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: 10.1145/3375637

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

- to include the article in the ACM Digital Library and in any Digital Library related services

- to allow users to make a personal copy of the article for noncommercial, educational or research purposes

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

Submit to AI Matters!

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

AI Matters Editorial Board

Amy McGovern, Co-Editor, *U. Oklahoma*
Iolanda Leite, Co-Editor, *KTH*
Sanmay Das, *Washington Univ. in Saint Louis*
Alexei Efros, *Univ. of CA Berkeley*
Susan L. Epstein, *The City Univ. of NY*
Yolanda Gil, *ISI/Univ. of Southern California*
Doug Lange, *U.S. Navy*
Kiri Wagstaff, *JPL/Caltech*
Xiaojin (Jerry) Zhu, *Univ. of WI Madison*

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 5(4)

Amy McGovern, co-editor (University of Oklahoma; aimatters@sigai.acm.org)
Iolanda Leite, co-editor (Royal Institute of Technology (KTH); aimatters@sigai.acm.org)
Anuj Karpatne, co-editor (Virginia Tech; aimatters@sigai.acm.org)
DOI: 10.1145/3375637.3375638

Issue overview

Welcome to the fourth issue of the fifth volume of the AI Matters Newsletter. [DO WE NEED TO SAY ANYTHING MORE HERE?]

This issue is full of great new articles and stories for you! In our regular articles, we open with an event report by Louise Dennis on recent conferences relevant to SIGAI organized around the world. Louise Dennis then provides a summary of upcoming AI events. In the policy column, Larry Medsker covers recent policies on national strategy for AI laid down by the U.S. government, implications of AI on workplace transitions, President Trump's executive order on re-establishing the President's Council of Advisors on Science and Technology (PCAST), and more. Our final regular column is our AI crosswords from Adi Botea. Enjoy!

We have three contributed articles in our current issue. In the first article, Sharad Sinha and Clint George provide a summary of the events and activities of a recent workshop conducted in India on AI for high school students. In the second article, Trevor Watkins describes the Cosmology of Artificial Intelligence Project (CAIP) that uses innovative visualization techniques to present the evolution of the field of AI. In the third article, Virginia Dignum discusses how AI is progressively trending to become a multi-disciplinary field of reseach and the implications of this trend on the future of AI.

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.

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

Iolanda Leite is co-editor of AI Matters. She is an Assistant 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.

Conference Reports

Louise A. Dennis (University of Liverpool; L.A.Dennis@liverpool.ac.uk)

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

11th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 19)

Vienna, Austria, September 17-19, 2019
fhttp://www.ic3k.org/?y=2019
IC3K 2019 was held in Vienna, Austria from 17 to 19 of September, 2019. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC) and co-organized by the Vienna University of Technology (TU Wien). IC3K 2019 was organized in cooperation with the ACM SIGAI and several other national societies. IC3K 2019 is composed of three co-located complementary conferences, each specialized in at least one of the aforementioned main knowledge areas. Namely: KDIR, 11th International Conference on Knowledge Discovery and Information Retrieval; KEOD, 11th International Conference on Knowledge Engineering and Ontology Development; and KMIS, 11th International Conference on Knowledge Management and Information Systems. IC3K received 220 paper submissions from 57 countries. To evaluate each submission, a double blind paper review was performed by the Program Committee. 22.27% of the papers were published and presented as full papers. Four invited talks were delivered by internationally distinguished speakers, namely: Robert Pergl, Czech Technical University in Prague, Czech Republic; Eduard Babkin, National Research University "Higher School of Economics", Russian Federation; Heimo Müller, Medical University of Graz, Austria; Kayyali Mohamed, International Federation of Global & Green Information Communication Technology, United States. The IC3K program also included a tutorial on How to Extract Knowledge from Communications: Combining Pragma-linguistics

and Knowledge Engineering Techniques, lectured by Nada Matta, Hassan Atifi and Nelson Tenório and a Doctoral Consortium chaired by Jorge Bernardino, David Aveiro and Ana Salgado. Additionally, a Best Paper Award, a Best Student Paper Award and a Best Poster Award were conferred at the conference.

11th International Joint Conference on Computational Intelligence (IJCCI 19)

Vienna, Austria, September 17-19, 2019
http://www.ijcci.org/?y=2019
IJCCI 2019 was held in Vienna, Austria from 17 to 19 of September, 2019. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC) and co-organized by the Vienna University of Technology (TU Wien). It was also technically co-sponsored by the International Federation of Automatic Control (IFAC) and by the World Federation on Soft Computing. IJCCI 2019 was organized in cooperation with the ACM SIGAI and several other national societies. IJCCI 2019 is composed of three co-located complementary conferences, each specialized in at least one of the aforementioned main knowledge areas. Namely: ECTA, 11th International Conference on Evolutionary Computation Theory and Applications; FCTA, 11th International Conference on Fuzzy Computation Theory and Applications; and NCTA, 11th International Conference on Neural Computation Theory and Applications. IJCCI received 86 paper submissions from 32 countries. 24.42% of the papers were published and presented as full papers. Four invited talks were delivered by internationally distinguished speakers, namely: Pietro Oliveto, University of Sheffield, United Kingdom; Vesna Sesum-Cavic, TU Vienna, Austria; Andreas Holzinger, Medical University Graz, Austria; Jonathan Garibaldi, University of Nottingham, United Kingdom. Additionally, a Best Paper Award, a Best Student Paper Award and a Best Poster Award were conferred in each of the co-located complimentary conferences (ECTA, FCTA and NCTA).

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

📅 Events

Louise A. Dennis (University of Liverpool; L.A.Dennis@liverpool.ac.uk)

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.

17th IEEE/ACM International Conference on Mining Software Repositories (MSR 2020)

Seoul, South Korea, May 25-26, 2020
https://2020.msrconf.org/
The Mining Software Repositories (MSR) conference is the premier conference for data science, machine learning, and artificial intelligence in software engineering. The goal of the conference is to improve software engineering practices by uncovering interesting and actionable information about software systems and projects using the vast amounts of software data such as source control systems, defect tracking systems, code review repositories, archived communications between project personnel, question-and-answer sites, CI build servers, and run-time telemetry.

