lightening talk/demo sessions for posters and tool demonstrations. Post-conference videos of talks, where author permission was given, are now hosted on the ACM YouTube channel with a ASE2020 playlist.

Overall the conference was successful, as highlighted in the post-conference survey.

2020 ACM International Conference on Intelligent Virtual Agents

Online, October 20-22, 2020
https://iva2020.psy.gla.ac.uk

The 2020 ACM International Conference on Intelligent Virtual Agents, (IVA) was the 20th meeting of the interdisciplinary conference. It was originally planned to be hosted at the University of Glasgow, Scotland but was instead hosted online

IVA is the leading scientific forum for presenting research on the modeling, development and evaluation of Intelligent Virtual Agents (IVAs) with a specific focus on communicative abilities and social behavior. IVAs are interactive characters that exhibit human-like qualities including communicating using natural human modalities such as facial expressions, speech and gesture. IVAs are also capable of real-time perception, cognition, emotion and action that allows them to participate in dynamic social situations.

The specific focus of IVA 2020 was to further increase participation, submissions and interdisciplinary participation. The special theme of the conference was *Exploring Connections across Disciplines*. The goal was to explore connections between scientific communities interested in IVAs but often not participating in joint forums, specifically AI/Computer Science, Robotics and Psychology. IVA 2020 undertook several strategies to realize these goals, including allowing abstract only submissions to receive full presentations in order to encourage psychologists to participate without negatively impacting their ability to publish a journal article and putting together a highly Interdisciplinary list of Invited speakers including a researcher in human-robot interaction, a neuroscientist, and a panel Artists/Architects using AI to realize interactive work. Specifically,

- AI Powered Interactive Art/Architecture Panel
  - Behnaz Farahi, Ph.D.: Emotive Matter: Affective Computing from Fashion to Architecture
  - Güvenç Özel: Persuasion Machines: Networked Architectures in the Age of Surveillance Capitalism
- Professor Jodi Forlizzi: HRI and HAI: Merging Perspectives from Two Fields
- Professor Lars Muckli: Emotions and other contextual signals in early visual cortex and the computational role for AI

Despite the Covid-19 crisis, attendance and submission were higher than usual, interdisciplinary participation was much higher than the norm. IVA 2020 was complimented on the social interaction, the interdisciplinary participation and the invited talks.

Louise Dennis is the Conference Coordination Officer for ACM SIGAI, and a faculty member at the University of Manchester. Her research is in Verification of Autonomous Systems, Cognitive Agents and Machine Ethics. Contact her at louise.dennis@manchester.ac.uk.

Unplugged Assignments for K-12 AI Education

Duri Long (Expressive Machinery Lab, Georgia Institute of Technology; duri@gatech.edu)

Jonathan Moon (Expressive Machinery Lab, Georgia Institute of Technology; jonathanmoon-design@gmail.com)

Brian Magerko (Expressive Machinery Lab, Georgia Institute of Technology; magerko@gatech.edu)

DOI: 10.1145/3465074.3465078

Introduction

In this column, we introduce our two "unplugged" (i.e. no technology needed) Model AI Assignments: Introducing AI and Semantic Networks and Knowledge Representations. We also reflect on the potential benefits of unplugged activities for broadening access to AI-related learning experiences.

Why "Unplugged" Activities to Learn about AI?

Resources for computer science (CS) education that do not require technology have become valuable in computing education for a variety of reasons, including their low cost, ease of implementation, incorporation of physical/embodied interaction, and often playful nature (Nishida et al., 2009). Inspired by these "CS Unplugged" materials (Bell, Rosamond, & Casey, 2012), there have been a few existing online resources for AI education developed in the past year or two that do not require technology. Ali et al. have developed an unplugged middle-school curriculum for AI ethics (Ali, Payne, Williams, Park, & Breazeal, 2019) and Lindner et al. have developed a six-lesson unplugged curriculum for teaching about concepts like decision trees and reinforcement learning (Lindner, Seegerer, & Romeike, 2019). A few other resources for unplugged AI resources that have not been formally published have been recently made available as lesson plans online (Microsoft, n.d.; Group, n.d.; Krueger, n.d.; Seegerer & Lindner, n.d.).

There remains a lot of space in the field for the development of additional "unplugged" AI activities. These activities have the potential to broaden access to AI-related learning experiences at a low cost to educators. They also have the potential to be more engaging for novice audiences, since they involve hands-on paper-based activities that typically do not require prerequisite coding knowledge.

The Model AI Assignments

In this column, we present two unplugged assignments that aim to teach introductory AI concepts to young learners with no prior experience in AI or computer science.

Introducing AI

The first assignment, Introducing AI, is intended as a high-level introduction to AI and can be used to kick off an AI-related class, unit, or workshop. Students engage in an interactive worksheet activity and explore questions such as: *What is artificial intelligence?*; *Where have you used it before?*; *How do you feel about it?*; *How does it work?*. This assignment can be completed as a worksheet activity or the worksheet can be used as a guide to lead an in-class activity.

The activity requires the use of a printable deck containing cards with *examples of AI technologies* and possible *inputs*, *algorithms*, and *outputs* for AI devices. The worksheet activity begins by prompting students to consider where they have seen AI before and how they feel about AI. Then, they are asked to look at the examples of AI in the card deck and select cards with technologies they have interacted with previously. Finally, students are walked through a high-level explanation of how AI works and are guided by the worksheet to create an imaginary AI using the *input*, *algorithm*, and *output* cards. Our assignment also includes additional questions for learners to reflect on after the activity, such as prompting learners to reflect on the strengths and weaknesses of their imaginary AI device or asking learners to discuss whether they think their feelings about AI might change in the future.

In our experience, this activity has been engaging for family groups with children ages 6 and up. Younger children (6-9) needed support from adults or older siblings to read the worksheet instructions/card descriptions but were able to actively engage in identifying AI examples they had previously interacted with, discussing their feelings about AI, and finding sensor/dataset/algorithm cards to create an imaginary AI (mostly using the pictures on the front as a guide).

This activity aims to equip learners to be able to a) identify AI technologies that they have used before; b) distinguish between technology that uses AI and technology that does not use AI; c) identify their preconceptions and discuss their feelings about AI; d) define the terms sensor, dataset, and algorithm and recognize several examples of each; and e) explain that AI takes an input, processes that input using an algorithm, and produces an output.

Semantic Networks and Knowledge Representations

The second assignment, Semantic Networks and Knowledge Representations, is focused on communicating concepts related to knowledge representations and reasoning. AI agents store and organize information in their memory using structures known as *knowledge representations*. One type of knowledge representation is a *semantic network*. Semantic networks are a way of representing relationships between objects and ideas. For example, a network might tell a computer the relationship between different animals (e.g. a cat IS-A mammal; a cat HAS whiskers). In this assignment, learners can create their own semantic networks (Figure 1) by gluing down printable cards containing concepts (e.g. cat, mom, friend) and arrows containing relationships (is, has, likes, dislikes). Provided card decks contain concepts related to animals, family, and musical instruments. Blank cards are also provided to allow learners to make networks on custom topics.

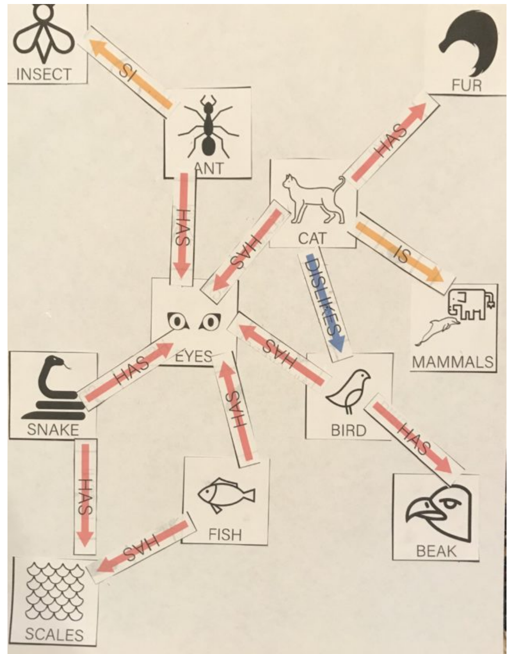

Figure 1: Example of a completed semantic network activity

Students can simulate an AI-user interaction using their semantic networks. Two students can trade completed semantic networks and ask their partner questions about the network they created (e.g. "What is a cat?"). The student's partner should answer the questions using the semantic network as their only guide (simulating an AI agent whose only knowledge is based on the semantic network). Learners are then encouraged to reflect on the networks they create and consider the strengths and limitations of the knowledge representation using a provided list of questions that can be used to foster discussion or as a written activity.

In our experience, this activity has been engaging for learners ages 6 and up, although younger learners many need some adult support during the latter half of the activity when they are asked to simulate an AI agent using their network. We have also observed that the activity is engaging and fosters learning for learners with little to no prior knowledge about AI. This assignment could be adapted as either a take-home written activity or an in-class group project, and cards/arrows could be customized to foster interdisciplinary connections.

This activity aims to help learners a) understand that one way computers store commonsense knowledge is using networks of connected concepts and relationships; b) explain at a high-level how a computer would use the network they built to answer questions (e.g. follow the HAS arrows connected to "cat" to answer "What does a cat have?"); and c) reflect on the strengths and limitations of semantic networks as a way of representing knowledge.

References

Ali, S., Payne, B. H., Williams, R., Park, H. W., & Breazeal, C. (2019). Constructionism, ethics, and creativity: Developing primary and middle school artificial intelligence education. In *International workshop on education in artificial intelligence k-12 (eduai'19).*

Bell, T., Rosamond, F., & Casey, N. (2012). Computer science unplugged and related projects in math and computer science popularization. In *The multivariate algorithmic revolution and beyond* (pp. 398–456). Springer.

Group, C. E. R. (n.d.). *Artificial Intelligence.* Retrieved from https://classic.csunplugged.org/artificial-intelligence/

Krueger, N. (n.d.). *3 unplugged activities for teaching about AI.*

Lindner, A., Seegerer, S., & Romeike, R. (2019). Unplugged Activities in the Context of AI. In *International Conference on Informatics in Schools: Situation, Evolution, and Perspectives* (pp. 123–135). Springer.

