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:

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

Editor Emerita

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(2)

Iolanda Leite, co-editor (Royal Institute of Technology (KTH); aimatters@sigai.acm.org)
Anuj Karpatne, co-editor (Virginia Tech; aimatters@sigai.acm.org)
DOI: 10.1145/3478369.3478370

Issue overview

Welcome to the second 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 officer. In our regular Education column, Carolyn Rosé discusses the role of AI in education in a post-pandemic reality. We then bring you our regular Policy column, where Larry Medsker covers interesting and timely discussions on AI policy, for example whether governments should play a role in reducing algorithmic bias. This issue closes with an article contribution from Li Dong, one of the runner-ups in the latest AAIS/SIGAI dissertation award, on the use neural models to build natural language interfaces.

The editorial team wishes a great summer to all SIGAI community members!

Special Issue: AI For Social Good

Recognizing the potential of AI in solving some of the most pressing challenges facing our society, we are excited to announce that the next Newsletter of AI Matters will be a special issue on the theme of "AI for Social Good." We solicit articles that discuss how AI applications and/or innovations have resulted in a meaningful impact on a societally relevant problem, including problems in the domains of health, agriculture, environmental sustainability, ecological forecasting, urban planning, climate science, education, social welfare and justice, ethics and privacy, and assistive technology for people with disabilities. We also encourage submissions on emerging problems where AI advances have the potential to influence a transformative change, and perspective articles that highlight the challenges faced by current standards of AI to have a societal impact and opportunities for future research in this area. More details to be coming soon on http://sigai.acm.org/aimatters. Please get in touch with us if you have any questions!

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: 10.1145/3478369.3478371

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.

The 5th International Conference on Information and Knowledge Systems (ICIKS'21)

Online Conference, June 22-23, 2021
https://iciks.org/
ICIKS 2021 aims to gather both researchers and practitioners in the fields of Information Systems, Artificial Intelligence, Knowledge Management and Decision Support in promoting discussions on various organizational, technological, and socio-cultural aspects of research in the design and use of information and knowledge systems in organizations. The conference expects to feature research related to the design, implementation, and use of information and knowledge systems that take into consideration the cultural context of the organizations, and the viewpoints of their shareholders. Researchers and practitioners are invited to share recent advances in methodologies, models and tools concerning decision processes and information and knowledge systems that contribute to the success of digital transformation in organizations. The conference program includes three tracks on: privacy, legal and ethical issues in information and knowledge systems; knowledge transfer in the age of artificial intelligence; and E-learning in pandemic time: challenges and opportunities.

Due to the ongoing Covid-19 pandemic, the conference will be held online this year. Registrations to the conference are currently open and please refer to the conference website for more detail on registration and the program.

The 18th International Conference on Artificial Intelligence and Law 2021 (ICAIL 2021)

Online Conference, June 21-25, 2021
https://icail.lawgorithm.com.br/
The ICAIL conference is the primary international conference addressing research in Artificial Intelligence and Law, and has been organized biennially since 1987 under the auspices of the International Association for Artificial Intelligence and Law (IAAIL). It provides a forum for the presentation and discussion of the latest research results and practical applications in the area fostering interdisciplinary and international collaboration. The conference will feature a main track for technical papers, a demonstration track, workshops, tutorials, a doctoral consortium and best paper prizes. The research contributions focus on a variety of topics with particular attention on three tracks: innovative applications in AI and Law; ethical and legal issues of AI technology and its applications; and reproducible results and data.

The conference will be held online due to the ongoing Covid-19 pandemic. Registrations to the conference are currently open and please refer to the conference website for more detail on registration and the program.

The 11th International Conference on Knowledge Capture (KCAP 2021)

Online Conference, December 2-3, 2021
https://www.k-cap.org/2021/
KCAP is aimed at attracting researchers from diverse areas of Artificial Intelligence, including knowledge representation, knowledge acquisition, intelligent user interfaces, problem-solving and reasoning, planning, agents, text extraction, machine learning, information enrichment, and visualization. It also encourages the participation of researchers interested in cyber-infrastructures to foster publication, retrieval, reuse, and integration of data. The conference highlights the importance of

precise and scalable knowledge capture given the increasing demand for higher quality, accurate information that can be effectively and unambiguously used.

In light of promoting discussions around this area, the conference seeks to publish cutting-edge research on topics of interest that include but not limited to work on: knowledge capture for mulitimodal data and semantic web; problem-solving knowledge and methods; knowledge engineering and knowledge governance; ontology design patterns; and addressing scalability issues for distributed knowledge graphs.
Submission deadline: September 15, 2021

The 13th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2021)

Online Conference, 25-27 October, 2021
http://www.ic3k.org/
The purpose of the IC3K is to bring together researchers, engineers and practitioners on the areas of knowledge discovery, knowledge engineering and knowledge management. IC3K is composed of three co-located conferences, each specialized in at least one of the three main knowledge areas. The conference seeks high quality research work submissions for the regular track, position papers, workshops, doctoral consortium, special sessions, tutorials, and demos.