The MSR conference has several tracks: the technical track for the latest research, experiences, and innovations; the data showcase to highlight datasets; the education track; and a mining challenge on the Software Heritage Graph Dataset, a very large dataset containing the development history of publicly available software. In addition, to increase transparency and replicability of empirical research, the MSR conference is piloting a track for registered reports.

Submission deadlines: Technical Papers (abstracts: January 9, papers: January 16, 2020); Mining Challenge, Data Showcase (abstracts: January 30, papers: February 6, 2020); Registered Reports (January 31, 2020).

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

AI Policy Matters

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

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.

National AI Strategy

The National Artificial Intelligence Research and Development Strategic Plan, an update of the report by the Select Committee on Artificial Intelligence of The National Science & Technology Council, was released in June, 2019, and the President's Executive Order 13859 Maintaining American Leadership in Artificial Intelligence was released on February 11. The Computing Community Consortium (CCC) recently released the AI Roadmap Website, and an interesting industry response is "Intel Gets Specific on a National Strategy for AI, How to Propel the US into a Sustainable Leadership Position on the Global Artificial Intelligence Stage" by Naveen Rao and David Hoffman. Excerpts follow and the accompanying links provide the details:

"AI is more than a matter of making good technology; it is also a matter of making good policy. And that's what a robust national AI strategy will do: continue to unlock the potential of AI, prepare for AI's many ramifications, and keep the U.S. among leading AI countries. At least 20 other countries have published, and often funded, their national AI strategies. Last month, the administration signaled its commitment to U.S. leadership in AI by issuing an executive order to launch the American AI Initiative, focusing federal government resources to develop AI. Now it's time to take the next step and bring industry and government together to develop a fully realized U.S. national strategy to continue leading AI innovation . . . to

sustain leadership and effectively manage the broad social implications of AI, the U.S. needs coordination across government, academia, industry and civil society. This challenge is too big for silos, and it requires that technologists and policymakers work together and understand each other's worlds." Their call to action was released in May 2018.

Four Key Pillars

"Our recommendation for a national AI strategy lays out four key responsibilities for government. Within each of these areas we propose actionable steps. We provide some highlights here, and we encourage you to read the full white paper or scan the shorter fact sheet.

- Sustainable and funded government AI research and development can help to advance the capabilities of AI in areas such as healthcare, cybersecurity, national security and education, but there need to be clear ethical guidelines.

- Create new employment opportunities and protect people's welfare given that AI has the potential to automate certain work activities.

- Liberate and share data responsibly, as the more data that is available, the more "intelligent" an AI system can become. But we need guardrails.

- Remove barriers and create a legal and policy environment that supports AI so that the responsible development and use of AI is not inadvertently derailed.

Work Transitions

AI and other automation technologies have great promise for benefitting society and enhancing productivity, but appropriate policies by companies and governments are needed to help manage workforce transitions and make them as smooth as possible. The McKinsey Global Institute report AI, automation, and the future of work: Ten things to solve for states that "There is work for everyone today and

there will be work for everyone tomorrow, even in a future with automation. Yet that work will be different, requiring new skills, and a far greater adaptability of the workforce than we have seen. Training and retraining both mid-career workers and new generations for the coming challenges will be an imperative. Government, private-sector leaders, and innovators all need to work together to better coordinate public and private initiatives, including creating the right incentives to invest more in human capital. The future with automation and AI will be challenging, but a much richer one if we harness the technologies with aplomb and mitigate the negative effects." They list likely actionable and scalable solutions in several key areas:

1. Ensuring robust economic and productivity growth

2. Fostering business dynamism

3. Evolving education systems and learning for a changed workplace

4. Investing in human capital

5. Improving labor-market dynamism

6. Redesigning work

7. Rethinking incomes

8. Rethinking transition support and safety nets for workers affected

9. Investing in drivers of demand for work

10. Embracing AI and automation safely

In redesigning work and rethinking incomes, we have the chance for bold ideas that improve the lives of workers and give them more interesting jobs that could provide meaning, purpose, and dignity. Some of the categories of new jobs that could replace old jobs are

1. Making, designing, and coding in AI, data science, and engineering occupations

2. Working in new types of non-AI jobs that are enhanced by AI, making unpleasant old jobs more palatable or providing new jobs that are more interesting; the gig economy and crowd sourcing ideas are examples that could provide creative employment options

3. Providing living wages for people to do things they love; for example, in the arts through dramatic funding increases for NEA and NEH programs. Grants to individual

artists and musicians, professional and amateur musical organizations, and informal arts education initiatives could enrich communities while providing income for millions of people. Policies to implement this idea could be one piece of the future-of-work puzzle and be much more preferable for the economy and society than allowing large-scale unemployment due to loss of work from automation.

Executive Order on The President's Council of Advisors on Science and Technology (PCAST)

President Trump issued an executive order re-establishing the President's Council of Advisors on Science and Technology (PCAST), an advisory body that consists of science and technology leaders from the private and academic sectors. PCAST is to be chaired by Kelvin Droegemeier, director of the Office of Science and Technology Policy, and Edward McGinnis, formerly with DOE, is to serve as the executive director. The majority of the 16 members are from key industry sectors. The executive order says that the council is expected to address "strengthening American leadership in science and technology, building the Workforce of the Future, and supporting foundational research and development across the country." For more information, see this Inside Higher Education article.

Please join our discussions at the SIGAI Policy Blog.

Larry Medsker is a Research Professor and was founding director of the Data Science graduate program at The George Washington University. He is 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 the Public Policy Officer for the ACM SIGAI.

Artificial Intelligence for All Using the R Programming Language

Sharad Sinha (Indian Institute of Technology (IIT) Goa; sharad@iitgoa.ac.in)
Clint P. George (Indian Institute of Technology (IIT) Goa; clint@iitgoa.ac.in)
DOI: 10.1145/3375637.3375642

Introduction

This article describes a public event, specifically a workshop, conducted by the authors to create awareness about Artificial Intelligence among high school students in the town of Ponda in the state of Goa, India. This event was funded by the ACM SIGAI AI Activities Fund and hence this article describes the event, its impact and follow up actions that the authors are engaged in.

