acm SIGAI

AI Matters

Annotated Table of Contents

Links

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

Join SIGAI

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

Also consider joining ACM.

Our mailing list is open to all.

Notice to Contributing Authors to SIG Newsletters

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

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

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

Submit to AI Matters!

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

AI Matters Editorial Board

Contact us: aimatters@sigai.acm.org

Contents Legend


Book Announcement


Ph.D. Dissertation Briefing


AI Education


Event Report


Hot Topics


Humor


AI Impact


AI News


Opinion


Paper Précis


Spotlight


Video or Image

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

Welcome to AI Matters 4(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)
DOI: 10.1145/3299758.3299759

Welcome to our new co-editor!

Welcome to the fourth issue of the fourth volume of the AI Matters Newsletter! Our biggest news is that we would like to extend a big welcome to our new co-editor! Iolanda Leite is an Assistant Professor at the School of Computer Science and Electrical Engineering at the KTH Royal Institute of Technology in Sweden. Iolanda is excited to join the AI Matters team and to help to continue to grow the newsletter for ACM SIGAI. She is also our guest for the interview this month so please go to our interview and learn more about our new co-editor!

Announcements

Joint AAAI/ACM SIGAI Doctoral Dissertation Award

The Special Interest Group on Artificial Intelligence of the Association for Computing Machinery (ACM SIGAI) and the Association for the Advancement of Artificial Intelligence (AAAI) are happy to announce that they have established the Joint AAAI/ACM SIGAI Doctoral Dissertation Award to recognize and encourage superior research and writing by doctoral candidates in artificial intelligence. This annual award is presented at the AAAI Conference on Artificial Intelligence in the form of a certificate and is accompanied by the option to present the dissertation at the AAAI conference as well as to submit one 6-page summary for both the AAAI proceedings and the newsletter of ACM SIGAI. Up to two Honorable Mentions may also be awarded, also with the option to present their dissertations at the AAAI conference as well as submit one 6-page summary for both the AAAI proceedings and the newsletter of ACM SIGAI. The award will be presented for the first time at the AAAI conference in 2020 for dissertations that have been successfully defended (but not necessarily finalized) between October 1, 2018 and September 30, 2019. Nominations are welcome from any country, but only English language versions will be accepted. Only one nomination may be submitted per Ph.D. granting institution, including large universities. Dissertations will be reviewed for relevance to artificial intelligence, technical depth and significance of the research contribution, potential impact on theory and practice, and quality of presentation. The details of the nomination process will be announced in early 2019.

ACM SIGAI Student Essay Contest

The ACM Special Interest Group on Artificial Intelligence (ACM SIGAI) supports the development and responsible application of Artificial Intelligence (AI) technologies. We therefore invite essay submissions to the 2018 ACM SIGAI Student Essay Contest, to be published in the ACM SIGAI newsletter *AI Matters*, addressing one or both of the following topic areas:

- What requirements, if any, should be imposed on AI systems and technology when interacting with humans who may or may not know that they are interacting with a machine?
- What requirements, if any, should be imposed on AI systems and technology when making decisions that directly affect humans?

Win one of several $500 monetary prizes or a Skype conversation with a leading AI researcher, including Joanna Bryson, Murray Campbell, Eric Horvitz, Peter Norvig, Iyad Rahwan, Francesca Rossi, or Toby Walsh.

The deadline for submissions is January 10th, 2019.

For details, please visit the blog announcement.

*New Conflict of Interest and Retractions Policies

The ACM Publications Board has approved new Conflict of Interest and Retractions policies. Both of these new policies are now posted on ACM's website at:

https://www.acm.org/publications/policies/conflict-of-interest

Issue overview

This issue opens with an interview from Marion Newmann of our new co-editor: Iolanda Leite. We also have a summary of recent AI conferences and events from Michael Rovatsos. Our educational column this issue provides a model AI assignment for teaching ethics in AI focused on autonomous vehicles.

Larry Medsker's policy column summarizes a variety of policy issues of interest to the SIGAI community including the creation of a new Roadmap for Artificial Intelligence, FTC hearings on AI, and recent discussions on how to design AI to be safe as well as issues surrounding legality of AI. He closes with a short discussion of privacy for AI Matters. ACM's agreement with EasyChair specifies that we collect age and gender data and we welcome any discussion on our proposed response to this data collection.

Finally, we close with two paper contributions. The first is a continuation of the discussion of an AI "cosmology" by our contributing editors Cameron Hughes and Tracey Hughes. They propose the idea of epochs of AI in their current article. The second paper discusses the future of AI and how it will affect jobs.

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 Computer Science and Electrical Engineering 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.

Call for Proposals: Artificial Intelligence Activities Fund

Sven Koenig (elected; ACM SIGAI Chair)
Sanmay Das (elected; ACM SIGAI Vice-Chair)
Rosemary Paradis (elected; ACM SIGAI Secretary/Treasurer)
Michael Rovatsos (appointed; ACM SIGAI Conference Coordination Officer)
Nicholas Mattei (appointed; ACM SIGAI AI and Society Officer)
DOI: 10.1145/3299758.3299760

ACM SIGAI invites funding proposals for artificial intelligence (AI) activities with a strong outreach component to either students, researchers, or practitioners not working on AI technologies or to the public in general.

The purpose of this call is to promote a better understanding of current AI technologies, including their strengths and limitations, as well as their promise for the future. Examples of fundable activities include (but are not limited to) AI technology exhibits or exhibitions, holding meetings with panels on AI technology (including on AI ethics) with expert speakers, creating podcasts or short films on AI technologies that are accessible to the public, and holding AI programming competitions. ACM SIGAI will look for evidence that the information presented by the activity will be of high quality, accurate, unbiased (for example, not influenced by company interests), and at the right level for the intended audience.