Submission deadlines: Position Papers: June 25, 2021; Doctoral Consortium: July 30, 2021; Special Sessions: July 30, 2021; Tutorials/Demos/Panels: September 13, 2021

The First International Conference on AI-ML-Systems (AIMLSystems'21)

Bangalore, India, October 21-23, 2021
https://www.aimlsystems.org/2021/
AIMLSystems'21 is organized for the first time with the intention to cover the white-space at the intersection of research on data-intensive AI/ML techniques and systems engineering. The conference aims to bring out and highlight the natural connections between the two fields. Specifically, it is expected to explore how immense strides in AI/ML techniques are made possible through computational systems research, how the use of AI/ML can help in the continuous and workload-driven design space exploration of computational systems and, the use of AI/ML in the design of socio-economic systems such as public healthcare, and security. The conference series is an initiative of the COMSNETS Association, a not-for-profit organization, which has organized the prestigious COMSNETS networking conference each year since 2009.

Topics of interest for research papers are broadly categorized into three streams: systems for AI/ML such as CPU/GPU architectures, embedded hardware, and hardware efficient ML methods; AI/ML for systems such as AI/ML for compiler optimization, networks and power management; and applications of AI/ML in socio-economic systems design such as applications of AI/ML in financial systems, and cyber-physical systems.
Submission deadline: July 13, 2021

Dilini Samarasinghe is the Assistant Conference Coordination Officer for ACM SIGAI, and a 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: 10.1145/3478369.3478372

This section is compiled from reports of recent events sponsored or run in cooperation with ACM SIGAI. In general these reports were written and submitted by the conference organisers.

The ACM- SIGAI POST CONFERENCE REPORT 12th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K '21)

Online 11/02/2020-11/04/2020
http://www.ic3k.org/?y=2020
The purpose of the IC3K is to bring together researchers, engineers and practitioners on the areas of Knowledge Discovery, Knowledge Engineering and Knowledge Management. IC3K is composed of three co-located conferences, each specialized in at least one of the aforementioned main knowledge areas.

IC3K 2020 was exceptionally held as an online streaming event, due to covid-19, from November 2 to 4, 2020. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC).

IC3K 2020 was also organized in cooperation with the ACM SIGAI - ACM Special Interest Group on Artificial Intelligence, the Portuguese Association for Artificial Intelligence, the Associação Portuguesa de Reconhecimento de Padrões (APRP), and the Associazione Italiana per l'Intelligenza Artificiale.

IC3K received 133 paper submissions from 44 countries. To evaluate each submission, a double-blind paper review was performed by the Program Committee. After a stringent selection process, 21.05% of the papers were published and presented as full papers, i.e. completed work (12 pages/25' oral presentation).

Four invited talks were delivered by internationally distinguished speakers, namely:

- Alexander Smirnov, SPIIRAS, Russian Federation
- Manfred Reichert, Ulm University, Germany
- Frank van Harmelen, The Hybrid Intelligence Center & Vrije Universiteit Amsterdam, Netherlands Stefan Decker, RWTH Aachen University, Germany

Additionally, the conference acknowledged "Best Paper Awards", "Best Student Paper Awards", "Best Poster Awards", one "Best PhD Project Award", and in some cases a "Best Industrial Paper Award", which were all conferred during the conference. More information can be found at: http://www.ic3k.org/PreviousAwards.aspx

The International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC '21)

Online 02/11/2021-02/13/2021
http://www.biostec.org/?y=2021
The purpose of BIOSTEC is to bring together researchers and practitioners, including engineers, biologists, health professionals and informatics/computer scientists, interested in both theoretical advances and applications of information systems, artificial intelligence, signal processing, electronics and other engineering tools in knowledge areas related to biology and medicine. BIOSTEC is composed of five co-located conferences, each specialized in a different knowledge area.

BIOSTEC 2021 was exceptionally held as an online streaming event, due to covid-19, from February 11-13, 2021. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC) and technically co-sponsored by the IEEE Engineering in Medicine and Biology Society.

BIOSTEC 2021 was also organized in cooperation with the ACM SIGAI - ACM Special Interest Group on Artificial Intelligence, the ACM Special Interest Group on Management Information Systems (SIGMIS), the ACM

Special Interest Group on Accessible Computing (SIG ACCESS), the ACM Special Interest Group on Bioinformatics, Computational Biology, and Biomedical Informatics (SIGBio), the Association for the Advancement of Artificial Intelligence (AAAI), the International Society for Computational Biology (ISCB), the Finnish Society for Medical Physics and Medical Engineering, the Swiss Society for Biomedical Engineering, the German Society for Biomedical Engineering (VDE DGBMT), the Societe Française de Genie Biologique et Medical (SFGBM), and TELIGHT.