The workshop was titled **"Artificial Intelligence for All using R Programming Language"** and was conducted on August 31, 2019. A total of forty (40) high school students belonging to Class XI in six schools in Ponda attended the workshop alongwith their teachers. The workshop was split into morning and afternoon sessions. The workshop was inaugurated by Prof. B K Mishra, Director of IIT Goa where he emphasized the need for early awareness among high school students about latest developments in science and technology, especially disruptive technologies like Artificial Intelligence. The morning session included one talk each by the authors and the afternoon session included hands-on exercises for students and their teachers to learn about a basic practical way of carrying out artificial intelligence related experiments using computers.

The six schools which attended the workshop were:

- *The Kamakshi Education Society's Higher Secondary School*, Curti- Ponda

- *G.V.M.'s S.N.J.A. Higher Secondary School*, Farmagudi, Ponda

- *S.S.Samiti's S.A.M.D. Higher Secondary School of Science*, Kavlem, Ponda

- *G.V.M.'s Adarsh Higher Secondary School, Dr. Dada Vidya Sankul*, Ponda

- *Swami Vivekanand Vidyaprasarak Mandal's Higher Secondary School*, Bori, Ponda

- *P.E.S. S.R.S.N. Higher Secondary School*, Farmagudi, Ponda

The authors had prepared an **"Artificial Intelligence (AI) Storybook"**, copies of which were made available to the participants. All participants were also given certificates of participation.

Forenoon Session: Technical Talks

Both the authors conducted a technical talk each. In the first such talk, the students were introduced to the very basic idea of artificial intelligence (AI). This session started with newspaper reports about AI related initiatives in India and the world, with specific examples from the fields of agriculture (smartphone based AI app for pest control) and smart traffic lights. Thereafter, the students were introduced to the basic idea of intelligence with a rudimentary classification that the first author came up with: analytical intelligence, emotional intelligence, moral and social intelligence. Video examples of a robot learning to flip pancakes and a car that parks itself were shared with the participants to excite them about the real world possibilities of AI. Figure 1 shows the participants before the beginning of the technical talks.

Figure 1: Students, teachers and authors during technical talks

The history of the development of machines

is exciting. Hence, the historical development from the six classical machines (wedge, ramp, wheel, lever, screwdriver and pulley) to modern machines that can play games like Go and defeat human players, was traced. This helped students appreciate how the science and technology behind modern artificial intelligence based machines has developed because of the inventions of programming languages, silicon transistors, integrated circuits and developments in mathematics of learning machines. Considering that the participants were primarily high school students, the technical interactive session was planned based on inputs from their teachers regarding their level of understanding of computers, programming, science and technology. The students were also introduced to the role played in artificial intelligence by the voluminous amount of data being generated these days.

The second technical talk was delivered by the second author. He began with a very simple example of recognition of basic geometrical shapes like circle, square etc. This was followed by a very gentle introduction to supervised and unsupervised learning, with a greater emphasis on supervised learning. The Iris dataset (Dua & Graff, 2017) was used to briefly introduce them to the task of automatically classifying iris flowers.

Afternoon Session: Hands-on Exercises

The afternoon session was conducted by both the authors with the help of their doctoral students. In this session, the students were introduced to very basic programming using R language. They were then introduced to the concept of datasets and their importance. The participants learned how to look for and download datasets. With the example of pre-downloaded iris dataset, they were taught step-by-step how to write code in R language to use the dataset to create a simple model to classify iris flowers. The concept of features of an object was introduced using this dataset - highlighting that the sepal and the petal lengths can be considered as object features which can help in clustering of iris flowers of similar dimensions. The hands-on exercises were conducted for 2.5 hours with a small break in between. Since the participants

were just learning to code in their school curriculum, a good amount of time was spent in making them familiar with the basics of the R programming language as relevant to the hands-on exercises. Figure 2 shows the participants during the hands-on session.
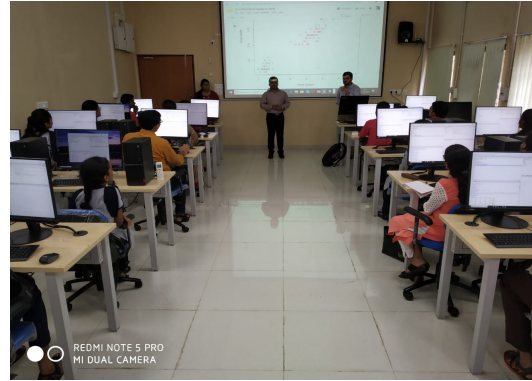

Figure 2: Students, teachers and authors during hands-on lab session

Artificial Intelligence (AI) Storybook

The authors had prepared an AI storybook that has two parts: the first part is a conversation between a boy and a girl on what is artificial intelligence. Written in a colloquial style, it was aimed at generating interest in the readers. The second part introduced supervised learning and k-nearest neighbor classification. Minimal mathematics was used to convey the gist of these two concepts. However, this part was more technical than the first part.

The authors were motivated with the idea of making the storybook accessible to multiple people in a house: students, their parents and perhaps grandparents. Hence, the first part was a conversation style discussion. The first part was also translated into Hindi, again with the help of Google translate followed by manual copyediting, to make this part accessible to grandparents many of who may not speak or read English fluently. Hindi is a language that is widely understood in Goa and in India. The storybook is available from th websites of the authors for download by anyone in the world. Figure 3 shows the participants with the AI Storybook.

Feedback and Closing Session

At the end of the hands-on exercises, all the participants were given participation certifi-

Figure 3: Students and teachers with the AI Storybook

cates and printed copies of the AI storybook. The students and the teachers very much appreciated the workshop and mentioned that they felt excited to learn more about this topic. A few of the schools also invited the authors to interact with their students on this subject. All of them expressed a desire to make the soft copy of the storybook available to them so that they could share with all the students in their schools. Figure 4 shows one of the participant groups with the authors.