ACM SIGAI has set aside $10,000 to provide grants of up to $2,000 each, with priority given to a) proposals from ACM affiliated organizations other than conferences (such as ACM SIGAI chapters or ACM chapters), b) out-of-the-box ideas, c) new activities (rather than existing and recurring activities), d) activities with long-term impact, e) activities that reach many people, and f) activities co-funded by others. We prefer not to fund activities for which sufficient funding is already available from elsewhere or that result in profit for the organizers. Note that expert talks on AI technology can typically be arranged with financial support of the ACM Distinguished Speaker program (https://speakers.acm.org/) and then are not appropriate for funding via this call.

A proposal should contain the following information on at most 3 pages:

- a description of the activity (including when and where it will be held); a budget for the activity and the amount of funding requested, and whether other organizations have been or will be approached for funding (and, if so, for how much); an explanation of how the activity fits this call (including whether it is new or recurring, which audience it will benefit, and how large the audience is);
- a description of the organizers and other participants (such as speakers) involved in the activity (including their expertise and their affiliation with ACM SIGAI or ACM);
- a description of what will happen to the surplus in case there is, unexpectedly, one; and
- the name, affiliation, and contact details (including postal and email address, phone number, and URL) of the corresponding organizer.

Grantees are required to submit a report to ACM SIGAI following completion of their activities with details on how they utilized the funds and other information which might also be published in the ACM SIGAI newsletter “AI Matters”.

The deadline for submissions is 11:59pm on March 15, 2019 (UTC-12). Proposals should be submitted as pdf documents in any style at

https://easychair.org/conferences/?conf=sigaiaaf2019.

The funding decisions of ACM SIGAI are final and cannot be appealed. Some funding earmarked for this call might not be awarded at the discretion of ACM SIGAI, for example, in case the number of high-quality proposals is not sufficiently large. In case of questions, please first check the

ACM SIGAI blog for announcements and clarifications: https://sigai.acm.org/aimatters/blog/. Questions should be directed to Sven Koenig (skoenig@usc.edu).

AI Profiles: An Interview with Iolanda Leite

Marion Neumann (Washington University in St. Louis; m.neumann@wustl.edu)
DOI: 10.1145/3299758.3299761

Introduction

This column is the seventh in our series profiling senior AI researchers. This month we are happy to interview Iolanda Leite, Assistant Professor at the School of Computer Science and Electrical Engineering at the KTH Royal Institute of Technology in Sweden. This is a great opportunity to get to know Iolanda, the new AI Matters co-editor in chief. Welcome on board!

Figure 1: Iolanda Leite

Biography

Iolanda Leite is an Assistant Professor at the School of Computer Science and Electrical Engineering at the KTH Royal Institute of Technology in Sweden. She holds a PhD in Information Systems and Computer Engineering from IST, University of Lisbon. Prior to joining KTH, she was a Research Assistant at the Intelligent Agents and Synthetic Characters Group at INESC-ID Lisbon, a Postdoctoral Associate at the Yale Social Robotics Lab and an Associate Research Scientist at Disney Research Pittsburgh. Iolanda's 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.

Getting to Know Iolanda Leite

When and how did you become interested in CS and AI?

I became interested in CS at the age of 4 when the first computer arrived at our home. It is more difficult to establish a time to define my interest in AI. I was born in the 80s and have always been fascinated by toys that had some level of "intelligence" or "life-likeness" like the Tamagotchi or the Furby robots. During my Master's degree, I chose the Intelligent Systems specialization. That time was probably when I seriously considered a research career in this area.

What professional achievement are you most proud of?

Seeing my students accomplish great things on their own.

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

I always loved to work with children so maybe something related to child education.

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

As an early researcher I often had a hard time dealing with the rejection of papers, applications, etc. What I wish the "past me" could know is that if one keeps working hard, things will eventually work out well in the end. In other words, keeping faith in the system.

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

All of them! If I have to highlight one, we are working with elementary schools that have classes of newly arrived children in a project where we are using social robots to promote inclusion between newly arrived and local children. This is part of an early career fellowship awarded by the Jacobs Foundation.

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

I believe that AI can be used to complement teachers and provide personalized instruction to students of all ages and in a variety of topics. Robotic tutors can play an important role in education because the mere physical presence of a robot has shown to have a positive impact on how much information students can recall, for example when compared to a virtual agent displayed in a computer screen delivering the exact same content.

How can we make AI more diverse? Do you have a concrete idea on what we as (PhD) students, researchers, and educators in AI can do to increase diversity our field?

Something we can all do is to participate in outreach initiatives targeting groups underrepresented in AI to show them that there is space for them in the community. If we start bottom-up, in the long-term I am positive that our community will be more diverse at all levels and the bias in opportunities, recruiting, etc. will go away.

What was your most difficult professional decision and why?

Leaving my home country (Portugal) after finishing my PhD to continue my research career because I miss my family and friends, and also the good weather!

How do you balance being involved in so many different aspects of the AI community?

I love what I do and I currently don't have any hobbies :-)

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

If AI could fully address any of the Sustainable Development Goals established by the United Nations, it would be (more than) great. Although there are excellent research and funding initiatives in that direction, we are still not there yet.

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

One of my favorite ones recently was the Westworld TV Series because of the power relationships between the human and the robotic characters. I find it hard to believe that humans will treat robots the way they are treated in the series, but it makes me reflect on how our future interactions with technology that is becoming more personalized and "human-like" might look like.

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

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

Events

Michael Rovatsos (University of Edinburgh; mrovatso@inf.ed.ac.uk)
DOI: 10.1145/3299758.3299762

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 organizing an event and would be interested in cooperating with SIGAI, or if you have announcements relevant to SIGAI. For more information about conference support visit sigai.acm.org/activities/requesting_sponsorship.html.