BIOSTEC received 265 paper submissions from 52 countries. To evaluate each submission, a double-blind paper review was performed by the Program Committee. After a stringent selection process, 25% of the papers were published and presented as full papers, i.e. completed work (12 pages/25' oral presentation).

Four invited talks were delivered by internationally distinguished speakers, namely:

- Athanasios Tsanas, University of Edinburgh, United Kingdom
- Thomas Ostermann, Universität Witten/Herdecke, Germany
- Mireya Fernández Chimeno, Universitat Politècnica de Catalunya, Spain
- Tiago Guerreiro, Faculdade de Ciências, Universidade de Lisboa, Portugal

Additionally, a "Best Paper Award", a "Best Student Paper Award", and a "Best Poster Award" were conferred for each of the sub-conferences that compose BIOSTEC during the event. Furthermore, a "Best Industrial Paper Award" was granted in HEALTHINF, and a "Best PhD Project Award" was granted in the Doctoral Consortium of the BIOSTEC joint conference. More information can be found at: http://biostec.org/PreviousAwards.aspx

IEEE AIVR 2020 (IEEE International Conference on Artificial Intelligence and Virtual Reality 2020)

Online 12/14/2021-12/18/2021
https://aivr.science.uu.nl/index.html
The main intention of IEEE AIVR is to bring together the communities of AI researchers and VR/AR researchers and everyone working on the intersection between these two areas. While there are plenty events in either of these fields where related researchers could publish their work, IEEE AIVR was founded four years ago based on the observation that having a unique event combining those fields would be beneficial because it could bring together these people distributed over different conferences where their work might only be considered a niche area, it could foster and encourage collaboration between different fields, and in the long run it could create a new, interdisciplinary community. There is still a long way to go, but after its third incarnation now, we see promising developments in this direction.

Like pretty much all conferences scheduled after March 2020, IEEE AIVR 2020 had to be held virtually due to the COVID-19-enforced restrictions. To accommodate for different time zones, the organizers decided to switch from a three full-day schedule to a five half-day schedule with online presentations using video conferencing and interactive poster and demo sessions using the web-based Mobilla Hubs VR environment. Despite the special circumstances due to the pandemic, IEEE AIVR 2020 saw a steady growth in registrations, passing the 100 mark for the first time, and in submissions, which also showed an increase compared to previous years. It is particularly pleasing to see that now in its third incarnation, IEEE AIVR seems to develop a unique community attracting people from different areas. While the majority of submissions is coming from VR/AR-related domains, the relevance that AI does or potentially can play in these works was clearly visible. The event was complemented by five workshops that also addressed interesting and relevant topics at the intersection between AI and VR/AR: One on ethics in AI and VR/AR, one on capturing and rendering of digital humans for AR/VR, one on Immersive Analytics, one on the animation of realistic crowds and humans, and one on VR/AR in medical imaging.

15th INTERNATIONAL CONFERENCE ON THE FOUNDATIONS OF DIGITAL GAMES (FDG 2020)

Online 09/15/2020-09/18/2020
http://fdg2020.org
The 15th International Conference on the Foundations of Digital Games (FDG'20) took place during 15-18 September as a virtual conference hosted by the University of Malta. FDG is an interdisciplinary conference on technology used to develop digital games and the study of digital games and their design, where academics can present their work to a diverse audience, share new ideas, and find collaborations with different backgrounds. The target audience primarily consists of games researchers and game AI researchers. The theme for FDG'20 was "Games and their Heritage". FDG'20 was originally planned to be hosted in Bugibba, on the Mediterranean island of Malta. Due to the global pandemic and to ensure everyone's well-being, FDG'20 took place in a virtual format with all presenters and attendees participating remotely. Remote attendance permitted a low admission price which in turn led to a much higher attendance than previous years, reaching over 300 registrations on the days of the conference. The conference took place during 15-18 September, with the 11th and 14th of September dedicated to the workshops of FDG'20 and its Doctoral Consortium.

FDG'20 received a total of 143 research paper submissions in eight different research tracks. Of these, 51 were accepted (acceptance rate: 36%). In addition, 44 short papers were accepted from the 143 research paper submissions and from 25 late-breaking paper submissions (overall acceptance rate: 57%). FDG'20 also hosted five workshops, four of which had papers published in the proceedings. Some workshops had a long history within FDG, such as the 11th workshop on Procedural Content Generation, while there were also two new workshops at FDG'20: the Workshop on Forgetting and Remembering and the Workshop on Digital Games for Digital Literacy and Computational Thinking. FDG'20 also included a procedural content generation competition, two panels, six demonstrations, and five presentations at the Doctoral Consortium. Five keynotes framed the conference:

- AI in Content Production Automation, by Sahar Asadi
- Procedural Content Generation as Art, Design and Play, by Tom Betts
- Photography Workshops for Spacemen & Parties in Spreadsheets, by Marie Foulston
- From the Ivory Tower to the Skunk Works - Bringing Research to Life in Industry, by Luke Dicken
- Back to school with Minecraft, by Deirdre Quarnstrom

A best paper ballot included six papers from different tracks. Attendees' voting decided the best paper of FDG'20: "Player-Centered AI for Automatic Game Personalization: Open Problems", by Jichen Zhu and Santiago Ontanon. The runner-up best paper was "Crafting is So Hardcore: Masculinized Making in Gaming Representations of Labor" by Anne Sullivan, Mel Stanfill and Anastasia Salter.