Figure 4: A student group with the authors during closing session

Post Workshop Follow up & Media Coverage

The authors have already visited S.S.Samiti's S.A.M.D. Higher Secondary School of Science, Kavlem on 26 September 2019 to interact with students from classes XI and XII on the topic of artificial intelligence. During this interaction, the authors also emphasized on learning basic mathematics and programming

and answered questions on careers and education in computer science and engineering. Figure 5 shows the authors with the students during the interaction session.

Figure 5: Authors with students and teachers during the interaction at S.S.Samiti's S.A.M.D. Higher Secondary School of Science, Kavlem

The workshop was also highlighted as a major news item in the Times of India, a prominent English newspaper in India. The news item is accessible at: http://tiny.cc/r96eez

References

Dua, D., & Graff, C. (2017). *UCI machine learning repository.* Retrieved from http://archive.ics.uci.edu/ml

Sharad Sinha is an Assistant Professor in the School of Mathematics and Computer Science at IIT Goa. He received PhD degree (2014) in Computer Engineering from NTU, Singapore and B.Tech degree (2007) in Electronics and Communicaiton Engineering from CUSAT, India. He is an Associate Editor of ACM Ubiquity and IEEE Potentials magazines and is very active in engaging with schools and educational institutions in promoting computer science and engineering education. His research and teaching interests are in computer architecture, computing system design for machine and deep learning, reconfigurable computing, computer security and assistive technologies. He has received a **Best Paper Award** at **ICCAD 2017** and a **Best Paper Award nomination** at **CASES 2018** for his research work.

Clint George is an Assistant Professor in the School of Mathematics and Computer Science at IIT Goa. He was previously a postdoc at the Informatics Institute and Department of Statistics, University of Florida with Dr. George Michailidis. He received Ph.D. and M.S. degrees from the Department of Computer and Information Science and Engineering, University of Florida and completed B. Tech. in Computer Science and Engineering from the College of Engineering, Trivandrum, India. His teaching and research interests are in the areas of machine learning and applied statistics, including topic models, empirical Bayes methods, Markov Chain Monte Carlo (MMC) methods etc.

Cosmology of Artificial Intelligence Project: Libraries, Makerspaces, Community and AI literacy

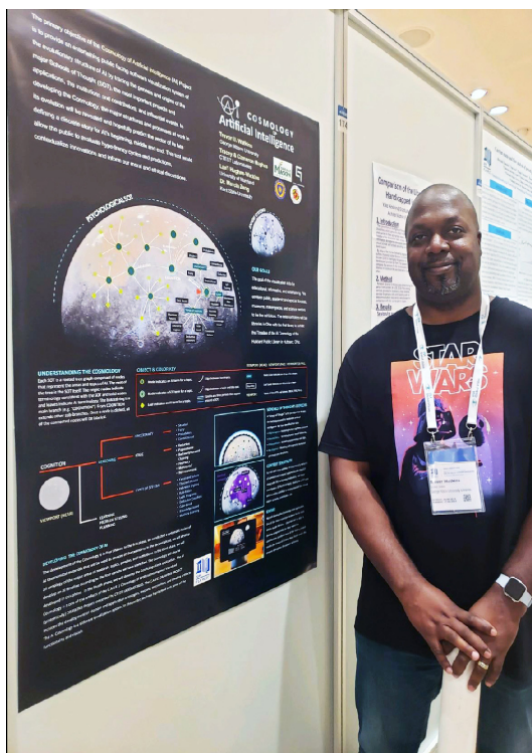**Trevor Watkins** (George Mason University; twatkin8@gmu.edu)

Figure 1: Trevor Watkins at Poster Session

Introduction

Artificial Intelligence (AI) has transcended beyond buzzwords, keywords or trends. The ubiquity of AI is so profound that it has managed to seep into popular culture. It breeds throughout social media platforms. It dominates the airwaves. It is impossible to watch television without mention of the acronym, the word, and how everyone is using it. However, if we look a little closer, take a deep dive into the AI pool where everyone seems to be swimming, we began to learn that AI has been misinterpreted, misrepresented, and incorrectly defined. "In the Laws of Thought and Thinking Machines," Hughes and Hughes (2019) talk about how uncovering the truth of the definition, meaning, applications, and im-

plications of Artificial Intelligence would be a noteworthy goal. Today we are pursuing that goal.

The primary objective of the Cosmology of Artificial Intelligence Project (CAIP) is to develop and provide a friendly outward facing visualization of the evolutionary structure of the field of Artificial Intelligence by tracing the genesis and origins of the major Schools of Thought (SOT), their major influences, and its interactions with the subject matter, the most important institutions, people, and techniques. The project aims to elucidate the cosmology in the form of a visualization map of a solar system used as a reference map. The graphs of the major SOTs, its topics, methods, and key contributors will be visually represented as astronomical objects in a solar system (e.g., planets, moons, rings, etc.). Historical events in AI and other fields of science that has had significant influence will be visualized as meteors, asteroids, comets or eclipses. The goal of the visualization is to be educational, informative, and entertaining. The visualization will be a 3D interactive visualization map kiosk where users will be able to explore the driving paradigms throughout the history of AI from its origins to the present. Typically, the term 'Cosmology' is used in conjunction with the study of the Universe. One convenient definition for Cosmology is the study, origin, evolution and the eventual fate of the Universe". Any valid AI Cosmology would force us to define its origin, uncover the goals and the meaning of the term "Artificial Intelligence" within the context of that history. In developing the Cosmology, the major structures and processes at work in its evolution would be revealed and hopefully predict the vector of its fate revealing a discrete story for AI's beginning, middle, and end. This tool would allow the public and committed others in the field to evaluate hype-frenzy cycles and predictions, contextualize innovations and inform our moral and ethical discussions. Much of the content for the AI cosmology will be historical in nature. We rely on

materials from digital libraries such as seminal books, papers, journal articles, interviews, oral histories, taxonomies, etc., to discern and identify the concepts, topics, contributors and projects and their relationships. The development of the Cosmology is in 4 phases:

- Phase 1: Review of AI Taxonomical Materials and Development of Visual Taxonomy
- Phase 2: Development of an Ontology of the Major Areas of Research, Topics, Peoples and Institutions
- Phase 3: Development of the AI Cosmology Part 1: Development of Cosmology Graphs Part 2: Mapping to Reference Map
- Phase 4: Development of the Cosmology Kiosk Application

The project will produce six artifacts:

- A 2D Taxonomical Visualization (Phase 1)
- A Digital Ontology using ISO 24707 format (Phase 2)
- A 2D version of the Cosmological map appropriate for a high-resolution printed poster (Phase 4)
- A kiosk application for libraries (public and academic), makerspaces, universities, museums, and science center exhibits (Phase 4)
- An AI literacy Framework
- An AI Curation Framework

The CAIP project team consists of Trevor Watkins from George Mason University (information literacy, AI literacy, metadata, systems and software engineering), Cameron and Tracey Hughes from Ctest Laboratories (software and visualization epistemology, data visualization, and AI historians), University Archivist Lae'l Hughes-Watkins from the University of Maryland College Park (information literacy, information appraisal, curation) and Dr. Marcia Zeng from Kent State University's I-School (knowledge organization and representation, thesaurus, taxonomy and ontology, and linked data and metadata).

We are now in phase 2 of the project.In phase 1, we produced a digital visualization of a complete taxonomy of the major research areas of Artificial Intelligence in addition to a timeline. This digital representation will soon be available for display as an exhibit in libraries, makerspaces, science centers, high schools, and universities. From the digital representation, a full color 42" x 90" visual taxonomy poster mounted under 1/8" (14 mm) thick acrylic glass will be generated. We have also created a full color AI timeline poster as well. The Hubbard Public Library in Hubbard, Ohio, and the Oakhill Makerspace in Youngstown, Ohio are the first public library and community makerspace to agree to serve as the first test site for this project. George Mason University Libraries will be the first academic library to display the digital representation. The digital file containing the visualization will be made available to ACM and IEEE Computer Society digital libraries, as well as a general download for the public from Ctest labs next year. This project is especially important because of recent national plans to invest in AI.
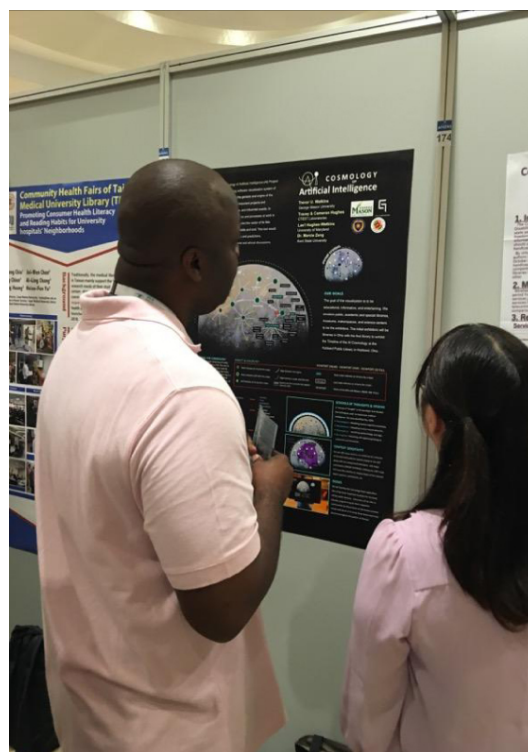

Figure 2: Poster Session

National AI Strategy

In an "AI Matters" blog post on September 9, 2019 titled "National AI Strategy," Larry Medsker (2019) provides references to the The National Artificial Intelligence Research

and Development Strategic Plan, the Computing Community Consortium's (CCC) Artificial Intelligence Roadmap Report, and Intel Corporation's recommendation for a National AI strategy, a white paper they released in March of 2019. Each of these resources respectively document and highlight strategies, recommendations, findings, and best practices regarding the future of the field of Artificial Intelligence.

In the National Artificial Intelligence Research and Development Strategic Plan (2019), the United States Government outlined eight strategic priorities:

- Make long-term investments in AI research.
- Develop effective methods for human-AI collaboration.
- Understand and address the ethical, legal, and societal implications of AI.
- Ensure the safety and security of AI systems.
- Develop shared public datasets and environments for AI training and testing.
- Measure and evaluate AI technologies through standards and benchmarks.
- Better understand the national AI Research and Development workforce needs.
- Expand public-private partnerships to accelerate advances in AI.

Intel Corporation described four responsibilities of governments initiating an AI strategy:

- Sustainable and funded government AI research and development
- Create new employment opportunities and protect people's welfare
- Liberate and share data responsibly
- Remove barriers and create a legal and policy environment that supports AI.

CCC (2019)has three core recommendations:

- Create and Operate a National AI Infrastructure
- Re-Conceptualize and Train an All-Encompassing Workforce
- Core Programs for AI Research (Increased Funding)

CCC's report was the only report that discussed the importance of creating an AI curricula, more specifically, "AI curricula guidelines should be developed to start at an early age to encourage interest in AI, understanding of the associated issues and implications, and curiosity to pursue careers in the field." We believe that AI literacy should be at the core of any AI curricula design. AI Literacy should be at the heart of any national AI strategy. Moreover, a national AI strategy should include an initiative to fund collaborations between libraries, school systems, makerspaces, etc. to serve as AI teaching and learning centers. The artifacts produced from CAIP, which includes the creation of an AI literacy framework, will serve as great tools for this type of initiative.