21st International Conference on Entreprise Information Systems

Heraklion, Crete, Greece, May 3-5, 2019
www.iceis.org
The purpose of the 21st International Conference on Enterprise Information Systems (ICEIS) is to bring together researchers, engineers and practitioners interested in the advances and business applications of information systems. Six simultaneous tracks will be held, covering different aspects of Enterprise Information Systems Applications, including Enterprise Database Technology, Systems Integration, Artificial Intelligence, Decision Support Systems, Information Systems Analysis and Specification, Internet Computing, Electronic Commerce, Human Factors and Enterprise Architecture. ICEIS 2019 will be held in conjunction with ENASE 2019, GISTAM 2019 and FEMIB 2019.

2018 IEEE/WIC/ACM International Conference on Web Intelligence

Santiago, Chile, December 3-6, 2018
webintelligence2018.com
Web Intelligence (WI) aims to achieve a multi-disciplinary balance between research advances in the fields of collective intelligence, data science, human-centric computing, knowledge management, and network science. It is committed to addressing research that deepens the understanding of computational, logical, cognitive, physical as well as business and social foundations of the future Web, and enables the development and application of intelligent technologies. WI'18 features high-quality, original research papers and real-world applications in all theoretical and technological areas that make up the field of WI. WI'18 welcomes research, application as well as Industry/Demo-Track paper submissions. Tutorial, Workshop and Special-Session proposals and papers are also welcome.

14th ACM/IEEE International Conference on Human-Robot Interaction

Daegu, Korea, March 11-14, 2019
humanrobotinteraction.org/2019
HRI 2019 is the 14th annual conference for basic and applied human-robot interaction research. Researchers from across the world attend and submit their best work to HRI to exchange ideas about the latest theories, technology, data, and videos furthering the state-of- the-art in human-robot interaction. Each year, the HRI Conference highlights a particular area through a theme. The theme of HRI 2019 is "Collaborative HRI". The HRI conference is a highly selective annual international conference that aims to showcase the very best interdisciplinary and multidisciplinary research in human-robot interaction with roots in and broad participation from communities that include but not limited to robotics, human-computer interaction, human factors, artificial intelligence, engineering, and social and behavioral sciences.

11th International Joint Conference on Biomedical Engineering Systems and Technologies

Prague, Czech Republic, February 22-24, 2019
www.biostec.org
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 2019 will be held in conjunction with MODELSWARD 2019 and ICISSP 2019. Registration to BIOSTEC allows free access to the MODELSWARD and ICISSP conferences (as a non-speaker).

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

Conference Reports

Michael Rovatsos (University of Edinburgh; mrovatso@inf.ed.ac.uk)
DOI: 10.1145/3299758.3299763

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

The 5th international Workshop on Sensor-based Activity Recognition and Interaction (iWOAR 2018)

Berlin, Germany, September 20-21, 2018
www.iwoar.org/2018

This conference-like workshop was initiated and organized by the Fraunhofer IGD and the University of Rostock. It offers scientists, interested parties, and users in the area of sensor-based activity recognition and interaction the possibility to an exchange of experiences and a presentation of best-practice examples, as well as technical and scientific results. The workshop focuses on technologies for human activity recognition and interaction via inertial sensors (accelerometers, gyroscopes etc.) and their scientific applications. In this year's edition of iWOAR there were 28 submissions, and the program committee chairs decided to accept 15 papers, resulting in an overall acceptance rate of 54%. For the first time at iWOAR, we invited 5 Posters, which were also peer-reviewed. Although the majority of authors are based in Germany, there were also submissions from Austria, Denmark, Greece, India, Japan, New Zealand, Spain, Sweden, Switzerland, Turkey, and USA.

The 31st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE-2018)

Montreal, Canada, June 25–28, 2018
ieaaie2018.encs.concordia.ca

IEA/AIE-2018 was held at Concordia University in Montreal, Canada this year. As in past years, the conference was sponsored by the International Society of Applied Intelligence (ISAI), held in cooperation with AAAI, SIGAI/ACM, and several other international organizations. The conference provides a medium for presenting and exchanging new scientific research ideas and technological achievements accomplished by the international community in applied intelligence. IEA/AIE 2018 received 146 submissions from Canada along with 44 other countries. 64 papers were selected for oral presentation at the conference and included in the related proceedings published by Springer in the Lecture Notes in Artificial Intelligence (LNCS/LNAI) series. In addition to the papers accepted at the main track, 22 papers were selected (from the papers submitted directly to the special session chairs) for presentation at the following four special sessions and for inclusion in the proceedings: Artificial Intelligence, Law and Justice Track, Data Science, Privacy, and Security Track, Intelligent Systems Approaches in Information Extraction Track, and Internet of Things and Ubiquitous Computing and Big Data Track.

The 15th International Conference on Informatics in Control, Automation and Robotics (ICINCO 2018)

Porto, Portugal, July 29-31, 2018
www.icinco.org/?y=2018

ICINCO 2018 was held on Porto from 29th to 31st July, 2018, and was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC). ICINCO 2018 was also organized in cooperation with the ACM SIGAI, INNS, AAAI, Euromicro, Associação Portuguesa de Controlo Automático, euRobotics AISBL, Sociedade Portuguesa de Robótica, Sociedade Brasileira de Automática and the Robotics Society of Japan. It was also technically co-sponsored by International Federation of Automatic Control (IFAC) and IEEE Robotics & Automation Society. ICINCO received 209 paper submissions from 40 countries, 22.49% of which were published and presented as full papers. Four invited talks were delivered by internationally distinguished speakers, namely Keith Clark (Imperial College London), Christian Schlette (Univer-

sity of Southern Denmark), and Bruno Siciliano (University of Naples Federico II). Additionally, a “Best Paper Award”, a “Best Student Paper Award” and a “Best Poster Award” was conferred at the conference venue. More information can be found at www.icinco.org/PreviousAwards.aspx.