Proceedings of the FDG'20 conference and its workshops can be found at: https://dl.acm.org/doi/proceedings/10.1145/3402942

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.

Enlichenment: Insights Towards AI Impact in Education through a Mycelial Partnership between Research, Policy, and Practice

Carolyn P. Rosé (Carnegie Mellon University; cprose@cs.cmu.edu)
DOI: 10.1145/3478369.3478373

Introduction

This column raises the question, as we begin to emerge from COVID 19, what is the role of the field of AI in this emerging reality? We specifically consider this in the face of tremendous learning loss and widening achievement gaps. In this wake, what specifically is the role of AI in the future of education as we move forward? This question bridges the worlds of basic research and the seemingly distant worlds of policy and practice.

History of AI in Education

Since Jaime Carbonell, pioneer in multiple areas of Artificial Intelligence, conducted landmark research on the SCHOLAR intelligent tutoring system in the early 70s, the field of AI and the fields of Human Learning and Teaching have partnered together to study how to use AI-enabled technology to understand and support human learning.

This partnership has birthed multiple interdisciplinary research societies, brought together under the umbrella of the International Alliance to Advance Learning in a Digital Era (IAALDE)[1], an international partnership involving at least 3,000 researchers worldwide. Nevertheless, despite tremendous growth in these research societies over the decades alongside advances in the core field of Artificial Intelligence, and notable large scale success of some learning technologies in large scale use, the emergency move to universal online learning at all levels over the past year has exposed gaps and breakdowns in the path from basic research into practice. As the new administration reacts by committing to invest substantial research dollars into addressing the COVID Melt, or learning loss, we must ask ourselves how to prepare for potentially future emergencies so that such tremendous and inequitable learning loss will not be repeated.

[1] https://alliancelss.com/

Addressing the Challenge through Public Engagement with Science

The researchers of IAALDE came to this project with the question: What does it mean for a meta-society dedicated to basic research like IAALDE to engage in partnership with the worlds of policy and practice? Adopting the metaphor of a mycelial network, IAALDE collaborated with the American Association for the Advancement of Science (AAAS)[2] to foster productive synergy between the worlds of research, policy, and practice.

Mycelial networks, not unlike neural networks in many ways, are the fungal networks that foster connection, communication, and ultimately the health and survival of forests and other large scale ecosystems. What might at the surface appear to be separate organisms living sometimes at great distance are in fact parts of an intricate and entangled system that thrives through symbiotic partnership. As a recent step, this IAALDE-AAAS collaboration has convened a public engagement project meant to engage the world of policy and practice in an emerging partnership to forge a new path forward in the image of a mycelium.

In the context of the IAALDE public engagement workshop[3], Administrators, policy makers, and implementors of policy were invited to engage with world class leading researchers across a broad spectrum of research in technology enhanced learning to accelerate the path from research into real educational impact through practice. The long term goal is for research, policy and practice going forward would benefit from more intensive, ongoing, discussion and coordination between communities. At the same time that greater awareness of research findings might offer opportunities to reflect and reconsider practices on the ground in schools. This discussion, involving over 100 delegates, was meant to lay the

[2] https://www.aaas.org/

[3] https://alliancelss.com/#publicengagement

foundation for documents, resources, and activities to move the conversation forward.

Highlights of Lessons Learned

Through interaction with colleagues across societies, in the AAAS, and other organizations, we designed the public engagement workshop to be collaborative and interactive. We began with a collaborative focus setting panel designed to offer a range of basic research topics within the expertise of the panelists and to engage the audience in ranking those topics and submitting specific pain points and questions related to the highlight ranked topics. The idea was to organize break out working sessions on the fly where there was a synergy of expertise, interest, and energy to bring together researchers and audience members to engage in problem finding and brainstorming related to shared concerns.