IFLA WLIC 2019 Greece

During the The International Federation of Library Associations and Institutions (IFLA) World Library and Information Congress (WLIC) 2019 held in Athens, Greece, we presented the Cosmology of AI project during a two-day poster session (See figure 1). In two days we had an accurate count of 367 visitors to our poster.
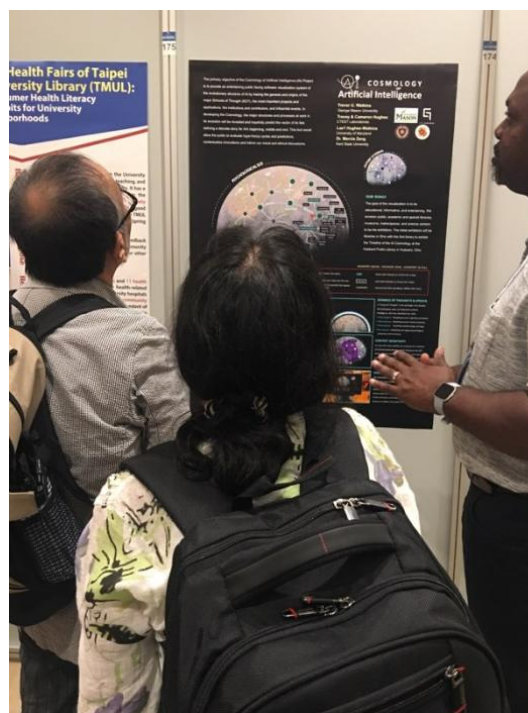

Figure 3: Poster Session Day 1

Some frequently asked questions and comments we received:

- When will this be available? Our library would be interested in having this.
- How much will it cost the library to have this?
- How will libraries use this as a tool?
- Why should libraries be concerned about AI?
- Will librarians be able to develop programming with this tool?
- Here is my contact information, let me know when this is available.
- This project may be too complex for the general public to understand.
- In what way do you think teens and younger children would benefit from this?
- We have a makerspace in our library, will there be different versions of this?
- Will it only be available in English?

The poster generated a lot interest from international public and academic libraries at the conference. Due to the feedback and questions we received, we plan to work with librarians to help develop content for AI specific events and programming, AI learning outcomes, and make it available in several languages.

References

Computing Community Consortium. (2019, August) Artificial Intelligence Roadmap. Retrieved From https://cra.org/ccc/visioning/visioning-activities/2018-activities/artificial-intelligence-roadmap/

Hughes, C., & Hughes, T. (2019). The laws of thought and thinking machines. AI Matters 5, 1 DOI: https://doi.org/10.1145/3320254.3320263

Medsker, L. (2019, September 9). National AI Strategy [Blog post]. Retrieved from https://sigai.acm.org/aimatters/blog/2019/09/09/national-ai-strategy/

WATKINS, T.U., HUGHES, T., HUGHES, C., HUGHES-WATKINS, L., & ZENG, M. (2019) Cosmology of Artificial Intelligence Project. Paper Presented at IFLA WLIC 2019 - Athens, Greece - Libraries: dialogue for change in Session 113c - IFLA Poster Session.

The White House, Executive Office of the President of the United States. (2019). National Artificial Intelligence Research and Development Strategic Plan. Retrieved from https://www.whitehouse.gov/wp-content/uploads/2019/06/ National-AI-Research-and-Development-Strategic-Plan-2019-Update-June-2019.pdf

Trevor Watkins is the Teaching and Outreach Librarian on the Teaching and Learning Team (Gateway Library) at George Mason University. He previously served five years as the STEM Librarian at Kent State University Libraries. He has over eight years of experience in Library IT, having previously served as a Systems Engineer and a Systems Librarian in public libraries in Ohio. He consults with academic and public libraries, makerspaces, and researchers on issues concerning privacy, information security, penetration testing, forensics, vulnerability assessment, data management, data science, data security, software development, and knowledge acquisition and dissemination. His research interests include Artificial Intelligence (AI) Literacy, developing systems using autonomous software agents and bio-inspired algorithms, ethical hacking, and human AI interaction. He is currently completing work on a GNU/Linux operating system for libraries and museums called Black Squirrel. He is a professional member of ACM (SIGAI, SIGCSE), IEEE, ALA, and IFLA.

AI is multidisciplinary

Virginia Dignum (Umeå University and WASP-HS; virginia@cs.umu.se)

Introduction

The last few years have seen a huge growth in the capabilities and applications of Artificial Intelligence (AI) and autonomous systems. Hardly a day goes by without news about technological advances and the societal impact of the use of AI. AI is changing the way we work, live and solve challenges. For example, it can improve healthcare, protect elephants from poachers, and work out how broadband should be distributed.

There are large expectations about the potential of AI to help to solve many current problems and to support human well-being, but we are also witnessing an increasing concern are about the potential of AI to do wrong. In particular issues of biased decisions, increased inequality, breaches of privacy, but also how the future of work will look like, and how can we keep AI under control, just to mention a few.

Many questions are arising from the development and use of AI. These are fundamental questions about what is intelligence, what is a fair and just decision, how to balance between individual and collective interests, how to solve ethical dilemmas or how automation will impact the labor market, are questions technology alone cannot answer. Answers to these questions require interdisciplinary approaches.

AI is not some kind of magic over which we have no control. We people are the ones that make AI happen. At the same time, it is increasingly more difficult to define AI (Monett & Lewis, 2017). Depending on the focus and the context, AI can refer to a computational technology (e.g. deep learning), a field of scientific research, or an autonomous entity capable of sensing and acting (Theodorou & Dignum, 2019). However, it is important to realise that eventhough AI is in fact a piece of software that we people design, understanding and guiding the impact of AI in society requires more than understanding its technical

design.

All technological change has traditionally been accompanied by fundamental societal changes. We are now at the brink of yet another one. It is up to us to make this a change for the better, for all of humankind and for the environment. This challenge is too large and too important to be left to engineers alone. All fields of knowledge from humanities and social science to art and design are needed to better build, understand and use AI.

A multidisciplinary perspective

Artificial Intelligence (AI) offers the potential for augmentation and potential replacement of human tasks and activities within a wide range of applications. The current pace of change for AI innovation is high, requiring societal, institutional and technological adjustments, and new opportunities for continued innovation across different domains including business and management, government, public sector, and science and technology. In order to navigate this potential, explore opportunities and mediate challenges, it is essential to integrate humanities and social science into the conversation about law, economics, ethics and impact of AI and digital technology. Only together we can chart a way forward into a beneficial and trustworthy future with our increasingly algorithmic societal systems.