The 20th International Conference on Enterprise Information Systems (ICEIS 2018)

Funchal, Portugal, March 21-24, 2018
www.iceis.org/?y=2018
ICEIS 2018 (20th International Conference on Enterprise Information Systems) was held on Funchal, Madeira, Portugal from 21-24 March, 2018. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication (INSTICC). ICEIS 2018 was also organized in cooperation with the ACM SIGMIS, ACM SIGAI, ACM SIGCHI, AAAI, and SWIM (the IEICE Special Interest Group on Software Interprise Modelling).

ICEIS received 242 paper submissions from 45 countries. To evaluate each submission, a double blind paper review was performed by the Program Committee. After a stringent selection process, 18% of the papers were accepted for publication and presented as full papers. Four invited talks were delivered by internationally distinguished speakers: Alexander Brodsky (George Mason University), Plamen Angelov (Lancaster University), Salvatore Distefano (Universitá degli Studi di Messina), and David Aveiro (University of Madeira / Madeira-ITI).

The conference also awarded a “Best Paper Award”, a “Best Student Paper Award” and a “Best Poster Award”, which were conferred at the conference. More information can be found at: www.iceis.org/PreviousAwards.aspx

The 10th International Conference on Agents and Artificial Intelligence (ICAART 2018)

Funchal, Portugal, January 16-18, 2018
www.icaart.org/?y=2018
ICAART 2018 was held on Funchal, Madeira - Portugal, from January 16 to 18, 2018. It was sponsored by the Institute for Systems and Technologies of Information, Control and Communication with the Madeira Interactive Technologies Institute (MITI) as a local partner. The conference was also technically co-sponsored by IEEE Computational Intelligence Society. The conference was organized in cooperation with the ACM SIGAI, the Portuguese Association for Artificial Intelligence, Associazione Italiana per l’Intelligenza Artificiale, the IberoAmerican Society of Artificial Intelligence, European Society for Fuzzy Logic and Technology, the Asociación Española para la Inteligencia Artificial, the International Fuzzy Systems Association, European Association for Artifical Intelligence, and AAAI. ICAART received 161 paper submissions from 34 countries, of which 27.95% were accepted as full papers. Four invited talks were delivered by internationally distinguished speakers: Luc Steels (ICREA and UPF Barcelona), Virginia Dignum (Delft University of Technology), Eduard Hovy (CMU), and Luís Antunes, Universidade de Lisboa, Portugal. A “Best Paper Award”, a “Best Student Paper Award” and a “Best Poster Award” were conferred at the conference venue. More information can be found at www.icaart.org/PreviousAwards.aspx

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

AI Education Matters: A Modular Approach to AI Ethics Education

Heidi Furey (Manhattan College; hfurey01@manhattan.edu)
Fred Martin (University of Massachusetts Lowell; fred_martin@uml.edu)
DOI: 10.1145/3299758.3299764

Introduction

In this column, we introduce our Model AI Assignment, A Module on Ethical Thinking about Autonomous Vehicles in an AI Course, and more broadly introduce a conversation on ethics education in AI education.

Why Ethics in an AI Course?

Recognition of the need for ethics education in the engineering-related disciplines goes back at least a hundred years (Layton Jr, 1986), but it's only been since the 1990s that expectations for ethics education have been adopted by accreditation bodies (Stephan, 1999).

The use of artificial intelligence and machine learning has become increasingly pervasive in our society, and this is accompanied by greatly expanding interest in the ethical considerations regarding their use. A crucial event was ProPublica's 2016 report investigating the use of AI tools by judges in criminal courts for determining the nature and lengths of sentences for convicts (Angwin, Larson, Mattu, & Kirchner, 2016). The article made a compelling case that the commercial software (which is in widespread use and employs secret algorithms) is racially biased.

International standards bodies are also tackling issues related to the ethical use of AI techniques, and specifically, the use of autonomous vehicles. The USA published a National Artificial Intelligence Research and Development (NAIRD) Strategic Plan, and the EU published a report, "Recommendations to EU Commission on Civil Law Rules on Robotics" (Martin & Makoundou, 2017).

Broadly, there is deep concern about the increasingly wide-reaching societal impact of AI approaches. It is urgent that our students are ready to think through the implications of their work, and make ethical choices.

Why A Modular Approach?

Here we are presenting a modular approach: a way of incorporating an introductory lesson about ethics into a one-semester AI course. Our module introduces students to the objectivist framework, and opens a specific conversation about the ethics of self-driving cars.

The modular approach is easily integrated into a single course. If the course is popular (or required), it reaches a large portion (or all) of the student population. Most importantly, students can connect the specific AI ideas they are learning to their ethical implications.

The limitations of the modular approach are mostly related to its short duration—there is only so much that can be accomplished with only one week of instruction.

A modular approach would be complementary with a whole-semester course on ethical thinking. Doing both would be more effective than only one or the other.

The Model AI Assignment

In our Module on Ethical Thinking about Autonomous Vehicles in an AI Course, we shared a set of resources for faculty use:

- Two days of lecture, in-class exercises, and discussion, introducing Utilitarianism and the Trolley Problem.
- A requirement that the course final project paper includes a discussion of the ethical implications of the project idea.
- A question on the final examination assessing students' understanding of the Trolley Problem.

These materials are elaborated in (Furey & Martin, 2018); here, we add further reflections and share more resources for introducing ethics to AI students.

Reflections on Ethics Teaching

One challenge to incorporating ethics into an AI course is helping students recognize the structure of ethical problems and their solutions. The field of ethics is both complex and far-ranging. Because of this, it can be difficult to meet this challenge within the space of an AI course without getting too far a field of the standard course material. With this concern in mind, we selected a topic that could serve as an example both of an ethical problem and ethical problem solving: the ethics of algorithm development for autonomous vehicles.