Microsoft. (n.d.). *Unplugged activity: Paper AI.* Retrieved from https://minecraft.makecode.com/courses/csintro/ai/unplugged

Nishida, T., Kanemune, S., Idosaka, Y., Namiki, M., Bell, T., & Kuno, Y. (2009). A CS unplugged design pattern. *ACM SIGCSE Bulletin*, *41*(1), 231–235. (Publisher: ACM New York, NY, USA)

Seegerer, S., & Lindner, A. (n.d.). *AI Unplugged.* Retrieved from https://www.aiunplugged.org/

Duri Long is a Ph.D. candidate in Human Centered Computing at Georgia Institute of Technology. Her research focuses on issues related to human-centered AI, including fostering public AI literacy and human-AI interaction through co-creative, embodied experiences.

Jonathan Moon is an Industrial Designer with a background in user experience research and traditional product design. Jonathan was the lead designer of the card decks for both activities presented in this column.

Brian Magerko is a Professor of Digital Media, Director of Graduate Studies in Digital Media, and head of the Expressive Machinery Lab at Georgia Tech. His research explores how studying human and machine cognition can inform the creation of new human/computer creative experiences.

AI Policy Matters

Larry Medsker (The George Washington University; lrm@gwu.edu)
DOI: 10.1145/3465074.3465079

Abstract

AI Policy Matters 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/`). We welcome everyone to make blog comments so we can develop a rich knowledge base of information and ideas representing the SIGAI members.

FR and Bad Science: Should some research not be done?

Facial recognition issues continue to appear in the news, as well as in scholarly journal articles, while FR systems are being banned and some research is shown to be bad science. AI system researchers who try to associate facial technology output with human characteristics are sometimes referred to as machine-assisted phrenologists. Problems with FR research have been demonstrated in machine learning research such as work by Steed and Caliskan in "A set of distinct facial traits learned by machines is not predictive of appearance bias in the wild." Meanwhile many examples of harmful products and misuses have been identified in areas such as criminality, video interviewing, and many others. Some communities have considered bans.

Yet, journals and conferences continue to publish bad science in facial recognition.

Some people say the choice of research topics is up to the researchers – the public can choose not to use the products of their research. However, areas such as genetic, biomedical, and cybersecurity R&D do have limits. Our professional computing societies can choose to disapprove research areas that cause harm. Sources of mitigating and preventing irresponsible research being introduced into the public space include:

- Peer pressure on academic and corporate research and development
- Public policy through laws and regulations
- Corporate and academic self-interest – organizations' bottom lines can suffer from bad behavior and publicity
- Vigilance by journals about publishing papers that promulgate the misuse of FR.

A recent article by Matthew Hutson in The New Yorker discusses "Who should stop unethical AI." He remarks that "Many kinds of researchers—biologists, psychologists, anthropologists, and so on—encounter checkpoints at which they are asked about the ethics of their research. This doesn't happen as much in computer science. Funding agencies might inquire about a project's potential applications, but not its risks. University research that involves human subjects is typically scrutinized by an I.R.B., but most computer science doesn't rely on people in the same way. In any case, the Department of Health and Human Services explicitly asks I.R.B.s not to evaluate the possible long-range effects of applying knowledge gained in the research, lest approval processes get bogged down in political debate. At journals, peer reviewers are expected to look out for methodological issues, such as plagiarism and conflicts of interest; they haven't traditionally been called upon to consider how a new invention might rend the social fabric."

Big Issues

Big Tobacco, Big Oil . . . and Big Tech

A larger discussion is growing out of the recent news about Dr. Timnit Gebru and Google. Big Tech is having a huge impact on individuals and society both for the many products and services we enjoy and for the current and potential cases of detrimental effects of unethical behavior or naiveté regarding AI ethics issues. How do we achieve AI ethics responsibility in all organizations, big and small? And, not just in corporations, but governmental and academic research organizations?

Some concerned people focus on regulation, but for a variety of reasons public and community pressure may be quicker and more acceptable. This includes corporations earning reputations for ethical actions in the design and development of AI products and systems. An article in MIT Technology Review by Karen Hao discusses a letter signed by nine members of Congress that "sends an important signal about how regulators will scrutinize tech giants." Ideally our Public Policy goal is strong AI Ethics in national and global communities that self-regulate on AI ethical issues, comparable to other professional disciplines in medical science and cybersecurity. Our AI Ethics community, as guidelines evolve, could provide a supportive and guiding presence in the implementation of ethical norms in the research and development in AI. The idea of a global AI Ethics community is reflected also in a recent speech by European Union President Ursula von der Leyen at the World Leader for Peace and Security Award ceremony. She advocates for transatlantic agreements on AI.

What Can Biden Do for Science?

A *Science—Business* Webcast presented a forum of public and private sector leaders discussing ideas about the need for the president-elect to convene world leaders to re-establish "rules of engagement" on science. Participants in the Webcast urged that a global assembly "should press leaders of the big industrial nations to open – or reopen – their research systems, while also ensuring that COVID-19 vaccines are freely available to everyone in the world." About an international summit, Robert-Jan Smits, former director-general of the European Commission's research and innovation directorate said it "would really show that senior leaders are turning the page."

HCAI for Policymakers

"Human-Centered AI" by Ben Shneiderman was recently published in *Issues in Science and Technology* 37, no. 2 (Winter 2021): 56–61. A timely observation is that Artificial Intelligence is clearly expanding to include human-centered issues from ethics, explainability, and trust to applications such as user interfaces for self-driving cars. The importance of the HCAI fresh approach, which can enable more widespread use of AI in safe ways that promote human control, is acknowledged by the appearance in NAS Issues in Science and Technology. An implication of the article is that computer scientists should build devices to enhance and empower—not replace—humans.

HCAI as described by Prof. Shneiderman represents a radically different approach to systems design by imagining a different role for machines. Envisioning AI systems as comprising machines and people working together is a much different starting point than the assumption and goal of autonomous AI. In fact, a design process with this kind of forethought might even lead to a product not being developed, thus preventing future harm. One of the many interesting points in the NAS Issues article is the observation about the philosophical clash between two approaches to gaining knowledge about the world—Aristotle's rationalism and Leonardo da Vinci's empiricism—and the connection with the current perspective of AI developers: "The rationalist viewpoint, however, is dominant in the AI community. It leads researchers and developers to emphasize data-driven solutions based on algorithms." Data science in particular unfortunately often focuses on the rationalist approach without including the contributions from and protection of the human experience.

From the NAS article, HCAI is aligned with "the rise of the concept of design thinking, an approach to innovation that begins with empathy for users and pushes forward with humility about the limits of machines and people. Empathy enables designers to be sensitive to the confusion and frustration that users might have and the dangers to people when AI systems fail. Humility leads designers to recognize the inevitability of failure and inspires them to be always on the lookout for what wrongs are preventable." Policymakers need to "understand HCAI's promise not only for our machines but for our lives. A good starting place is an appreciation of the two competing philosophies that have shaped the development of AI, and what those imply for the design of new technologies ... comprehending these competing imperatives can provide a foundation for navigating the vast thicket of ethical dilemmas now arising in the machine-learning space." An HCAI approach can incorporate creativity and innovation into AI systems by understanding and incorporating human insights about complexity into the design of AI systems and using machines to prepare data for taking advantage of human insight and experience. For many more details and enjoyable reading, see the NAS Issues article.

AI Centre of Excellence (AICE)

AICE conducted their inaugural celebration in December, 2020. Director John Kamara founded the AI Centre of Excellence in Kenya and is passionate about creating value and long term impact of AI and ML in Africa. The Centre aims to accomplish this by providing expert training to create skilled and employable AI and ML engineers. The Centre dives into creating sustainable impact through Research and Development. AI research and products are estimated to contribute over 13 trillion dollars to the global economy by 2030. This offers the Centre an opportunity to carry out research in selected sectors and build products based on the research. The world has around 40K AI experts in the world, with nearly half in the US and less than 5 percent in Africa. Oxford Insights estimates that Kenya ranks first in Africa, and AICE aims to leverage this potential and transform AICE into a full blown Artificial Intelligence Centre of Excellence. Please keep your eyes on Africa and ways our public policy can assist efforts there to grow AI in emerging education and research.

Data for AI: Interview with Dr. Eric Daimler

I recently spoke with Dr. Eric Daimler about how we can build on the framework he and his colleagues established during his tenure as a contributor to issues of AI policy in the White House during the Obama administration. Eric is the CEO of the MIT-spinout Conexus.com and holds a PhD in Computer Science from Carnegie Mellon University. Here are the interesting results of my interview with him. His ideas are important as part of the basis for ACM SIGAI Public Policy recommendations.

LRM: What are the main ways we should be addressing this issue of data for AI?

EAD: To me there is one big re-framing from which we can approach this collection of issues, prioritizing data interoperability within a larger frame of AI as a total system. In the strict definition of AI, it is a learning algorithm. Most people know of subsets such as Machine Learning and subsets of that called Deep Learning. That doesn't help the 99 percent who are not AI researchers. When I have spoken to non-researchers or even researchers who want to better appreciate the sensibilities of those needing to adopt their technology, I think of AI as the interactions that it has. There is the collection of the data, the transportation of the data, the analysis or planning (the traditional domain in which the definition most strictly fits), and the acting on the conclusions. That sense, plan, act framework works pretty well for most people.

LRM: Before you explain just how we can do that, can you go ahead and define some of your important terms for our readers?

EAD: AI is often described as the economic engine of the future. But to realize that growth, we must think beyond AI to the whole system of data, and the rules and context that surround it: our data infrastructure (DI). Our DI supports not only our AI technology, but also our technical leadership more generally; it underpins COVID reporting, airline ticket bookings, social networking, and most if not all activity on the internet. From the unsuccessful launch of healthcare.gov, to the recent failure of Haven, to the months-long hack into hundreds of government databases, we have seen the consequences faulty DI can have. More data does not lead to better outcomes; improved DI does.