The rapidly growing capabilities and increasing presence of AI-based systems in our lives raise pressing questions about the impact, governance, ethics, and accountability of these technologies around the world (Dwivedi et al., 2019). How can decisions about when, what for and how should AI be applied and how to integrate the the variety of views and requirements from people who use, interact with, and are impacted by these technologies? How do we harness the potential of AI systems while ensuring that they do not exacerbate existing inequalities and biases, or even create new ones? These questions cannot be answered from a computer science or engi-

neering perspective alone. In fact, we can say that Artificial intelligence is not any longer an engineering discipline but requires a broad involvement of different disciplines and participants.

Education plays here an important role. However, most current AI and robotics curricula worldwide deliver engineers with a too- narrow task view. The wide impact of AI on society requires a broadening of engineering education to include (Dignum, 2020):

1. analysis of the distributed nature of AI applications as these integrate socio-sociotechnical systems, and the complexity of human-agent interaction;

2. reflection on the meaning and global effect of the autonomous, emergent, decentralized, self-organizsing character of distributed learning entities and how they operate;

3. incremental design and development frameworks, and the unforeseen positive and negative influence of individual decisions at a system level, and as well as how these impact human rights, democracy, and education;

4. the consequences of inclusion and diversity in design, and how these inform processes and results;

5. understanding of governance and normative issues,: not only in terms of competences and responsibilities, but also in the case of views on health, safety, risks, explanations, and accountability; and,

6. the underlying societal, legal, and economic models of socio-sociotechnical systems.

Broadening AI curricula is possibly also a way to attract a more diverse student population. When AI curricula are known to be trans-disciplinary, it can be expected that female students, who traditionally choose humanities and social subjects over engineering ones, may be motivated to choose AI. In parallel, humanities and social sciences curricula also need to include subjects on the theory and practice of AI. For example, law curricula need to prepare law experts on how to address legal and regulatory issues around AI.

A Diversity Perspective

Impact of AI is now a global topic of discussion in academic and policy circles. At least 84 public–private initiatives have produced statements describing high-level principles, values and other tenets to guide the ethical development, deployment and governance of AI (Mittelstadt, 2019; Jobin, Ienca, & Vayena, 2019). Governance is necessary for the reduction of incidents, to ensure trust, and for society's long-term stability through the use of well-established tools and design practices. Well-designed regulations do not eliminate innovation but instead enhance it through the development and promotion of both socio-legal and technical means to enforce compliance [1]. Moreover, policy is needed to determine human responsibility in the development and deployment of intelligent systems, filling the gap that emerges from the increased automation of decision. Further, the ultimate aim of regulation is to ensure well-being for all in a sustainable world so it should can guide responsible research, development and use of AI (Theodorou & Dignum, 2019).

Besides disciplinary diversity, it is also important to consider cultural diversity, which includes factors such as education, religion, language. Artificial intelligence is increasingly pervasive and applied across cultures and geographic regions. Failure to understand cultural diversity impacts negatively the universal right to access to the advantages that the technology brings about. In an increasingly connected AI world, incentives and regulations can support awareness and commitment to a diverse perspective ensuring that AI applications are truly adaptable to a diverse cultural space, and thus enabling access to all. Or, as described in (Floridi et al., 2018) "Debates about technological challenges may lag behind the actual technical progress, but if they are strategically informed by a diverse, multistakeholder group, they may steer and support technological innovation in the right direction. Ethics should help seize opportunities and cope with challenges, not only describe them. It is essential in this respect that diversity infuses the design and development of AI, in terms of gender, class, ethnicity, discipline and other pertinent dimensions, in order to increase inclusivity, toleration, and the richness of ideas and perspectives."

Accountability is an important dimension of decision-making and therefore an essencial element for responsible and trustworthy AI. Inclusion, diversity and fairness are crucial to ensure that the impact of AI on individual and society is aligned with human rights and social values, and to analyse the nature and the role of biases that emerge from theoretical or empirical models that underpin AI algorithms and the interventions driven by such algorithms. While, the biases emerging from the theoretical and empirical models also affect human-controlled educational systems and interventions (e.g. hindsight and unconscious biases), the key mitigating difference between AI and human decision-making is that human decisions involve individual flexibility, context-relevant judgements, empathy, as well as complex moral judgements, missing from AI (Porayska-Pomsta & Rajendran, 2019). Again here, a multidisciplinary approach is essential to mitigate risks and increase benefits of AI applications.

AI, Humanities and Society: The Swedish approach

In the past years, the Swedish Wallenberg Foundations (https://www.wallenberg.org/en) have made significant investments in research on Ai, autonomous systems and software through the program WASP (http://wasp-sweden.org/). There is no doubt these technologies will have major effects on our society. Realising that fundamental research in AI must go hand in hand with a deep understanding of its societal, ethical, legal, economical and cultural impact, the Wallenberg Foundations are now investing over SEK 660 million in humanities and social sciences research, through the WASP-HS program (Wallenberg AI, Autonomous Systems and Software Program - Humanities and Society: wasp-hs.org). The WASP-HS program will analyze the impact of AI and autonomous systems in society, study the consequences of technology transition, and how these insights can contribute to the design and development of trustworthy and responsible systems. The WASP-HS program therefore extends and complements technological research and advances on AI with a strong investment on research in social science and humanities.

The WASP-HS program is planned to run 2019 – 2028 and will form an independent and parallel program to WASP, while maintaining a close dialogue with the WASP program. WASP-HS includes the following components:

- A research program aiming at forming an interdisciplinary community across the funded projects
- A national graduate school
- Recruitment of junior and visiting faculty across disciplines
- International partnerships and activities.

This multidisciplinary approach will advance our understanding of the challenges and impact of intelligent and autonomous technology, as well as contributing to the development of theory and practice of human and societal aspects of AI and autonomous systems. WASP-HS research tackles the challenges and impact of upcoming technology shifts as well as contributing to the development of theory and practice of human and societal aspects of AI and autonomous systems, and in particular, focus on potential ethical, economic, labor market, social, cultural and legal aspects of technological transition.