In particular, 23 experts, each representing one of the 10 partnering research societies, were divided into three panels of experts based on topic focus. For panel 1, the most interest from the audience was expressed in personalized learning at scale, the role of the teacher, and challenges in using data to improve instruction. Concerns included issues regarding orchestration challenges in heterogeneous classrooms, needed professional development to engage teachers in strategic thinking about use of data and adaptive systems in classroom teaching, fears of teachers regarding adoption of AI technologies such as facial recognition, and tensions between massification rather than personalization. For panel 2, the most support was expressed for topics related to offering support for teaching and learning, support for scientific reasoning and discourse, and providing instruction in low resource contexts. Concerns included a sense that available data might not be informative for the most important concerns, or that it might not seem directly meaningful to teachers, that deployed technologies might not be achieving the desired level of social presence, that the sets of deployed technologies have not yet been integrated into continua that support students along complete trajectories, and that in the absence to solutions for low resource conditions, the achievement gaps are widening. For panel 3, the most support was expressed for topics related to reducing barriers to supporting equity, learning technologies and interventions to support large, diverse populations, and optimizing human learning for all learners. Concerns related to inclusively, racism, and equal access to resources. Not surprisingly, concerns related to AI fairness were expressed. More surprising were expressed concerns about the rapid pace of change in the AI space and fears about the infrastructure not being capable of keeping up the pace.

The majority of researchers who came and engaged in the discussion represented basic research that already engaged with practice within their own research programs. Participants who came from policy efforts, funding agencies, companies, non-profit organizations, and schools who spoke up were mainly ones who were familiar enough with research to join in a discussion with academics. What became clear is that while these researchers and these policy people and these practitioners might not have been typical of any of these groups, these may be the boundary-spanning people to push the communication forward from a network perspective. The next important step, which began during the workshop, is to synthesize wisdom across these local efforts and package it for large scale use.

On June 24 at 12pm EDT, there will be an ACM SIGAI Learning Webinar, which will expand on insights gained from the public engagement workshop. A link to the recorded webinar will be found online[4].

Carolyn P. Rosé is a Professor of Language Technologies and Human-Computer Interaction in the School of Computer Science at Carnegie Mellon University. She is a 2020-2021 AAAS Fellow under the Leshner Institute for Public Engagement with Science, with a focus on public engagement with Artificial Intelligence.

[4] https://webinars.on24.com/acm/rose

AI Policy Matters

Larry Medsker (The George Washington University; lrm@gwu.edu)

DOI: 10.1145/3478369.3478374

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.

Should the government play a role in reducing algorithmic bias?

On March 12, the Center for Technology Innovation at Brookings hosted a webinar on the role of government in identifying and reducing algorithmic biases (video). Speakers discussed what is needed to prioritize fairness in machine-learning models and how to weed out artificial intelligence models that perpetuate discrimination. Questions included

> How do the European Union, U.K., and U.S. differ in their approaches to bias and discrimination?
>
> What lessons can they learn from each other?
>
> Should approaches to AI bias be universally applied to ensure civil and human rights for protected groups?

They observed that "policymakers and researchers throughout the world are considering strategies for reducing biased decisions made by machine-learning algorithms. To date, the U.K. has been the most forward in outlining a role for government in identifying and mitigating biases and their unintended consequences, especially decisions that impact marginalized populations. In the U.S., legislators and policymakers have focused on algorithmic accountability and the explanation of models to ensure fairness in predictive decision making."

The moderator was Alex Engler, Rubenstein Fellow, Brookings Governance Studies. The speakers and discussants were

> Lara Macdonald and Ghazi Ahamat, Senior Policy Advisors, UK Centre for Data Ethics and Innovation
>
> Nicol Turner Lee, Brookings Senior Fellow Governance Studies, and Director, Center for Technology Innovation
>
> Adrian Weller, Programme Director for AI at the Alan Turing Institute

Algo2021 Conference

On April 29, 2021, the University College London held online The Algo2021 Conference: Ecosystems of Excellence and Trust, building upon the success of their 2020 inaugural conference. They platformed all major stakeholders – academia, civil service, and industry – by showcasing the cutting-edge developments, contemporary debates, and perspectives of major players. The 2021 conference theme reflects the desire to promote public good innovation. Sessions and topics included the following:

> Machine Learning in Healthcare
>
> Trust and the Human-on-the-Loop
>
> Artificial Intelligence and Predictive Policing
>
> AI and Innovation in Healthcare Tech
>
> AI in Learning and Education Technologies
>
> Building Communities of Excellence in AI
>
> Human-AI and Ethics Issues.

Politico's Online AI Summit

The 2021 Summit on May 31 dissected Europe's AI legislative package, along with the impact of geopolitical tensions and tech regulations, on topics such as data and privacy concerns. The summit convened top EU and national decision makers, opinion formers, and tech industry leaders.

"The European Commission will soon introduce legislation to govern the use of AI, acting on its aim to draw up rules for the technology sector over the next five years and on its legacy as the world's leading regulator of digital privacy. At the heart of the issue is the will to balance the need for rules with the desire to boost innovation, allowing the old continent to assert its digital sovereignty. On where the needle should be, opinions are divided – and the publication of the Commission's draft proposal will not be the end of the discussion." Issues addressed are the following:

> How rules may fit broader plans to build European tech platforms that compete globally with other regions
>
> How new requirements on algorithmic transparency might be viewed by regular people
>
> What kind of implementation efforts will be required for startups, mid-size companies and big tech.