Fortunately, we have the technology and foresight to prevent future disasters, if we act now. Because AI is fundamentally limited by the data that feeds it, to win the AI race, we must build the best DI. The new presidential administration can play a helpful role here, by defining standards and funding research into data technologies. Attention to the need for better DI will speed responsiveness to future crises (consider COVID data delays) and establish global technology leadership via standards and commerce. Investing in more robust DI will ensure that anomalies, like ones that would have helped us identify the Russia hack much sooner, will be evident, so we can prevent future malfeasance by foreign actors. The US needs to build better data infrastructure to remain competitive in AI.

LRM: So how might we go about prioritizing data interoperability?

EAD: In 2016, the Department of Commerce (DOC) discovered that on average, it took six months to onboard new suppliers to a midsize trucking company—because of issues with data interoperability. The entire American economy would benefit from encouraging more companies to establish semantic standards, internally and between companies, so that data can speak to other data. According to a DOC report in early 2020, the technology now exists for mismatched data to communicate more easily and data integrity to be guaranteed, thanks to a new area of math called Applied Category Theory (ACT). This should be made widely available.

LRM: And what about enforcing data provenance?

EAD: As data is transformed across platforms—including trendy cloud migrations—its lineage often gets lost. A decision denying your small business loan can and should be traceable back to the precise data the loan officer had at that time. There are traceability laws on the books, but they have been rarely enforced because up until now, the technology hasn't been available to comply. That's no longer an excuse. The fidelity of data and the models on top of them should be proven—down to the level of math—to have maintained integrity.

LRM: Speaking more generally, how can we start to lay the groundwork to reap the benefits of these advancements in data infrastructure?

EAD: We need to formalize. When we built 20th century assembly lines, we established in advance where and how screws would be made; we did not ask the village blacksmith to fashion custom screws for every home repair. With AI, once we know what we want to have automated (and there are good reasons to not to automate everything!), we should then define in advance how we want it to behave. As you read this, 18 million programmers are already formalizing rules across every aspect of technology. As an automated car approaches a crosswalk, should it slow down every time, or only if it senses a pedestrian? Questions like this one—across the whole economy—are best answered in a uniform way across manufacturers, based on standardized, formal, and socially accepted definitions of risk.

LRM: In previous posts, I have discussed roles and responsibilities for change in the use of AI. Government regulation is of course important, but what roles do you see for AI tech companies, professional societies, and other entities in making the changes you recommend for DI and other aspects of data for AI?

EAD: What is different this time is the abruptness of change. When automation technologies work, they can be wildly disruptive. Sometimes this is very healthy (see: Schumpeter). I find that the "go fast and..." framework has its place, but in AI it can be destructive and invite resistance. That is what we have to watch out for. Only with responsible coordinated action do we encourage adoption of these fantastic and magical technologies. Automation in software can be powerful. These processes need not be linked into sequences just because they can. That is, just because some system can be automated does not mean that it should. – Too often there is absolutism in AI deployments when what is called for in these discussions is nuance and context. For example, in digital advertising my concerns are around privacy, not physical safety. When I am subject to a plane's autopilot, my priorities are reversed.

With my work in the US Federal Government, my bias remains against regulation as a first-step. Shortly after my time with the Obama Whitehouse, I am grateful to have participated with a diverse group for a couple of days at the Halcyon House in Washington D.C. We created some principles for deploying AI to maximize adoption. We can build on these and rally around a sort of LEED-like standard for AI deployment.

–

Dr. Eric Daimler is CEO and Founder of Conexus and Board Member of Petuum and WelWaze. He was a Presidential Innovation Fellow, Artificial Intelligence and Robotics. Eric is a leading authority in robotics and artificial intelligence with over 20 years of experience as an entrepreneur, investor, technologist, and policymaker. Eric served under the Obama Administration as a Presidential Innovation Fellow for AI and Robotics in the Executive Office of President, driving the agenda for U.S. leadership in research, commercialization, and public adoption of AI and Robotics. His newest venture, Conexus, is a groundbreaking solution for what is perhaps today's biggest information technology problem — data deluge. Eric works to empower communities and citizens to leverage robotics and AI to build a more sustainable, secure, and prosperous future. His academic research has been at the intersection of AI, Computational Linguistics, and Network Science (Graph Theory). He has studied at the University of Washington-Seattle, Stanford University, and Carnegie Mellon University, where he earned his Ph.D. in Computer Science.

Please join our discussions at the

SIGAI Policy Blog

Larry Medsker is a Research Professor at The George Washington University, where he was founding director of the Data Science graduate program. He is currently a faculty member in the GW Human-Technology Collaboration Lab and Ph.D. program. His research in AI includes work on artificial neural networks, hybrid intelligent systems, and the impacts of AI on society and policy. He is Co-Editor-in-Chief for the journal *AI and Ethics* and the Public Policy Officer for the ACM SIGAI.

Applied AI Matters
AI4Code: Applying Artificial Intelligence to Source Code

Kartik Talamadupula (IBM Research AI; krtalamad@us.ibm.com)
DOI: 10.1145/3465074.3465080

Introduction

The marriage of Artificial Intelligence (AI) techniques to problems surrounding the generation, maintenance, and use of source code has come to the fore in recent years as an important AI application area[1]. A large chunk of this recent attention can be attributed to contemporaneous advancements in Natural Language Processing (NLP) techniques and sub-fields. The *naturalness hypothesis*, which states that "software is a form of human communication" and that code exhibits patterns that are similar to (human) natural languages (Devanbu, 2015; Hindle, Barr, Gabel, Su, & Devanbu, 2016), has allowed for the application of many of these NLP advances to code-centric usecases. This development has contributed to a spate of work in the community – much of it captured in a survey by Allamanis, Barr, Devanbu, and Sutton (2018) that focuses on classifying these approaches by the type of probabilistic model applied to source code.

This increase in the variety of AI techniques applied to source code has found various manifestations in the industry at large. Code and software form the backbone that underpins almost all modern technical advancements: it is thus natural that breakthroughs in this area should reflect in the emergence of real world deployments.

AI4Code: Industrial Applications

There are several characterizations and groupings that can be made when considering applications of AI4Code. One that has already been discussed is predicated on the kinds of probabilistic models of code that are generated and exploited. Another is in terms of whether AI techniques are addressing code usecases (AI4Code); if code is being used to make AI problems easier to solve (Code4AI); or if AI/ML techniques are themselves being used to improve AI tools and lifecycles (AI4AI, or AutoAI/AutoML). Then there are classifications that are based on the specific location in the development-devops cycle where the AI techniques are being injected: this can be any stage starting from requirements gathering, through code generation and documentation, translation, testing, execution, and deployment. A final classification is in terms of the user role that is targeted by these AI4Code manifestations: some might target developers, while others target devops personas.

In the following, we present a non-exhaustive list of some recent AI4Code tools from a wide cross-section of industrial and applied research settings in order to introduce readers to the variety of AI applications in this space.

- **AutoAI** is a stream of work that applies AI techniques to automate machine learning and data science pipelines; recent work has looked at the issues inherent with humans in the loop in such end-to-end lifecycles (Wang et al., 2020).
- **CLAI**[2] is an open-source project that brings AI advances to the command line to automate and ease developer and devops usecases (Agarwal, Barroso, et al., 2020).
- **CodeGuru**[3] is a tool that provides developer oversight of code, particularly with an eye towards efficiency and cost.
- **CriticalHop**[4] is a commercial AI planning engine that seeks to help users minimize Kubernetes and cloud deployment issues.
- **DeepCode**[5] is an AI enabled code review engine that performs semantic code analysis to find critical issues and vulnerabilities.

[1] This report builds on the content of a panel at the NeurIPS 2020 Industrial Expo on *AI4Code at IBM + RedHat*, hosted by the author.

[2] https://github.com/ibm/clai

[3] https://aws.amazon.com/codeguru/

[4] https://www.criticalhop.com/

[5] https://www.deepcode.ai/

- **Graph4Code**[6] is a knowledge graph over code from Python programs that captures the semantics of that code.
- **IntelliCode**[7] is an IDE plugin that tries to enable faster and more efficient coding for developers.
- **Kite**[8] is an AI-powered code completion assistant that aims to minimize developer keystrokes.
- **ModelOps** is a cloud-based platform for end-to-end development and lifecycle management of AI applications (Hummer et al., 2019).
- **Mono2Micro**[9] uses AI techniques to offer recommendations to refactor monolithic code into microservices.
- **Thoth**[10] uses AI techniques to analyze and recommend software stacks for AI applications during deployment.

Code Translation with Neural Models: A Human in the Loop Case Study

While the transfer of AI techniques into industrial tools for code has kept up a brisk pace, there is still much work to be done in terms of evaluating the impact of all these groundbreaking techniques and enriched tools. Specifically, the effect that these tools have on developers – who are the users at the center of this AI4Code revolution – remains to be fully measured. There have been nascent efforts in this space, e.g. the HAI-GEN workshop series[11] at the ACM Intelligent User Interfaces (IUI) conference; and Xu, Vasilescu, and Neubig (2021)'s work.

One recent area of interest under the AI4Code umbrella has been the problem of code translation – that is, automatically translating source code in one language to another. Code translation has applications in many important scenarios, including the modernization of legacy code that runs critical infrastructure and applications in industries such as finance, travel, and government. There have been a number of AI4Code approaches (Nguyen, Nguyen, & Nguyen, 2014; Oda et al., 2015) that seek to cast the code translation problem as a special case of the more general *machine translation* task, which entails automatic translation between two (human) natural languages. Recently, unsupervised machine translation techniques have been applied to the code translation task, with great success. This is exemplified by the TransCoder (Roziere, Lachaux, Chanussot, & Lample, 2020) system, which trains a fully unsupervised neural transcompiler to translate functions between Java, C++, and Python.

The TransCoder model, which is a sequence to sequence (seq2seq) model (Sutskever, Vinyals, & Le, 2014), generates tokens at inference time that together make up the translation of a given input (source) function. The model is able to generate multiple translations using beam search decoding; this information can be used to compute token-level confidence scores for each token produced by the model. These confidence scores can then be shown to the human end-user in the absence of ground truth about a specific translation.