In the module, students were first introduced to the connection between ethical algorithms for autonomous vehicles and a classic ethical dilemma: the Trolley Problem. Afterword, students were guided through a worksheet driven group exercise designed to foster discussion and debate on the topic. This discussion helped prepare students to evaluate the potential benefits and shortcomings of an intuitive solution to ethical problems—Utilitarianism. Finally, students were offered examples of how one might construct a solution to these problems and challenged to continue thinking about the issue.

Thoughts on Implementation

The Trolley Problem presents a seeming intractable ethical dilemma — one in which every solution comes with an ethical cost. One benefit of exposing students to a somewhat eccentric philosophical example is that students become familiar with a key ethical problem-solving tool — the use of "thought experiment" — highly idealized hypothetical cases designed to test ethical theories and to isolate relevant moral variables. Another benefit is that students come to understand the complex nature of ethical dilemmas, and to recognize that solutions to such problems are rarely straightforward and may perhaps be equally complex. Too often, students who are unfamiliar with ethical problem solving resist thinking about ethical dilemmas because the answers appear "unknowable." Here, it is useful to draw parallels between difficult moral questions and difficult technical questions, reminding them that that complexity does not necessarily equal intractability.

It is useful to give them some general background in the field of ethics. For instance, It is important for students to recognize that there are different sorts of questions that one might ask regarding human behavior, only some of fall under the domain of ethical inquiry. Ethicists, as opposed to psychologists or sociologist, are interested in what people ought morally do to apart from what they in fact do or why they do it.

Other Resources

Burton et al.'s "Ethical Considerations in Artificial Intelligence Courses" (2017) provides an in-depth introduction to this conversation, including case studies and resource links. These authors also contribute the idea of using science fiction to engage students in thinking through ethical considerations (Burton, Goldsmith, & Mattei, 2015).

For a short video introduction to some ethics issues in AI, see Atlantic Magazine's Moral Code: The Ethics of AI.

For a variety of AI ethics resources including research, symposia, workshops, and reports, visit the AI Now Institute at New York University.

The AI Ethics Lab is a virtual organization that "brings together researchers and practitioners from various disciplines to detect and solve issues related to ethical design in AI."

Several articles from popular media address the challenges of introducing AI ethics into the classroom. For instance, see (Tugend, 2018) and (Holmes, 2018).

To date, more research has been done in engineering ethics pedagogy than in AI ethics pedagogy. For this reason, it can be useful to look for resources in engineering ethics. For a problem-solving approach to ethics instruction, see (Whitbeck, 1996). For a brief overview of some of the major ethical theories along with a discussion benefits and difficulties of teaching ethical theory, see (Bouville, 2008). For an article on teaching higher-order ethical concepts in engineering, see (Haws, 2004). For an example of an alternative to stand-alone ethics instruction in engineering, see (Riley, Davis, Jackson, & Maciukenas, 2009).

References

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May). Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica*, *23*.

Bouville, M. (2008). On using ethical theories to teach engineering ethics. *Science and Engineering Ethics*, *14*(1), 111–120.

Burton, E., Goldsmith, J., Koenig, S., Kuipers, B., Mattei, N., & Walsh, T. (2017). Ethical considerations in artificial intelligence courses. *arXiv preprint arXiv:1701.07769*.

Burton, E., Goldsmith, J., & Mattei, N. (2015). Teaching AI ethics using science fiction. In *Aaai workshop: Ai and ethics.*

Furey, H., & Martin, F. (2018). Introducing ethical thinking about autonomous vehicles into an AI course. In *AAAI-2018.*

Haws, D. R. (2004). The importance of meta-ethics in engineering education. *Science and Engineering Ethics*, *10*(2), 204–210.

Holmes, W. (2018). The ethics of artificial intelligence in education. *University Business*.

Layton Jr, E. T. (1986). *The revolt of the engineers. social responsibility and the american engineering profession.* ERIC.

Martin, C. D., & Makoundou, T. T. (2017). Taking the high road ethics by design in AI. *ACM Inroads*, *8*(4), 35–37.

Riley, K., Davis, M., Jackson, A. C., & Maciukenas, J. (2009, March). "ethics in the details": Communicating engineering ethics via micro-insertion. *IEEE Transactions on Professional Communication*, *52*(1), 95-108. doi: 10.1109/TPC.2008.2012286

Stephan, K. D. (1999). A survey of ethics-related instruction in U.S. engineering programs. *Journal of Engineering Education*, *88*(4), 459-464.

Tugend, A. (2018). Colleges grapple with teaching the technology and ethics of A.I. *New York Times*.

Whitbeck, C. (1996). Ethics as design: Doing justice to moral problems. *Hastings Center Report*, *26*(3), 9–16.

Heidi Furey is an Assistant Professor of Philosophy and Director of the Ethics Center at at Manhattan College. Furey researches applied ethics and ethical pedagogy. Her current projects include work on the ethics of autonomous vehicles.

Fred Martin is Professor of Computer Science and Associate Dean for Student Success in the Kennedy College of Sciences at the University of Massachusetts Lowell. Martin develops novel technologies for CS and data science education, particularly for K–12.

AI Policy Matters

Larry Medsker (George Washington University; lrm@gwu.edu)
DOI: 10.1145/3299758.3299765

Abstract

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

Introduction

The SIGAI Public Policy goals are to

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

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

AI Appointments

The Computing Community Consortium (CCC) announced a new initiative to create a Roadmap for Artificial Intelligence. SIGAI's Yolanda Gil (University of Southern California and President-Elect of AAAI) will work with Bart Selman (Cornell University) to lead the effort. The initiative will support the U.S. Administration's efforts in this area and involve academic and industrial researchers to help map a course for needed research in AI. They will hold a series of workshops in 2018 and 2019 to produce the Roadmap by Spring of 2019.