Conclusions

Many questions are arising from the development and use of AI. These are fundamental questions about what is intelligence, what is a fair and just decision, how to balance between individual and collective interests, how to solve ethical dilemmas or how automation will impact the labor market, are questions technology alone cannot answer. Answers to these questions require interdisciplinary approaches.

That is why the research program WASP-HS is of crucial importance for the future of AI, and in fact for the future of all of us. The aim of the WASP-HS program is to strengthen research and competence on the challenges and impact of AI, autonomous systems and software in the humanities and social sciences and to build a ground to utilise this expertise to inform society and industry on the design of socially aligned systems and on consequences for industry, society and humanity. We look forward

to collaboration with similar programs across the world.

References

References

Dignum, V. (2020). Responsibility and Artificial Intelligence. In *The Oxford Handbook of Ethics of AI, editor=markus dubber and frank pasquale and sunit das.* OUP.

Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., ... Williams, M. D. (2019). Artificial intelligence (ai): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*. Retrieved from http://www.sciencedirect.com/science/article/pii/S026840121930917X doi: https://doi.org/10.1016/j.ijinfomgt.2019.08.002

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... others (2018). Ai4people—an ethical framework for a good ai society: opportunities, risks, principles, and recommendations. *Minds and Machines*, *28*(4), 689–707.

Jobin, A., Ienca, M., & Vayena, E. (2019). Artificial intelligence: the global landscape of ethics guidelines. *arXiv preprint arXiv:1906.11668*.

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical ai. *Nature Machine Intelligence*, 1–7.

Monett, D., & Lewis, C. W. (2017). Getting clarity by defining artificial intelligence—a survey. In *3rd conference on" philosophy and theory of artificial intelligence* (pp. 212–214).

Porayska-Pomsta, K., & Rajendran, G. (2019). Accountability in human and artificial decision-making as the basis for diversity and educational inclusion. Springer Nature.

Theodorou, A., & Dignum, V. (2019). Towards ethical and socio-legal governance in AI. *Nature Machine Intelligence*, *to appear*.

Virginia Dignum is Professor of Social and Ethical Artificial Intelligence at Umeå University, Sweden and associated with the TU Delft in the Netherlands. She has a PHD in Artificial Intelligence from Utrecht University. Currently she is the scientific director of WASP-HS, the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program - Humanities and Society. She is a Fellow of the European Artificial Intelligence Association (EURAI), a member of the European Commission High Level Expert Group on Artificial Intelligence, of the World Economic Forum's Global Artificial Intelligence Council, of the Executive Committee of the IEEE Initiative on Ethically Alligned Design, and a founding member of ALLAI-NL, the Dutch AI Alliance. She is the author of "Responsible Artificial Intelligence: developing and using AI in a responsible way" published by Springer in 2019.

AI Fun Matters

Adi Botea (Eaton; adibotea@eaton.com)

1 A	I	4	5	6	■	7	8	9	10	11
12					■	13				
14		R	E	■	S	E	A	R	C	H
16			■	17	18			■	19	
20		■	21			■	22			
23		24			■	25				
■		I	S	■	27			■	■	
28	29	30		■	31			32	33	34
35			■	36			■	37		
38		■	F		■	■	40 A			
41		42	U	■	43	44			I	
45			N	■	46					
47			■	■	48					

Across: 1) Delivered by train. **7)** Give a false impression. **12)** With no parts missing. **13)** Unlimited warranty? **14)** Join an ideology. **15)** Core AI technology. **16)** Evolve as the Pisa Tower. **17)** Rewarded for work. **19)** Girl mentioned in absence. **20)** Apple computer. **21)** Queen goddess in Mt. Olympus. **22)** Unpleasant look. **23)** Catching with a cruel, inhumane trap. **25)** Puts together a book or a thesis. **26)** Suppliers of a product or service. **28)** Living plagues. **31)** Passes... into the fourth dimension. **35)** ___ = Distance/Time. **36)** Naval military prison. **37)** Set an upper bound. **38)** Deborah ___, Scotish journalist. **39)** Competitive speed-wise. **40)** Two in a binary logarithm. **41)** Old computer bus design. **43)** Spanish airline. **45)** Biggest prime digits. **46)** Sport from Electronic Arts' FIFA series. **47)** Fashion movement. **48)** Adopters of a greedy approach.

Down: 1) Old kingdoms. **2)** Linked to a big South American mountain chain. **3)** Town that hosts Cornell University. **4)** Right to keep someone else's property. **5)** Make a mistake. **6)** Advances research. **7)** Kept to create off-

spring. **8)** Chronos' major accomplishment. **9)** Kim G. ___, Computer Science professor at the Aalborg University. **10)** Moved along slowly. **11)** Ethyl oxides. **13)** Continent where IJCAI 2019 was held. **18)** Worthy opponents to IBM's Debater. **21)** React like a scaredy cat. **22)** Programming language designed by John McCarthy. **24)** It sounds like right. **25)** Proudly present personal achievements. **27)** Supporter of select social groups. **28)** Marcel ___, majestic French novelist. **29)** Income generator. **30)** Pursue a dream with passion. **32)** Hard to find. **33)** More attractive facility-wise. **34)** Britney ___, American diva. **36)** Low-range instrument. **39)** Support financially. **40)** Christopher ___, AI professor at the University of Toronto. **42)** ___ Goertzel, AI scientist featured in documentary *Machine of Human Dreams*. **44)** Large constrictor snake.

Previous solution: RESTED, AWASH, ALCOVE, STELLA, CLONES, POTTER, EIRE, CRAM, ADA, ROE, BEEN, RIGS, STRAINS, AIRES, MATTER, BAYES, ALTERED, USES, GRIS, EVA, RHO, MATT, AGER, SOMBER, INLINE, TREATY, SOONER, SENDS, TREADS

References

Botea, A. (2007). Crossword Grid Composition with A Hierarchical CSP Encoding. In *The 6th CP Workshop ModRef*.

Adi Botea is an AI senior specialist at Eaton, Ireland. He holds a PhD from the Univ. of Alberta, Canada. Adi co-authored over 80 scientific papers. He has published crosswords in Romanian, in national-level venues, for almost 3 decades.