However, this raises an interesting issue: what is a human user to infer from these confidences? Apart from being able to order them ordinally, there is no correlation between the task at hand (code translation) and the output of the model. Agarwal, Talamadupula, et al. (2020) address this problem, and seek to *anchor* the confidence scores by correlating them to linter errors that are generated by the translated code. The thesis underlying their work is that human users – particularly those who are going to interact with and ultimately use the output of the AI system – need to sufficiently ground their understanding of that output to some intermediate representation from their domain of expertise. A detailed analysis of a user study on this usecase is presented in Weisz et al. (2021).

Conclusion

In this article, we sought to briefly introduce and explore the area of *AI4Code* – i.e., the application of AI techniques to usecases revolving around source code. We showed that

[6] https://wala.github.io/graph4code/

[7] https://visualstudio.microsoft.com/services/intellicode/

[8] https://www.kite.com/

[9] http://ibm.biz/Mono2Micro

[10] https://thoth-station.ninja

[11] https://hai-gen2021.github.io/

in recent years, there has been an explosion of work in this area thanks to the naturalness hypothesis that draws a direct link between source code artefacts and the recent leaps in NLP that have occurred on human natural languages. We discussed the manifestation of some of these techniques in industrial and applied research scenarios, and provided a broad list of AI4Code deployments. This was followed by a deep dive into one specific effort centered around code translation; and an examination of one human in the loop issue around the deployment of this usecase. AI4Code remains a very exciting and fast-moving application area, and promises many breakthroughs in the days to come.

References

Agarwal, M., Barroso, J. J., Chakraborti, T., Dow, E. M., Fadnis, K., Godoy, B., & Talamadupula, K. (2020). CLAI: A Platform for AI Skills on the Command Line. *arXiv preprint arXiv:2002.00762*.

Agarwal, M., Talamadupula, K., Houde, S., Martinez, F., Muller, M., Richards, J., ... Weisz, J. D. (2020). Quality Estimation & Interpretability for Code Translation. *Computer Assisted Programming (CAP) Workshop at NeurIPS 2020*.

Allamanis, M., Barr, E. T., Devanbu, P., & Sutton, C. (2018). A survey of machine learning for big code and naturalness. *ACM Computing Surveys (CSUR)*, *51*(4), 1–37.

Devanbu, P. (2015). New initiative: The naturalness of software. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering* (Vol. 2, pp. 543–546).

Hindle, A., Barr, E. T., Gabel, M., Su, Z., & Devanbu, P. (2016). On the naturalness of software. *Communications of the ACM*, *59*(5), 122–131.

Hummer, W., Muthusamy, V., Rausch, T., Dube, P., El Maghraoui, K., Murthi, A., & Oum, P. (2019). Modelops: Cloud-based lifecycle management for reliable and trusted ai. In *2019 IEEE International Conference on Cloud Engineering (IC2E)* (pp. 113–120).

Nguyen, A. T., Nguyen, T. T., & Nguyen, T. N. (2014). Migrating code with statistical machine translation. In *Companion Proceedings of the 36th International Conference on Software Engineering* (pp. 544–547).

Oda, Y., Fudaba, H., Neubig, G., Hata, H., Sakti, S., Toda, T., & Nakamura, S. (2015). Learning to generate pseudo-code from source code using statistical machine translation (t). In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)* (pp. 574–584).

Roziere, B., Lachaux, M.-A., Chanussot, L., & Lample, G. (2020). Unsupervised translation of programming languages. *Advances in Neural Information Processing Systems*, *33*.

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *arXiv preprint arXiv:1409.3215*.

Wang, D., Ram, P., Weidele, D. K. I., Liu, S., Muller, M., Weisz, J. D., ... others (2020). Autoai: Automating the end-to-end ai lifecycle with humans-in-the-loop. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion* (pp. 77–78).

Weisz, J. D., Muller, M., Houde, S., Richards, J., Ross, S. I., Martinez, F., ... Talamadupula, K. (2021). Perfection Not Required? Human-AI Partnerships in Code Translation. In *26th Annual Conference on Intelligent User Interfaces (IUI).*

Xu, F. F., Vasilescu, B., & Neubig, G. (2021). In-IDE Code Generation from Natural Language: Promise and Challenges. *arXiv preprint arXiv:2101.11149*.

Kartik Talamadupula is a Research Scientist at IBM Research AI, and a AAAI Senior Member. He also serves as the SIGAI Applied AI Officer. His recent research interests are in applying AI to Business Automation, and in studying the implications of human users interacting with Generative AI models. Kartik can be reached at krtalamad@us.ibm.com.

Learning to See the Physical World

Jiajun Wu (Stanford University; jiajunwu@cs.stanford.edu)
DOI: 10.1145/3465074.3465081

Introduction

I am fascinated by how rich and flexible human intelligence is. From a quick glance at the scenes in Figure 1A, we effortlessly recognize the 3D geometry and texture of the objects within, reason about how they support each other, and when they move, track and predict their trajectories. Stacking blocks, picking up fruits—we also plan and interact with scenes and objects in many ways.

The goal of my dissertation research is to build machines that see, interact with, and reason about the physical world just like humans (Wu, 2020). This problem, **physical scene understanding**, involves three key topics that bridge research in computer science, AI, robotics, cognitive science, and neuroscience:

- **Perception** (Figure 1B): How can structured, physical object and scene representations arise from raw, multi-modal sensory input (e.g., videos, audios)?
- **Physical interactions** (Figure 1C): How can we build dynamics models that quickly adapt to complex, stochastic real-world scenarios, and how can they contribute to planning and motor control? Modeling physics helps robots to build bridges from a single image and to play games such as Jenga.
- **Reasoning** (Figure 1D): How can physical models integrate structured, often symbolic, priors such as symmetry and use them for commonsense reasoning?

Physical scene understanding is challenging, because it requires a holistic interpretation of scenes and objects, including their 3D geometry, physics, functionality, and modes of interactions, beyond the scope of a single discipline such as computer vision. Structured priors and representations of the physical world are essential: we need proper representations and learning paradigms to build data-efficient, flexible, and generalizable intelligent systems that understand physical scenes.

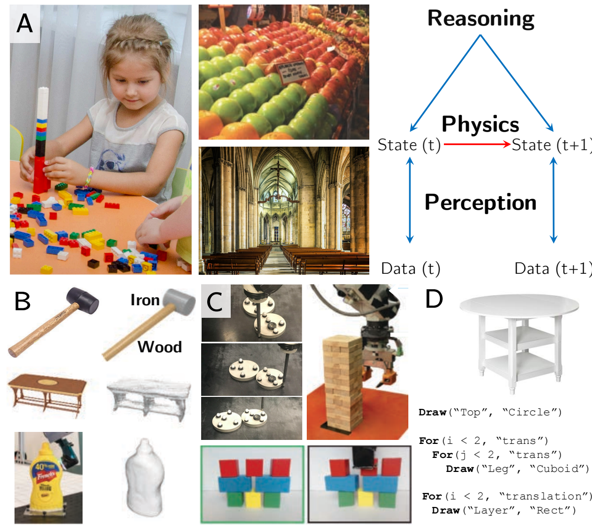

Figure 1: **Physical scene understanding** involves **(I) perception**, building physical representations from multi-modal data, **(II) physical interaction**, capturing scene dynamics for planning and control, and **(III) commonsense reasoning**, modeling high-level priors in scenes.

Our approach to constructing representations of the physical world is to integrate bottom-up **recognition models**, **deep networks**, and **efficient inference algorithms**, with top-down, structured **graphical models**, **simulation engines**, and **probabilistic programs**. In the dissertation, we develop and extend techniques in these areas (e.g., proposing new deep networks and physical simulators); we also explore ways to combine them, building upon studies across vision, learning, graphics, and robotics, with inspiration from cognitive science and neuroscience. Only by exploiting knowledge from all these areas and disciplines, may we build machines that have human-like, physical understanding of complex, real-world scenes.

Perception

Motivated by human perception—rich, generalizable, data-efficient—my research on perception focuses on building structured, object-based models to characterize the appearance

Figure 2: **Learning to see shapes, texture, and physics**. **A.** Reconstructing 3D shapes from a single color image via 2.5D sketches (Wu, Wang, et al., 2017). **B.** Generative modeling of 3D shapes and 2D images via a disentangled representation for object geometry, viewpoint, and texture (Zhu et al., 2018). **C.** 3D-aware representations for objects and scenes (Yao et al., 2018). **D.** Part-based object representations for its geometry and physics (Wu, Lu, et al., 2017).

and physics of objects.

My dissertation research covers various components of the appearance model. On bottom-up recognition, we have developed a general pipeline for 3D shape reconstruction from a single color image (Wu, Wang, et al., 2017) via modeling *intrinsic images*—depth, surface normals, and reflectance maps (Figure 2A). Our research is inspired by the classic research on multi-stage human visual perception (Marr, 1982), and has been extended to integrating learned priors of 3D shapes (i.e., 'what shapes look like?') for more realistic 3D reconstructions, and to tackling cases where the object in the image is not from the training categories.

Complementary to these bottom-up recognition models, we have also explored learning top-down graphics engines directly. We have proposed 3D generative adversarial networks, first applying generative-adversarial learning to 3D shapes for shape synthesis. We have later extended the model as visual object networks, which synthesize object shape and texture simultaneously, enforcing various consistencies with a distributed representation for object shape, 2.5D sketches, viewpoint, and texture (Figure 2B) (Zhu et al., 2018). We have generalized our models to scenes, recovering structured scene representations that not only capture object shape and texture, but enable 3D-aware scene manipulations (Figure 2C) (Yao et al., 2018).

Beyond object appearance, intuition of object physics assists humans in scene understanding. We have developed computational models that learn to infer object physics directly from visual observations. The Galileo model marries a physics engine with deep recognition nets to infer physical object properties (e.g., mass, friction). With an embedded physical simulator, the Galileo model discovers physical properties simply by watching objects move in unlabeled videos; it also predicts how they interact based on the inferred physical properties. The model was tested on a real-world video dataset, Physics 101, of 101 objects interacting in various physical events.