Ed Felten, Princeton Professor of Computer Science and Public Affairs, has been confirmed by the U.S. Senate to be a member of the U.S. Privacy and Civil Liberties Oversight Board, a bipartisan agency within the executive branch. He will serve as a part-time member of the board while continuing his teaching and research at Princeton. The five-person board is charged with evaluating and advising on executive branch anti-terrorism measures with respect to privacy and civil liberties. It is a very important issue, Felten said. Federal agencies, in the course of doing national security work, have access to a lot of data about people and they do intercept data. Its important to make sure they are doing those things in the way they should and not overstepping. Felten added that the board has the authority to review programs that require secrecy. The public has limited visibility into some of these programs, Felten said. The board's job is to look out for the public interest.

FTC Hearings on AI

On November 13 and 14, the Federal Trade Commission (FTC) held the seventh hearing in its series of nine planned Hearings on Competition and Consumer Protection in the 21st Century. From the FTC: "The hearing will examine competition and consumer protection issues associated with the use of algorithms, artificial intelligence, and predictive analytics in business decisions and conduct. See detailed agenda. The record of that proceeding will be open until mid-February. To further its consideration of these issues, the agency seeks public comment on the questions, and it welcomes input on other related topics not specifically listed in the 25 questions." See the agenda and more information in the press release. Comments can be submitted online no later than February 15, 2019.

AI Safety By Design

Policymakers need to think of AI and Autonomous Systems (AI/AS) as always needing varying degrees of the human role (hybrid human/machine systems). Understanding the potential and limitations of combining technologies and humans is important for re-

alistic policymaking. A key element, along with accurate forecasts of the changes in technology, is the safety of AI/AS-Human products as discussed in the IEEE report *Ethically Aligned Design*, which is subtitled "A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems", and Ben Shneidermans excellent summary and comments on the report as well as the YouTube video of his Turing Institute Lecture on "Algorithmic Accountability: Design for Safety".

In Shneidermans proposal for a National Algorithms Safety Board, he writes What might help are traditional forms of independent oversight that use knowledgeable people who have powerful tools to anticipate, monitor, and retrospectively review operations of vital national services. The three forms of independent oversight that have been used in the past by industry and governmentsplanning oversight, continuous monitoring by knowledgeable review boards using advanced software, and a retrospective analysis of disastersprovide guidance for responsible technology leaders and concerned policy makers. Considering all three forms of oversight could lead to policies that prevent inadequate designs, biased outcomes, or criminal actions.

Efforts to provide safety by design include work at Google on Human-Centered Machine Learning and a general human-centered approach "that foregrounds responsible AI practices and products that work well for all people and contexts. These values of responsible and inclusive AI are at the core of the AutoML suite of machine learning products ". Further work is needed to systemize and enforce good practices in human-centered AI design and development, including algorithmic transparency and guidance for selection of unbiased data used in machine learning systems.

Legal AI

AI is impacting law and policy issues as both a tool and a subject area. Advances in AI provide tools for carrying out legal work in business and government, and the use of AI in all parts of society is creating new demands and challenges for the legal profession.

Lawyers and AI Tools: In a recent study, "20 top US corporate lawyers with decades of experience in corporate law and contract review were pitted against an AI. "Their task was to spot issues in five Non-Disclosure Agreements (NDAs), which are a contractual basis for most business deals." The LawGeex AI system attempted correct identification of basic legal principles in contracts The results suggest that AI systems can produce higher accuracy in shorter times compared to lawyers. As with other areas of AI applications, issues include trust in automation to make skilled legal decisions, safety in using AI systems, and impacts on the workforce of the future. For legal work, AI systems potentially reduce the time needed for high-volume and low-risk contracts and give lawyers more time to work on less mundane work. Policies should focus on automation where possible and safe, and the AI for legal work is another example of the need for collaborative roles for human and AI systems.

AI Impact on Litigation: Litigation is emerging in all parts of society related to the use of AI. Understanding the nature of adaptive AI systems can be crucial for fact-finders and difficult to explain to non-experts. Smart policymaking needs to make clear the liability issues and ethics in cases involving the use of AI technology. Artificial Intelligence and the Role of Expert Witnesses in AI Litigation by Dani Alexis Ryskamp, writing for The Expert Institute, discusses artificial intelligence in civil claims and the role of expert witnesses in elucidating the complexities of the technology in the context of litigation. "Over the past few decades, everything from motor vehicles to household appliances has become more complex and, in many cases, artificial intelligence only adds to that complexity. For end-users of AI products, determining what went wrong and whose negligence was responsible can be bafflingly complex. Experts retained in AI cases typically come from fields like computer or mechanical engineering, information systems, data analysis, robotics, and programming. They may specialize in questions surrounding hardware, software, 3D-printing, biomechanics, Bayesian logic, e-commerce, or other disciplines. The European Commission recently considered the question of whether to give legal status to certain robots. One of the issues weighed in the decision involved legal liability: if an AI-based robot or

system, acting autonomously, injures a person, who is liable?"

Pew Research Report

Pew Research Center recently released a report "Public Attitudes Toward Computer Algorithms", by Aaron Smith, on Americans concerns about fairness and effectiveness in making important decisions. The report says "This skepticism spans several dimensions. At a broad level, 58 percent of Americans feel that computer programs will always reflect some level of human bias although 40 percent think these programs can be designed in a way that is bias-free. And in various contexts, the public worries that these tools might violate privacy, fail to capture the nuance of complex situations, or simply put the people they are evaluating in an unfair situation. Public perceptions of algorithmic decision-making are also often highly contextual the survey presented respondents with four different scenarios in which computers make decisions by collecting and analyzing large quantities of public and private data. Each of these scenarios were based on real-world examples of algorithmic decision-making and included: a personal finance score used to offer consumers deals or discounts; a criminal risk assessment of people up for parole; an automated resume screening program for job applicants; and a computer-based analysis of job interviews. The survey also included questions about the content that users are exposed to on social media platforms as a way to gauge opinions of more consumer-facing algorithms."