The Politico 4th edition of the AI Summit addressed important questions in panel discussions, exclusive interviews, and interactive round table discussions. Top regulators, tech leaders, startups, and civil society stakeholders examined the EU's legislative framework on AI and data flow while tackling uncomfortable questions about people's fundamental rights, misinformation, and international cooperation that will determine the future of AI in Europe and worldwide.

What exactly are AI and DS?

Maybe it's too late, but how about sharpening up the current terminology in discussions and descriptions of AI systems and products? From the news media and marketing to the intentional misinformation elements of our times, we could use some clarity, including in policymaker discussions about AI and understandings about AI in data science. John Launchbury's "A DARPA Perspective on Artificial Intelligence," DARPA, 4-7 (Feb. 2017), is a good resource. The final report of the National Security Commission on Artificial Intelligence (NSCAI) describes two waves of AI: "AI technologies and applications such as pattern recognition, machine learning, computer vision, natural language understanding, and speech recognition have evolved for many decades. In the early years of AI, the period the Defense Advanced Research Projects Agency (DARPA) describes as the 'first wave,' researchers explored many approaches, including symbolic logic, expert systems, and intelligent planning. Some of the most effective results were based on 'handcrafted knowledge' defined by humans and then used by machines to imitate the way humans reason and interact."

"Within the past 10 years, we have witnessed a 'second wave' of AI, propelled by large-scale statistical machine learning that enables engineers to create models that can be trained to specific problem domains if given exemplar data or simulated interactions. Learning from data, these systems are designed to solve specific tasks and achieve particular goals with competencies that, in some respects, parallel the cognitive processes of humans: perceiving, reasoning, learning, communicating, deciding, and acting. Today most fielded large-scale AI systems employ elements of both first- and second-wave AI approaches."

The new generation of data scientists is not always aware of the history of their craft. Insights and experience from the first wave, including issues of AI Ethics, are important for Data Science to be fully a new discipline.

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.

Learning Natural Language Interfaces with Neural Models

Li Dong (Microsoft Research; lidong1@microsoft.com)

DOI: 10.1145/3478369.3478375

Introduction

Language is the primary and most natural means of communication for humans. The learning curve of interacting with various services (e.g., digital assistants, and smart appliances) would be greatly reduced if we could talk to machines using human language. However, in most cases computers can only interpret and execute formal languages.

The research goal of my dissertation is to use neural models to build natural language interfaces which learn to map naturally worded expressions onto machine-interpretable representations. The task is challenging due to (1) structural mismatches between natural language and formal language, (2) the well-formedness of output representations, (3) lack of uncertainty information and interpretability, and (4) the model coverage for language variations. In this dissertation, we develop several flexible neural architectures to address these challenges.

The dissertation presents a framework based on encoder-decoder neural networks for natural language interfaces. Beyond sequence modeling, we propose a tree decoder to utilize the compositional nature and well-formedness of meaning representations, which recursively generates hierarchical structures in a top-down manner. To model meaning at different granularity levels, we present a structure-aware neural architecture which decodes semantic representations following a coarse-to-fine procedure.

The proposed neural models remain difficult to interpret, acting in most cases as a black box. We explore ways to estimate and interpret the model's confidence in its predictions, which we argue can provide users with immediate and meaningful feedback regarding uncertain outputs. We estimate confidence scores that indicate whether model predictions are likely to be correct. Moreover, we identify which parts of the input contribute to uncertain predictions allowing users to interpret their model.

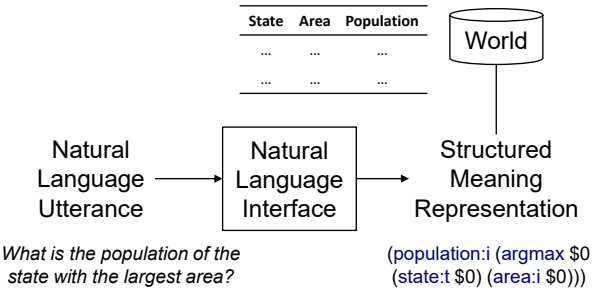

Figure 1: The goal of natural language interfaces is allowing users to interact with computers in human language. As shown by the example from the GEO dataset (Zelle & Mooney, 1996; Zettlemoyer & Collins, 2005), the model maps the input question to the λ-calculus meaning representation, and then execute it over the database to obtain the answer.

Model coverage is one of the major reasons resulting in uncertainty of natural language interfaces. Therefore, we develop a general framework to handle the many different ways natural language expresses the same information need. We leverage external resources to generate felicitous paraphrases for the input, and then feed them to a neural paraphrase scoring model which assigns higher weights to linguistic expressions most likely to yield correct answers. The model components are trained end-to-end using supervision signals provided by the target task.