The dissertation also involves integrating geometry and physics perception (Figure 2D), with two primary results as "physical primitive decomposition" (PPD) and "visual de-animation" (VDA) (Wu, Lu, et al., 2017). In PPD, we decompose an object into parts with distinct geometry and physics, by learning to explain both the object's appearance and its behaviors in physical events; in VDA, our model learns to jointly infer physical world states and to simulate scene dynamics, integrating both a physics engine and a graphics engine. Our recent work has extended these models to complex indoor scenes, exploiting stability for more accurate 3D scene parsing.

Physical Interactions

My dissertation research also includes learning to approximate simulation engines (forward models) themselves. We have explored building physical models in various forms—image-based, object-based, and particle-based; analytical, neural, and hybrid—and have demonstrated their power in challenging, highly underactuated control tasks (Figure 3).

Compared with off-the-shelf simulators, a learned dynamics simulator flexibly adapts to novel environments and captures the stochasticity in scene dynamics. Our visual dynamics model demonstrates this in the pixel domain, where it learns to synthesize multiple possible future frames from a single color image by automatically discovering independent movable parts and their motion distributions (Xue et al., 2016) (Figure 3A). We have later extended the

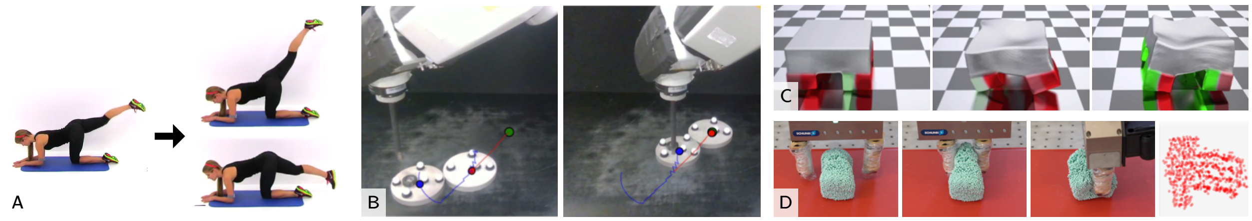

Figure 3: **Physical models for future prediction and control**. **A.** Modeling visual dynamics allows us to generate multiple possible future frames from a single image (Xue et al., 2016). **B.** We have developed a hybrid model that captures object-based dynamics by integrating analytical models and neural nets. It assists the robot in accomplishing a highly underactuated task: pushing the right disk to the target (green) by only interacting with the left disk (Ajay et al., 2019). **C. D.** Particle-based dynamics models support controlling soft robots (Hu et al., 2019) and manipulating deformable objects and liquids (Li et al., 2019).

model to additionally capture the hierarchical structure among object parts.

Modeling dynamics directly in the pixel space is universal but challenging due to the intricate interplay between physics and graphics; an alternative is to separate perception from dynamics modeling, and learn dynamics from object states. Our recent work along this line has shown that a model that learns to approximate object dynamics can be useful for planning, generalize to scenarios where only partial observations are available, and discover physical object properties without supervision. We have further extended our model to particle-based representations, so that it can characterize the dynamics of soft robots (Hu et al., 2019) (Figure 3C) and of scenes with complex interactions among rigid bodies, deformable shapes, and fluids (Li et al., 2019) (Figure 3D).

We have also explored the idea of learning a hybrid dynamics model, augmenting analytical physics engines with neural dynamics models (Ajay et al., 2019) (Figure 3B). Such a hybrid system achieves the best of both worlds: it performs better, captures uncertainty in data, learns efficiently from limited annotations, and generalizes to novel shapes and materials. These dynamics models can be used in various control tasks: they help to solve highly underactuated control problems (pushing disk A, which in turn pushes disk B to the target position), to control and co-design soft robots (Hu et al., 2019), to manipulate fluids and rigid bodies on a Kuka robot (Li et al., 2019), and to interact and play games such as Jenga that involve complex frictional micro-interactions.

Reasoning

The physical world is rich but structured: natural objects and scenes are compositional (scene are made of objects which, in turn, are made of parts); they often have program-like structure (e.g., symmetry). We have also been exploring ways to bridge structured, often symbolic, priors into powerful deep recognition models.

A test of these neuro-symbolic representations is how well they support solving various reasoning tasks such as analogy making and question answering. Our recent work demonstrated that, when combined with deep visual perception modules, a symbolic reasoning system achieves impressive performance on visual reasoning benchmarks, outperforming end-to-end trained neural models. We have also extended it to jointly learn visual concepts (e.g., colors, shapes) and their correspondence with words from natural supervision (question-answer pairs) via curriculum learning, without human annotations (Mao, Gan, Kohli, Tenenbaum, & Wu, 2019).

Next Steps

With big data, large computing resources, and advanced learning algorithms, the once separated areas across computer science (vision, learning, symbolic reasoning, NLP, rule learning and program induction, planning, and control) has begun to reintegrate. We should now take an more integrative view towards these areas and actively explore their interactions for a more general AI landscape.

One such direction is to achieve more fundamental integration of perception, reasoning, and planning. While most computational models have treated them as disjoint modules, we observe that having them communicate with each other facilitates model design and leads to better performance. Another direction is to integrate symbolic priors with deep representation learning via program synthesis for concept and structure discovery. Neuro-symbolic methods enjoy both the recognition power from neural nets and the combinatorial generalization from symbolic structure; therefore, they have great potential in scaling up current intelligent systems to large-scale, complex physical scenes in the real life, for which pure bottom-up, data-driven models cannot work well due to the exponentially increasing complexity. Beyond physical objects and scenes, I also want to build computational models that understand an agent's goals, beliefs, intentions, and theory of mind, and use these knowledge for planning and problem solving, drawing inspiration from intuitive psychology.

References

Ajay, A., Bauza, M., Wu, J., Fazeli, N., Tenenbaum, J. B., Rodriguez, A., & Kaelbling, L. P. (2019). Combining physical simulators and object-based networks for control. In *IEEE International Conference on Robotics and Automation (ICRA).*

Hu, Y., Liu, J., Spielberg, A., Tenenbaum, J. B., Freeman, W. T., Wu, J., ... Matusik, W. (2019). Chainqueen: A real-time differentiable physical simulator for soft robotics. In *IEEE International Conference on Robotics and Automation (ICRA).*

Li, Y., Wu, J., Tedrake, R., Tenenbaum, J. B., & Torralba, A. (2019). Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *International Conference on Learning Representations (ICLR).*

Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., & Wu, J. (2019). The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations (ICLR).*

Marr, D. (1982). *Vision*. W. H. Freeman and Company.

Wu, J. (2020). *Learning to see the physical world* (Doctoral dissertation). Massachusetts Institute of Technology.

Wu, J., Lu, E., Kohli, P., Freeman, W. T., & Tenenbaum, J. B. (2017). Learning to see physics via visual de-animation. In *Advances in Neural Information Processing Systems (NeurIPS).*

Wu, J., Wang, Y., Xue, T., Sun, X., Freeman, W. T., & Tenenbaum, J. B. (2017). MarrNet: 3D Shape Reconstruction via 2.5D Sketches. In *Advances in Neural Information Processing Systems (NeurIPS).*

Xue, T., Wu, J., Bouman, K. L., & Freeman, W. T. (2016). Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *Advances in Neural Information Processing Systems (NeurIPS).*

Yao, S., Hsu, T.-M. H., Zhu, J.-Y., Wu, J., Torralba, A., Freeman, W. T., & Tenenbaum, J. B. (2018). 3D-aware scene manipulation via inverse graphics. In *Advances in Neural Information Processing Systems (NeurIPS).*

Zhu, J.-Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J. B., & Freeman, W. T. (2018). Visual object networks: Image generation with disentangled 3D representations. In *Advances in Neural Information Processing Systems (NeurIPS).*

Jiajun Wu is an Assistant Professor of Computer Science at Stanford University, working on computer vision, machine learning, and computational cognitive science. Before joining Stanford, he was a Visiting Faculty Researcher at Google Research. He received his PhD in Electrical Engineering and Computer Science at Massachusetts Institute of Technology. Wu's research has been recognized through the AAAI/ACM SIGAI Doctoral Dissertation Award, the ACM Doctoral Dissertation Award Honorable Mention, the MIT George M. Sprowls PhD Thesis Award in Artificial Intelligence and Decision-Making, the 2020 Samsung AI Researcher of the Year, the IROS Best Paper Award on Cognitive Robotics, and fellowships from Facebook, Nvidia, Samsung, and Adobe.

Novel Practices and Highlights from the Fourth Workshop on Mechanism Design for Social Good

Francisco J. Marmolejo Cossio (University of Oxford; francisco.marmolejo@cs.ox.ac.uk)
Faidra Monachou (Stanford University; monachou@stanford.edu)
DOI: 10.1145/3465074.3465082

Abstract

The *Fourth Workshop on Mechanism Design for Social Good* was held virtually in August 2020, with a focus on work bridging research and policy. This article represents the experience of the chairs and discusses novel conference-organizing practices aimed at promoting *multi-disciplinary research for social good* and increasing *racial, linguistic*, and *geographic diversity and inclusion*.

Introduction

The *Fourth Workshop on Mechanism Design for Social Good* (MD4SG'20) was held online on August 16-19, 2020. The authors of this article co-chaired and organized the workshop.

The goal of MD4SG'20 was to highlight research where techniques from algorithms, optimization, and mechanism design, along with insights from other disciplines, have the potential to improve access to opportunity for historically underserved and marginalized communities. The workshop featured five keynote presentations, forty contributed talks including problem pitches and demos, two poster sessions, a panel discussion, and networking events, with a focus on *bridging research and policy*. To this end, participants included researchers as well as practitioners in various government and non-government organizations and industry. Due to the COVID-19 pandemic, the workshop was fully virtual, taking place on the online platforms Zoom and Gather.