EasyChair Policy Issue

We are revisiting an issue raised last June involving data privacy for SIGAI and ACM members using EasyChair to submit articles for publication, including the AI Matters Newsletter. As part of entering a new submission through EasyChair, the following message appears: "AI Matters, 2014-present, is an ACM conference. The age and gender fields are added by ACM. By providing the information requested, you will help ACM to better understand where it stands in terms of diversity to be able to focus on areas of improvement. It is mandatory for the submitting author (but you can select prefer not to submit) and it is desirable that you fill it out for all authors. This information will be deleted from EasyChair after the conference."

We recommend that to evaluate the likelihood of privacy protection, one should pay attention to the EasyChair Terms of Service, particularly Section 6 "Use of Personal Information" to assess better the level of risk if our members choose to enter personal information. Your Public Policy Officer recommends further discussion among the SIGAI membership about requesting changes in ACM policy. Please send your views on this topic to SIGAI officers and contribute comments to the policy blog.

Larry Medsker is Research Professor of Physics and was founding director of the Data Science graduate program at The George Washington University. Dr. Medsker is a former Dean of the Siena College School of Science, and Professor in Computer Science and in Physics. His research in AI includes work on artificial neural networks, hybrid intelligent systems, and the impacts of AI on society and policy.[a]. He is the Public Policy Officer for the ACM SIGAI.

[a] http://humac-web.org/

Epochs of an AI Cosmology

Cameron Hughes (Northeast Ohio ACM Chair; cameronhughes@acm.org)
Tracey Hughes (Northeast Ohio ACM Secretary; tracey.hughes@neoacmchapter.org)
DOI: 10.1145/3299758.3299868

Epochs and Hype Cycles

We began the discussion about the educational benefits to the general public of a standard AI Cosmology in Volume 4 Issue 2 of "AI Matters". But the mere notion of such a standard immediately raises lots of questions. In this column, we only focus on a very narrow set of those questions: First and foremost: Is there an AI Cosmology? If in fact the notion of cosmology applies to AI, is there only one valid AI Cosmology or are there competing valid cosmologies? If there is only one valid cosmology, what is its structure? How far can we take an AI Cosmology analogy? Like our Universe, did AI have the equivalent of a Big Bang? Is the field or structure of AI expanding or contracting? In an interview with Nikola Danaylov, podcaster of Singularity.FM, the late Marvin Minsky[1], one of the early AI pioneers, suggested that the evolution of AI in the 21st century maybe be contracting or even stalled! We posit that an AI Cosmology would help to educate the general public about our research and work, and its impact on society. Internally, it would help us to fix the objects, structure, and dynamics of the field. A standard cosmological model would help us all distinguish spurious AI mythos, pseudo science, hype, and fake news from the genuine article. A true cosmology of Artificial Intelligence would help us fix its origins, understand its trajectory, and predict its fate.

We are currently at the beginning of what we could consider the epoch of *Machine Learning*. Of course how the field will ultimately measure the epochs of AI remains an open question. But dividing the history of AI into epochs might help us better understand the cosmology of AI a little better. For instance, we could say that before the epoch of *Machine Learning* was the epoch of *Expert Systems*, and prior to that, the epochs of *Cybernetics* and *Mathematical Logic Machines*. We could ask ourselves what causes one AI epoch to start and another to end? We might look at an AI epoch from the points of view of (research, investment, utilization) funding, media, industry, and social impact. For critical or perhaps even cynical purposes, we might ask how or if AI epochs correspond or correlate to technology hype cycles. Figure 1 shows how a generic tech hype cycle curve.

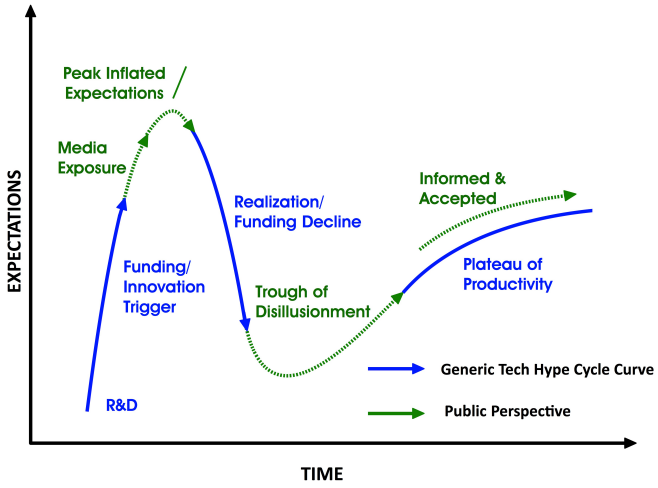

Figure 1: A Generic Tech Hype Curve.

Notice that a hype cycle reaches a peak where expectations are unrealistically high, the public is at its height of misinformation, and claims of technologys applications are inflated. Following the curve in Figure 1 is the realization that the inflated claims are not coming true, funding and investment decreases, and then the public becomes disillusioned. Is it the case that AI epochs happen to coincide with AI hype cycles? Are AI hype cycles typically components of epochs? Are they one in the same? One of our current interests with respect to the general public's vulnerability to AI hype cycles, how much do technical terminology misnomers, misconstrued AI concepts, and ill-defined AI vocabulary contribute? In

[1] https://www.youtube.com/watch?v=3PdxQbOvAlI

Figure 2, we show our first rough draft of AIs Epochs we've discussed so far.

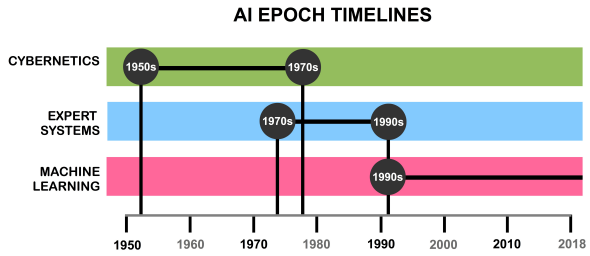

Figure 2: Rough Draft of AI's Epochs.