Experimental results show that the proposed neural models can be easily ported across tasks. Moreover, the robustness of natural language interfaces can be enhanced by considering the output well-formedness, confidence modeling, and improving model coverage.

Neural Semantic Parsing

Semantic parsing is the task of translating text to a formal meaning representation such as logical forms or structured queries, which is one of the core components of natural language interfaces (as shown in Figure 1).

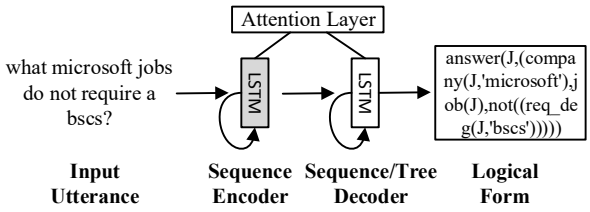

Figure 2: Input utterances and their logical forms are encoded and decoded with neural networks. An attention layer is used to learn soft alignments.

There has recently been a surge of interest in developing machine learning methods for semantic parsing, due in part to the availability of corpora containing utterances annotated with formal meaning representations. In order to predict the correct logical form for a given utterance, most previous systems rely on predefined templates and manually designed features, which often render the parsing model domain- or representation-specific.

In the dissertation, we aim to use a portable method to bridge the gap between natural language and logical form with minimal domain and linguistic knowledge (Dong & Lapata, 2016). The proposed framework (as shown in Figure 2) is portable as the models can be end-to-end trained by giving annotated data, namely, natural language utterances paired with their meaning representations. So we can easily adapt the models to different applications with minimal efforts. The structural gap between inputs and outputs is bridged by neural encoder-decoder networks augmented with attention mechanisms.

Constrained Decoding

After obtaining the output meaning representations from natural language interfaces, we usually need to execute them to obtain user intentions. Because the downstream executors only accept grammatical programs, it is beneficial to explicitly model the structure of predictions. The structural information of the output should be taken into consideration so that the models can generate well-formed meaning representations.

The fact that meaning representations are typically structured objects prompts efforts to develop neural architectures which explicitly account for their structure. In order to guarantee

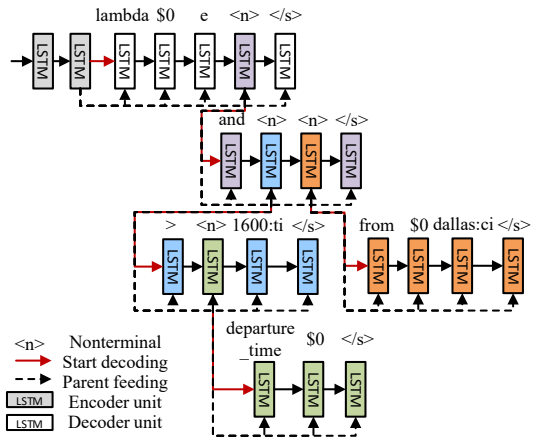

Figure 3: Sequence-to-tree model (Dong & Lapata, 2016) with a hierarchical tree decoder.

the well-formedness of the output we explicitly model the hierarchical and compositional nature of meaning representations, and thus develop a sequence-to-tree model (Dong & Lapata, 2016), and a coarse-to-fine decoding algorithm (Dong & Lapata, 2018).

In the first solution, the proposed tree decoder defines a placeholder to indicate nonterminal nodes as shown in Figure 3. Tree structures are recursively generated in a top-down, and left-to-right manner. The explicit modeling of hierarchical structures constrains results in the space of well-formed trees. In other words, ill-formed logical forms can be pruned from the candidate set. In order to model meaning at different levels of granularity, the second solution uses a structure-aware neural architecture to decode semantic representations from coarse to fine. The coarse meaning decoder first generates a rough sketch of the meaning representation, which omits low-level details, such as arguments and variable names. Then, the fine meaning decoder fills in missing details by conditioning on the input utterance and the sketch itself.

Confidence Modeling

Neural semantic parsing models map natural language text to a formal meaning representation (e.g., logical forms or SQL queries). However, despite achieving promising results, the neural semantic parsers remain difficult to interpret, acting in most cases as a black box, not providing any information about what

made them arrive at a particular decision. My dissertation explores ways to estimate and interpret the model's confidence in its predictions, which we argue can provide users with immediate and meaningful feedback regarding uncertain outputs (Dong, Quirk, & Lapata, 2018).

An explicit framework for confidence modeling would benefit the development cycle of neural semantic parsers which, contrary to more traditional methods, do not make use of lexicons or templates and as a result the sources of errors and inconsistencies are difficult to trace. Moreover, from the perspective of application, semantic parsing is often used to build natural language interfaces, such as dialogue systems. In this case it is important to know whether the system understands the input queries with high confidence in order to make decisions more reliably. For example, knowing that some of the predictions are uncertain would allow the system to generate clarification questions, prompting users to verify the results before triggering unwanted actions. In addition, the training data used for semantic parsing can be small and noisy, and as a result, models do indeed produce uncertain outputs, which we would like our framework to identify.