The current workshop was the fourth annual workshop (since 2017) in a series of workshops on *Mechanism Design for Social Good* (MD4SG). This was the first independent MD4SG workshop, as the three previous iterations of the workshop were organized alongside the annual ACM Conference on Economics and Computation. All four workshops were organized as a part of the larger MD4SG initiative, which is a multidisciplinary and multi-institutional online research initiative that promotes research at the intersection of computer science, operations, economics, humanities and other disciplines, with the mission of bringing together a range of expertise to tackle problems impacting disadvantaged communities around the world. Since its foundation in 2016, MD4SG has grown to a community of more than 2,000 participants and organizes workshops, tutorials, colloquium series and working groups covering topics such as developing nations, discrimination and equity in algorithmic decision-making systems, environment and climate, inequality, civic participation, as well as newly-formed regional groups (Asia-Pacific, Latin America and the Caribbean).

Workshop Objectives and Main Contributions

Our goal, through the main program of the workshop and novel practices we adopted, was to organize a workshop that is *multidisciplinary*, *diverse* and *thought-provoking*. More specifically, the workshop focused on and achieved three objectives:

– **Multi-disciplinary research for social good.** Following the tradition of the MD4SG initiative and the previous workshops, MD4SG'20 brought together researchers, policymakers and other domain experts and professionals interested in improving equity and developing solutions for problems in a variety of application domains such as education, labor, environment, healthcare, algorithmic fairness, and digital platforms. Due to its interdisciplinary nature, MD4SG'20 attracted a very diverse and large group of members with backgrounds in *computer science, AI, operations research, economics, public policy* and *humanities*, while a great number of papers combined methodologies and insights from multiple fields.

  Beyond science, policy, and humanities, the workshop also explored the intersection between *art* and mechanism design for social good. Our keynote speaker Stephanie Dinkins, Artist Fellow at the Berggruen Institute, talked about her practice as a transmedia artist working with AI, which "employs lens-based practices, emerging technologies and community engagement to confront questions of bias in AI, consciousness, data sovereignty and social equity".

– **Bridging research and policy.** Our workshop emphasized the application- and policy-oriented character of mechanism design for social good by including four different tracks which ranged from technical content (AI/ML, Theory, Empirical Studies and Policy) to more practical applications (Problems and Demonstrations).

  As a novel highlight from our main program, we encouraged, reviewed, and accepted law and policy papers. Joint research at the intersection of law, policy, economics and computation is underdeveloped. Our session on *Technology, Law, and Policy* included work from law experts on topics such as privacy, security, and gender equality.

  Several of our award-winning papers focused on *policy design*, by presenting policy-oriented research on topics such as feminicide in Latin America, educational policies for admissions at University of California or school choice in Peru and San Francisco, HIV prevention methods for homeless youth, and discrimination in labor markets.

– **Diversity and inclusion: race, language and geography.** One of the main goals of the workshop was to reach out to a broader audience and be inclusive of underserved communities in academia. Such inclusion, especially with regard to *gender, racial and geographic diversity and inclusion*, has always been a core value and strength of MD4SG.

  This year's workshop achieved a record number of *submissions (138)*, attracting more than double the number of submissions in 2019, and more than *700 registrations* from *59 countries* around the world.

  The workshop successfully reached out to communities and institutions in *Africa* and *Latin America and Caribbean* (LAC). More specifically, 20.9% of our registrants identified as Black, African American, or African, while 18% were of Hispanic, Latino, or Spanish origin. Out of 138 submissions in total, we received 7 and 22 submissions from authors based in Africa and LAC respectively.

  One of the novel practices of the workshop was *linguistic diversity*. A significant barrier for the participation of Latin American communities in similar initiatives has been the language. We released the call for participation both in *English* and *Spanish*, and our Spanish-speaking Program Committee members reviewed and shepherded 6 submissions written entirely in Spanish. In addition, our plenary speaker Natalia Ariza Ramírez (Economist at National University of Colombia and former Vice Minister of Education in Colombia) gave her plenary talk in Spanish, and conducted a discussion panel with experts from LAC entirely in Spanish. We provided real-time interpretation from Spanish to English for the two events in Spanish, and from English to Spanish for the remaining plenary talks.

  For many of our participants, including students and attendees from the Global South, finances were a barrier to attending this workshop. To assist such participants, we provided financial assistance in the form of *registration fee waivers* to 190 participants and 21 *data plan scholarships* to participants without Internet access. 18 of the data plan scholarships were awarded to participants located in Africa.

  We hope that our outreach to Latin America and Africa will set a positive example for other conferences and research initiatives like MD4SG, and have a long-term impact on increasing the representation and the participation of Latin America in computer science, operations, and related fields.

Novel Conference Practices: Outcomes and Broad Impact

Under the current backdrop of the global COVID-19 pandemic, most academic events such as MD4SG'20 have been forced to migrate to a virtual setting. Although the specific details of effective implementation – from an infrastructural and logistical perspective – are important in their own right, we would like to highlight the most salient aspect of our experience: the very virtual nature of the MD4SG initiative (which has held year-round virtual events since its inception in 2016) along with targeted outreach massively increased the degree of *diversity and inclusion* we could foster. We intend to host virtual events well beyond the pandemic to continue to engage with academics, local stakeholders and relevant policymakers who may otherwise be unable or less willing to attend in-person events.

MD4SG has traditionally had strong representation from international communities (in particular from within the African continent), and for many such members, attending our workshop in the past has been difficult due to: the large financial burden induced by transportation and registration costs; the logistical difficulties in travelling long distances to workshops; and/or visa difficulties when preparing for travel to events that tend to occur in the United States. Given the virtual nature of MD4SG'20, these issues were either minimized or eliminated completely, paving the way for increased participation from the communities that make MD4SG unique as an organization.

Financial Assistance: Registration Fee Waivers and Data Plan Scholarships

An important lesson from our experience was that registration infrastructure and financial assistance can amplify the increased participation achieved in virtual events. Although we set a 10 USD registration fee, this was mainly to prevent possible spamming attacks that occur at similar virtual events. We provided financial assistance in the form of fee waivers, and data scholarships which provided individuals with 20 USD in the form of local mobile data to participate in the workshop. We received 190 registration fee waiver requests and 87 data plan scholarship applications. The most common reasons for applying for some form of financial aid were:

(1) loss of income due to the COVID-19 pandemic, and

(2) being a student in a developing country and/or without personal income.

We granted all *190 registration fee waivers*, while we also purchased and provided *21 data plan scholarships* (of 20 USD each, in the form of local mobile data) to participants without Internet access; 18 of the data scholarships were awarded to individuals located in Africa. More information about the geographic region of participants who requested financial assistance is provided in Figure 1; our collected data showed that almost half of the financial assistance applications came from the African continent.

Figure 1: *Financial assistance*: Information about the type of financial assistance requested (left) and the geographic region of the applicants (right).

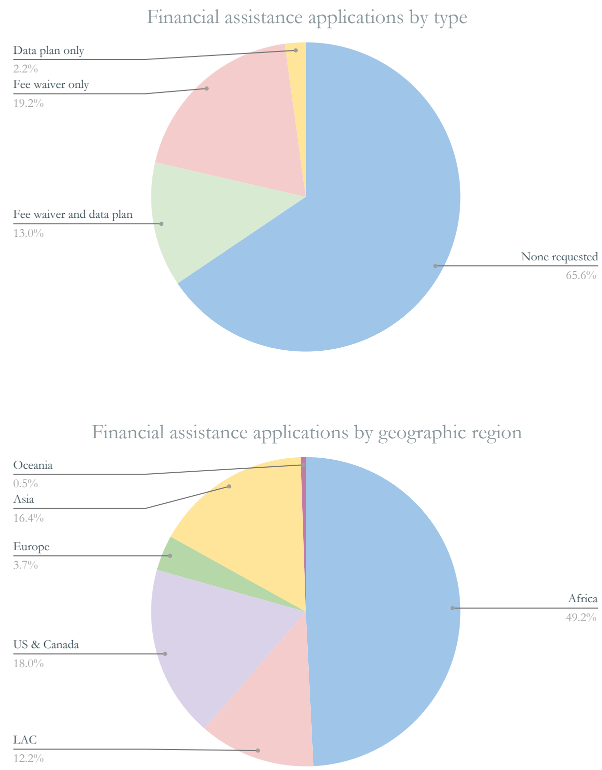

Furthermore, a key innovation within MD4SG'20 registration was the functionality for paying participants to donate towards registration waivers of others in need of assistance. We were pleasantly surprised by the generosity of the MD4SG community, as we obtained 923 USD in *donations*, many of which came in small amounts from other individuals within the very international communities we were striving to include.

Cultural and Linguistic Inclusion

One of our key goals for MD4SG'20 was to foster participation in the workshop and the research organization as a whole from within the *Latin American and Caribbean* (LAC) community. Early on we noticed that a language divide often creates a barrier for participation for individuals from LAC. Thus, in order to engage the LAC community and bridge this divide, some key changes to the agenda as well as the promotion of the workshop helped us expand its inclusiveness for presenters as well as attendees.

For example, prior to the event, we translated our call for papers to *Spanish* to enable Spanish-speaking participants to share their submissions with us. We also ensured that all our marketing materials for the event were bilingual to encourage further awareness and participation from the community. Not only did this result in us receiving 22 submissions from authors affiliated to institutions within LAC regions, but we also received 6 papers entirely in Spanish—a first for a major technical workshop like MD4SG'20. Many of these submissions addressed key issues in education and policy in LAC, opening up a new spectrum of perspectives to all the workshop participants. Correspondingly, we arranged for the submissions to be reviewed by Spanish speaking members of our Program Committee to ensure that they were impartially and accurately reviewed.

Several of the Spanish submissions were shaped into poster presentations via a *shepherding process* involving mentorship from the Program Committee. Each poster had an associated lightning talk, the material of which was presented in both English and Spanish. The sessions were recorded and hosted on YouTube to increase visibility.

From a social perspective, we tried to enhance the participation experience of our Spanish speaking attendees by providing social spaces for Spanish speakers within the Gather platform.

Finally, and perhaps most importantly, we engaged an official *real-time translation and interpretation service* to ensure that all our plenary sessions and panel discussions were live-translated from English into Spanish (and vice versa). Feedback from participants was highly positive, and indicates that this sort of initiative is not often observed in events within our community.

Once more, we note that these outreach efforts alongside the virtual nature of MD4SG'20 allowed us to increase participation from the LAC community to an unprecedented degree. We hope that other established academic events adopt similar practices to decrease linguistic barriers. Such practices have the long-term potential to drastically increase the inclusion of underserved communities in STEM.

Racial, geographic, and gender diversity

The workshop received *more than 700 registrations*, out of which 650 registrants responded to our pre-registration survey. Based on the survey responses, we collected several useful statistics.

MD4SG'20 was a truly global workshop. In terms of geography, our participants came from *59 countries* around the world (see Figure 2), with most of them coming from the following countries, in order of participation: *United States*, *Ethiopia*, *Mexico*, *India*, *Nigeria*, *Canada*, *United Kingdom*, and *Tanzania*. Thus, despite the time constraints of any online workshop, MD4SG'20 managed to have a very good level of representation around the world, especially from Africa and Latin America. Figure 3 includes the relevant statistics for each geographic region.

Figure 2: *Registration map*: The countries (current location) of registrants are highlighted in blue.

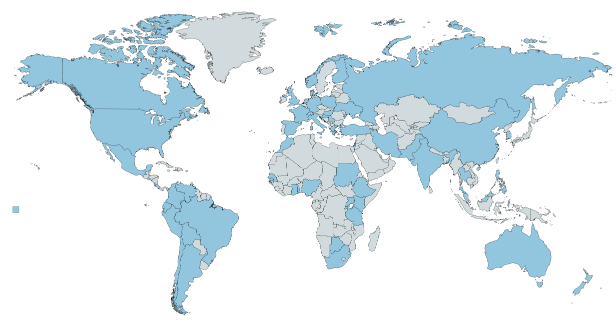

Figure 3: *Geographic representation*: Percentage of registrants based on their reported country of current location, grouped by geographic location. N/A denotes that registrants did not report their current location.

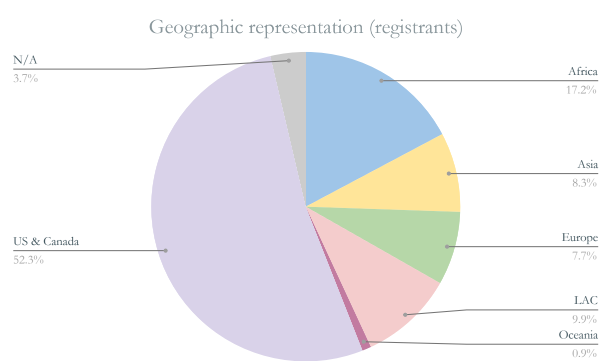

From the responses to our pre-registration survey, we also collected some useful demographic information about the workshop registrants. Regarding the *diversity and inclusion* outcomes of the workshop, some statistics are particularly encouraging as Figure 5 illustrates. For example, in terms of racial and ethnic diversity, we had *20.9%* registrants who identified as *Black, African American, or African*, while *18%* of registrants were of *Hispanic, Latino, or Spanish origin*. With respect to gender diversity, *37.5%* of our registrants identified as *female*. Furthermore, the registrants consisted mostly of younger individuals, with 74% being less than 34 years old, and 38.2% of our registrants were graduate students while 7.5% were undergraduate students.

Academic Inclusion of Other Disciplines

One of the key pillars of the MD4SG initiative is fostering participation from a variety of disciplines to ensure well-rounded perspectives on the key issues that our community works on.

A big step towards this in our workshop was the inclusion of a number of participants from different backgrounds—most notably, *law*. To this end, we held a session with four contributed talks on the intersection of Law and Computer Science which was well-received by participants. While there was scope to further align the content of these talks to the backgrounds of the vast majority of the participants, we believe this session was a step in the right direction towards bringing these two fields together.

A great number of papers also focused on policy interventions informed by empirical methods, as well as the empirical validation of policy interventions. The Problems and Demonstrations track targeted papers from government and policy, as well as non-government organizations and industry, and included white papers documenting open problems or demonstrating prototyped and/or deployed software systems and mobile platforms.

We were particularly encouraged by the wide range of topics studied by the papers we received. As Figure 4 shows in greater detail, the most common research area was fairness in algorithmic design and resource allocation settings, followed by papers in the intersection of law and policy.

In the future, we would like to further work alongside our participants from other disciplines to ensure that their work translates more effectively into the language of our community.

Figure 4: *Research topics at MD4SG'20*: Research topics studied by accepted papers (posters, talks) at the workshop. A subset of papers focus on more than one research area.

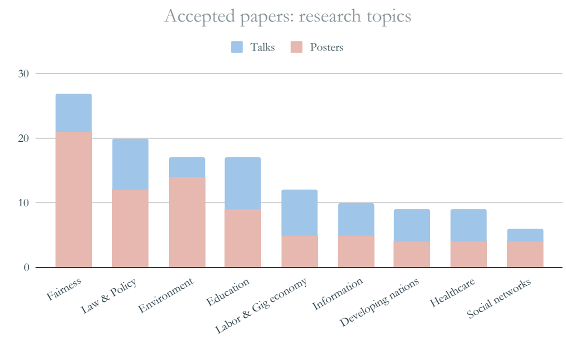

Partnerships and Collaborations Developed

The workshop helped MD4SG solidify our existing partnerships with relevant organizations such as: Schmidt Futures and SIGAI who so kindly helped fund our event. Furthermore, the engagement we received from individuals from LAC has persisted in various ways. First of all, many working groups have increased their membership to include individuals from LAC countries who were previously unaware of MD4SG research. This is especially true of individuals from academic institutions which were heavily represented in the submission process, such as the University of Chile (who's authors led 4 different submissions). In addition, the success of the LAC outreach from MD4SG'20 has created sufficient momentum for existing MD4SG members to prepare a bilingual work agenda around regional issues for a new LAC working group.

Overall Lessons

The current pandemic is a situation completely unforeseen by all of us within the community. In the midst of these difficult times, we have also learned how valuable the social angle of a workshop like MD4SG'20 can be. A common remark on our user survey is that participants enjoyed the social nature of the Gather platform driven by the social events prepared by our workshop chairs. MD4SG is currently in the process of providing a similar social space on a regular basis for MD4SG members throughout the year.

Figure 5: *Diversity and Inclusion*: Demographic information about the MD4SG'20 registrants.

(a) Race

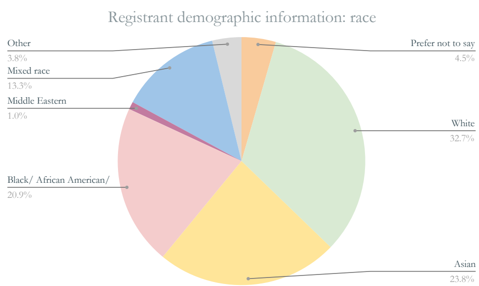

(b) Hispanic/Latino origin

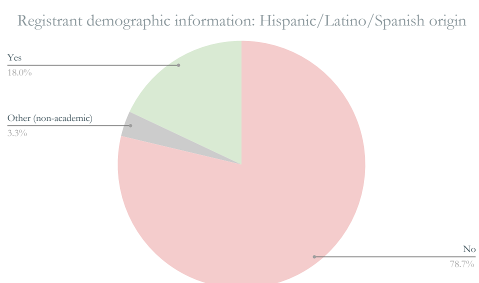

(c) Gender

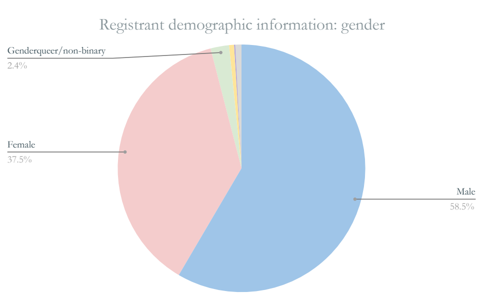

(d) Age

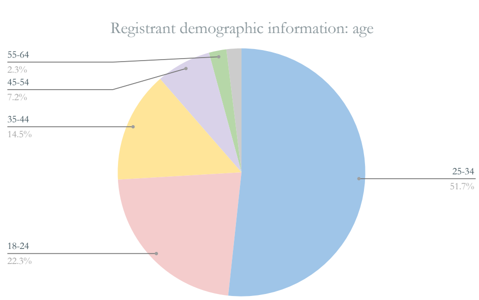

Program Highlights

Together with the whole MD4SG community, the participation at our annual workshop has also grown year by year. The first MD4SG workshop took place in 2017 and had 20 submissions. The workshop continued to grow with 30 submissions in 2018 and 65 submissions in 2019. Our virtual workshop MD4SG'20 received *more than double* the number of submissions in 2019, reaching the record number of *138 submissions* and exceeding our initial expectations.

As already mentioned, the main program of the workshop (on August 17-19) ran for three days and included 5 *keynote talks*, 40 *contributed long and short oral presentations*, a *panel discussion*, and two *poster sessions*. We also organized 4 *networking events*, starting with an informal reception on August 16.

Keynote talks

The workshop hosted 5 invited presentations which focused on topics connecting research for social good and policy design: markets without money, civil liberties and extremism, technology in support of care-giving, educational policies in Colombia, and the dialogue between art and artificial intelligence. Each talk was followed by a short presentation and an open discussion with the participants led by invited discussants. The talks were as follows:

"Research and Policy Challenges in Implementing Colombia's *Ser Pilo Paga* Program," by *Natalia Ariza Ramírez* (Economist at National University of Colombia and former Vice Minister of Education in Colombia);

"Community, Craft, and the Vernacular in Artificial Intelligence," by *Stephanie Dinkins* (Artist Fellow at the Berggruen Institute);

"Tech in Support of Caregiving: Innovation Opportunities and Ecosystem Challenges," by *Deborah Estrin* (Associate Dean for Impact at Cornell Tech);

"Maximizing the Social Good: Markets without Money," by *Nicole Immorlica* (Senior Researcher at Microsoft Research);

"How to Fight White Supremacist Extremism While Protecting Civil Liberties: A Multidisciplinary Approach Using Technology, Research, and Policy," by *Anjana Rajan* (Chief Technology Officer at Polaris).