We categorize the causes of uncertainty into three types, namely *model uncertainty*, *data uncertainty*, and *input uncertainty* and design different metrics to characterize them. Furthermore, we propose a method based on backpropagation which allows to interpret model behavior by identifying which parts of the input contribute to uncertain predictions.

Query Paraphrasing

One of the challenges to build a robust natural language interface is model coverage. Due to the limited size of training data, it is challenging to handle the many different ways natural language expresses the same information need. As a result, small variations in semantically equivalent inputs may yield different results. For example, a hypothetical natural language interface must recognize that the questions "*who created microsoft*" and "*who started microsoft*" have the same meaning and that they both convey the `founder` relation in order to obtain the correct answer from a knowledge base. Moreover, one of the main causes of uncertainty is defined as data uncertainty, namely, uncertainty of predictions affected by the coverage of training data. If the pattern of input is unseen by the model on training data, it is difficult to predict reliable outputs. We leverage external resources to rewrite the natural language input during both training and test, so that model coverage can be increased by augmenting the original expression with its variations (Dong, Mallinson, Reddy, & Lapata, 2017).

Next Steps

This dissertation focuses on a single domain, and were trained on English utterances paired with their meaning representations. It is worth studying and exploring how to improve the proposed models' scalability in terms of supporting many different domains, languages, and supervision signals:

- **Cross-Domain Sharing.** For real-world applications (such as voice assistants), a system usually needs to handle many different domains. It is helpful to transfer and share knowledge across domains and meaning representations, especially when the data size of each domain is not large enough. We can share the common operators, predicates, and composition structures for similar examples.
- **Zero/Few-Shot Learning.** The training data size for a new domain is often small or even non-existent. It is valuable to conduct few-shot or zero-shot learning for a cold start, so that the model can be quickly adapted to a new similar domain.
- **Data Collection.** Model performance can usually be improved if more annotated data is fed to the model. However, sometimes it is difficult to directly annotate meaning representations for ordinary users. A new paradigm of training data collection is critical for the acceleration of model deployment.
- **Weakly Supervised Learning.** The models presented in this dissertation were trained on natural language utterances paired with their meaning representations. Given the data paucity of such parallel corpora, we can learn models from weak supervision signals (e.g., question-answer pairs),

which would reduce the annotation burden. Weakly supervised learning also provides a way to utilize online user feedback.

- **Multilingual Semantic Parsing.** Natural language interfaces should accept multiple languages, so users from different world regions can freely use their native language. A typical problem is that for some languages there is less data compared to others, which inevitably results in inferior semantic parsing performance.
- **Multi-Turn Interactions.** Sometimes users express their intentions in multiple utterances or update the requests according to the system responses. Users often tend to omit information which has been expressed in the conversation history. So the prediction should be conditioned on the current utterance as well as the interaction history.

References

Dong, L., & Lapata, M. (2016, August). Language to logical form with neural attention. In *Proceedings of the 54th annual meeting of the association for computational linguistics* (pp. 33–43). Berlin, Germany.

Dong, L., & Lapata, M. (2018). Coarse-to-fine decoding for neural semantic parsing. In *Proceedings of the 56th annual meeting of the association for computational linguistics* (pp. 731–742). Association for Computational Linguistics.

Dong, L., Mallinson, J., Reddy, S., & Lapata, M. (2017, September). Learning to paraphrase for question answering. In *Proceedings of the 2017 conference on empirical methods in natural language processing* (pp. 875–886). Copenhagen, Denmark: Association for Computational Linguistics.

Dong, L., Quirk, C., & Lapata, M. (2018). Confidence modeling for neural semantic parsing. In *Proceedings of the 56th annual meeting of the association for computational linguistics* (pp. 743–753). Association for Computational Linguistics.

Zelle, J. M., & Mooney, R. J. (1996). Learning to parse database queries using inductive logic programming. In *Proceedings of the 13th national conference on artificial intelligence* (pp. 1050–1055). Portland, Oregon.

Zettlemoyer, L., & Collins, M. (2005). Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Proceedings of the 21st conference on uncertainty in artificial intelligence* (pp. 658–666). Edinburgh, Scotland.

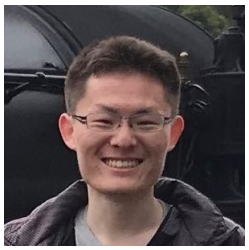

Li Dong is a Researcher at Microsoft Research, working on natural language processing, and representation learning. He received his PhD in School of Informatics at University of Edinburgh. Li's research has been recognized through the AAAI/ACM SIGAI Doctoral Dissertation Award Runner Up, the ACL-2018 Best Paper Honourable Mention, the AAAI-2021 Best Paper Runner Up, and fellowship from Microsoft.