Discussion panel

On the second day of the main program, the workshop highlighted topics from Latin America. To that end, we also hosted a discussion panel in Spanish (with live interpretation to English) which followed after the keynote talk by Natalia Ariza Ramírez and the session on Education in Practice. The focus of our panel was the intersection of policy and academia within the scope of education in Latin America. Our panelists were our keynote speaker *Natalia Ariza Ramírez*, *José R. Correa* (Professor at Universidad de Chile), and *Rafael Obregón* (UNICEF Paraguay). The engaging discussion introduced the MD4SG community to new policy-oriented problems and relevant experts, and helped the participants understand the unique challenges that policy makers and researchers face with respect to the educational system in Latin American countries, in comparison to the rest of the world and the United States in particular.

Contributed talks

The technical program included *12 long talks*, *28 short talks*, and *75 poster presentations*. The 10 sessions for contributed talks represented the wide range of topics and application domains of interest to the MD4SG community, the combination of novel and diverse methodologies as well as the strong connections of many papers to policy design. In the spirit of the workshop's theme of bridging research and policy, and to encourage the academic exchange of ideas between law and mechanism design for social good, we included a session dedicated to non-technical papers on the connections among law, technology and policy.

More specifically, the sessions were as follows: (1) *Education Policy and Diversity*, (2) *Technology, Law and Policy*, (3) *Labor Markets*, (4) *Environment, Agriculture and Food Consumption*, (5) *Education in Practice*, (6) *Healthcare*, (7) *Fairness and Inequality*, (8) *Algorithms for Policy and Governance*, (9) *Online Platforms and Civic Participation*, and, finally, (10) *Information*.

Awards

The workshop included three categories of technical awards for exemplary work (*paper awards*, *poster awards*, and *participant awards*), while *Most Popular Poster Awards* recognized the most successful poster presentations based on participants' votes. Three *Best Tweet Awards* were given to the participants with the most engaging tweets or the most active social media coverage of the workshop.

Related to the workshop theme of bridging research and policy, the award-winning papers highlighted new research directions for policy-oriented work in the MD4SG community. They spanned various critical application domains such as education, labor, healthcare, and criminal justice in the United States and worldwide, successfully combining different methodologies in a novel manner. The winning papers that equally shared the **Best Paper Award** were:

- **"Top Percent Policies and the Return to Postsecondary Selectivity"**, by Zachary Bleemer, and
- **"Competition under Social Interactions and the Design of Education Policies"**, by Claudia Allende.

Both papers focused on the design of effective educational policies and demonstrated exceptional, policy-driven research that can have a positive impact on the lives of thousands of students from less privileged backgrounds. The former paper used novel data from a "Top Percent" admissions policy implemented by the University of California to analyze the impact on barely-eligible applicants on their university admission and future career outcomes, while the latter studied the role of peer preferences in school choice and the design of optimal assignment policies using data from elementary schools in Peru.

The **Best Student Paper Award** was awarded to two papers with student leading authors:

- **"Large-scale Clinical Trial of an AI-augmented Intervention for HIV Prevention in Youth Experiencing Homelessness,"** by Bryan Wilder, Laura Onasch-Vera, Graham Diguiseppi, Robin Petering, Chyna Hill, Amulya Yadav, Eric Rice and Milind Tambe, and
- **"All Things Equal? Social Networks as a Mechanism for Discrimination,"** by Chika Okafor.

The paper by Wilder et al. focused on the issue of HIV prevalence in homeless youth and the design and successful implementation of a related clinical trial, with a particular highlight on community engagement and informational bottlenecks. The other award-winning paper by Okafor developed a theoretical labor market model with referrals and showed that the combination of homophily and different group size can lead to disparities across different social groups.

This year, we also introduced the **New Horizons Award**, to highlight promising, ongoing work in an emerging area of research. Two working papers shared this award:

– **"Feminicide and Machine Learning: Detecting Gender-based Violence to Strengthen Civil Sector Activism,"** by Catherine D'Ignazio, Helena Suarez Val, Silvana Fumega, Harini Suresh, Isadora Cruxen, Wonyoung So, Maria De Los Angeles Martinez and Mariel Garcia-Montes, and

– **"Modeling Assumptions Clash with the Real World: Configuring Student Assignment Algorithms to Serve Community Needs,"** by Samantha Robertson, Tonya Nguyen, and Niloufar Salehi.

The former paper highlighted a highly relevant topic to the Latin American region by adopting machine learning tools to understand the complex issue of feminicide and inform policy at all levels. The latter paper studied how the theoretical guarantees of the San Francisco Unified School District's student assignment algorithm can differ from the practical behaviour of parents using the algorithm.

Finally, the workshop featured awards for papers accepted for poster presentation. Based on participants' votes, three papers (one written and presented in Spanish) shared the **Most Popular Poster Award**:

– **"Laboratorio de Derecho y Política Local, propone: Red de Monitoras y Monitores Derecho para Todos,"** by Lorayne Finol Romero, Cecilia González Jeria and Maximiliano Núñez Gómez,

– **"Guaranteeing Maximin Shares: Some Agents Left Behind,"** by Hadi Hosseini, Andrew Searns and Sawyer Welden, and

– **"A Comparison of Living Standards Across the States of America,"** by Vegard Nygaard and Elena Falcettoni.

Acknowledgments

It was a wonderful learning experience to preside over MD4SG'20 as co-chairs and work together with so many people to organize this workshop. We are indebted to our *Steering Committee* members Rediet Abebe, Kira Goldner, Jon Kleinberg, Illenin Kondo, Sera Linardi, Irene Lo, and Ana-Andreea Stoica, for their guidance and support in organizing the workshop. We are also very grateful to our *Area Chairs* Zoe B. Cullen, Daniel Freund, Abhishek Gupta, Dina Machuve, Robert Manduca, Araba Sey, Ana-Andreea Stoica, Sam Taggart, Matt Weinberg, Bryan Wilder, and—of course—all the 74 members of our *Program Committee* for their invaluable help with the reviewing process, as well as our *Organization Chairs* Jessie Finocchiaro, Michelle Gonzalez, Meareg Hailemariam, Wanyi Li, Duncan McElfresh, Amita Shukla, Logan Stapleton, and Lily Xu, for their hard work in planning and organizing the workshop. We would also like to thank Jenny Lam and the department of Management Science and Engineering at Stanford University for their assistance throughout the registration process and financial planning of the workshop, *Virtual Chair* for their assistance in event planning, platform design and organization, and Pat Gutierrez for her wonderful translation work. Finally, we are extremely grateful to the sponsors of our workshop, *ACM SIGAI Activities Fund* and *Schmidt Futures*, for their generous financial support.

Francisco J. Marmolejo Cossío is a Fellow in Computer Science at Balliol College (University of Oxford) and a Research Fellow at IOHK. His research lies at the intersection of Algorithmic Game Theory, Decentralised Consensus Protocols and Computational Learning Theory. Francisco is originally from Mexico, and has been co-organizing the MD4SG research group since January 2020.

Faidra Monachou is a Ph.D. candidate in Management Science and Engineering at Stanford University. She is interested in market design and the study of discrimination, diversity, and information in socioeconomic problems that arise in matching markets. Faidra received her undergraduate degree in Electrical and Computer Engineering from National Technical University of Athens in Greece.

Decoding AI – Summary of an ACM SIGAI Sponsored Event

Vandana Srivastava (ViSER; viserllc@gmail.com)
DOI: 10.1145/3465074.3465083

Event Details

The "Decoding AI" event was conducted online by ViSER for high school students and adults between Oct 10 to Oct 31, 2020 and was sponsored by ACM SIGAI . The proposal for the event was selected by ACM SIGAI from a pool of entries. It was advertised on different social media platforms and registration was available using Google forms and Eventbrite.

A total of 67 people registered for the event comprising of adults (26), high school students (22), college students (5), and others who did not declare their background (14). The registrants belonged to different states in the USA (NY, SC, TX, CA, NE) and from different countries (USA, India, England, Germany, Kenya, Nepal).

The event highlights are as follows:

- The event comprised of 4 sessions (on Saturdays), each session being 2 hours long and centered around different topics related to AI.
- The following topics were covered in four weeks: AI and Machine Learning, Data Analysis, Text Analysis and Creating Chatbots using Dialogflow. Each session included an introduction of the topic, accompanied by the hands-on project.
- The participants were introduced the basics of Python programming in the first session of the event. In the second session, we trained them to use Jupyter notebook online (https://jupyter.org/try – binder) so that they can follow the programming projects.
- In session 1, the participants took part in a short quiz. In sessions 2 and 3, two high school students, who did their internship with ViSER, presented their Data Analysis and Text Analysis project. In session 4, an expert in AI gave a 10 min talk on AI and Ethics.

Study Material For The Event

The video recording of all the sessions, Jupyter notebook with programs, and data files were shared with the participants. These files were shared either during or after each session on github. The recordings are available at:

- Session 1:Intro to AI and Machine Learning (video - part 1) and (video - part 2)
- Session 2: Python for Data Analysis (video)
- Session 3: Machine Learning for Text Analysis (video)
- Session 4: How to Create Chatbots using Dialogflow? (video)

The event feedback was very positive and encouraging. All the respondents showed interest in having more events like this in the future. ViSER would like to thank ACM-SIGAI for sponsoring this event.

Vandana Srivastava is Founder and CEO of ViSER LLC, a services company for individuals - providing learning resources, and businesses - providing data driven solutions. She has more than 10 years of experience in both industry and academia in USA and India. Vandana has MBA in Financial Management from Pace University, NY and MS in Computational Mathematics from Arizona State University, AZ. She has worked as Vice President (Tantiv4), Incentive Analyst (IBM), in USA and as Assistant Professor of Computer Science, Mathematics and Finance in different Colleges in India, including Indian Institute of Technology (IIT), Delhi.