In Figure 2, each AI epoch has a timeline and within each epoch, there has been a hype cycle. For now we conjecture that hype cycles and epochs are not the same. Epochs start prior to the hype cycle and continues after the hype cycle is over.

Machine Learning and Back to Cybernetics

We are in the Machine Learning epoch. In AI Matters Volume 4 Issue 2, we talked about the age of Expert Systems. But before AI manifested itself in the form of Expert Systems, Cybernetics represented the promise of Artificial Intelligence. In June 1975, our own SIGART Newsletter No. 52 contained an article by Edward M. Riseman entitled "AAI AND BRAIN THEORY AT THE UNIV. OF MASSACHUSETTS". Riseman defines Cybernetics accordingly:

At the University of Massachusetts at Amherst, we use the name to describe the study of Artificial Intelligence (AI) and Brain Theory (BT). At the University of Massachusetts in 1975, the Computer Science and Information Science Program was roughly divided into three major areas: Computer Systems, Theory of Computation and Cybernetics.

In Riseman's SIGART article, he describes the 1974-1975 graduate courses in the Fundamentals of Cybernetics as including:

- Computational Cybernetics
- Artificial Intelligence
- Neurocybernetics of the Cerebral Cortex
- Scene analysis
- Natural vs. Artificial Intelligence
- Theorem Proving
- Implementing A Question-Answering systems

AI Magazine Volume 3 Number 1 (1982) reports that AI and Brain Theory at UMASS was still intermingling the notions of AI and Cybernetics.

Our program in AI is part of the larger departmental focal area of Cybernetics which integrates both AI and Brain Theory (BT). Our research also draws upon a new and expanding interdepartmental program in Cognitive Science that brings together researchers in Cybernetics, Linguistics, Philosophy, and Psychology.

The same fervor that Pop Culture now has for Machine learning was once had for Expert Systems. Prior to Expert Systems, feedback communication and control in animals and machines was all the craze and we were in the epoch of Cybernetics. It was just like the parallels between the AI/Machine Learning and the AI/Expert System explosions. Cybernetics was the talk of the town complete with scientific breakthroughs, industrial innovations, massive social change and a strong dose of Cyber-Techno-Hype. In 1969, Georges Boulanger, the president of the International Association of Cybernetics noted that "To the general public cybernetics conjures up visions of some fantastic world of the future peopled by robots and electronic brains" (Boulanger, 1969). Cybernetics represented the embodiment, incarnation, goals, aspirations, and funding targets of some of the same subject matter that is reserved for AI/Machine Learning. Cybernetics was a staple of science fiction and a fad among artists, musicians, and intellectuals in the 1950s and 1960s. These facts are well documented in Ronald Kline's 2015 tour de force book "The Cybernetics Moment" (Kline, 2015) where he painstakingly captures the history of Cybernetics, Information Theory, and the early AI Ethos. Upon a close reading of Kline's "The Cybernetics Moment", one can almost see how the Artificial Intelligence baton would eventually pass from Cybernetics to Expert Systems.

The intersection between the actual artifacts

of AI research and the general publics understanding and expectations with regard to those artifacts can be perilous. It appears that this has been the case for each of the AI epochs! As AI researchers, educators, and practitioners, we have an obligation to minimize (as best we can) the misconstructions, misunderstandings, and misrepresentations of our work and our field. At Ctest Laboratories, we are positing the notion of devising a "Standard Model" for an AI Cosmology that would standardize the structure, taxonomy, and ontology in a form that is suitable for consumption by early college STEM programs, high schools, and the general public. The moniker AI remains for much of the general public a reference to vague, ambiguous, futuristic, even threatening subject matter. This was also the case for the moniker Cybernetics. But after all what is cybernetics? Or rather what is it not, for paradoxically the more people talk about cybernetics the less they seem to agree on a definition, George Boulanger 1969 (Boulanger, 1969).

In this column, we would like to instigate a conversation with our fellow SIGAI members on this notion of a standard model for AIs Cosmology. So far we've offered conjectured notions of standardized epochs of AI, AI Hype Cycles, standard AI taxonomies, and a standard AI domain ontology as a starting point. What do we really mean by "Artificial Intelligence"? If the evolution of Artificial Intelligence can be divided into epochs, where are we now, and what was the first one? Can we understand the evolution of Artificial Intelligence, its technologies and applications in the context of a AI cosmological standard model? If we can construct an AI cosmological standard model, can we represent it in way that is suitable for consumption by the general public? If the epoch of Cybernetics preceded the epoch of Expert Systems, was there an AI epoch that preceded Cybernetics? We'll try to fill in more of the timeline in the next article.

References

Kline, R. (2015). *The Cybernetics Moment.* Baltimore, MD, USA: John Hopkins University Press.

Boulanger, G.R. (1969) Prologue: What is Cybernetics?, in *Survey of Cybernetics: A Tribute to Dr. Norbert Weiner*, ed. J. Rose. London, England. Iliffe Books.

Cameron Hughes is a computer and robot programmer. He is a Software Epistemologist at Ctest Laboratories where he is currently working on A.I.M. (Alternative Intelligence for Machines) and A.I.R (Alternative Intelligence for Robots) technologies. Cameron is the lead AI Engineer for the Knowledge Group at Advanced Software Construction Inc. He is a member of the advisory board for the NREF (National Robotics Education Foundation) and the Oak Hill Robotics Makerspace. He is the project leader of the technical team for the NEOACM CSI/CLUE Robotics Challenge and regularly organizes and directs robot programming workshops for varying robot platforms. Cameron Hughes is the co-author of many books and blogs on software development and Artificial Intelligence.

Tracey Hughes is a software and epistemic visualization engineer at Ctest Laboratories. She is the lead designer for the MIND, TAMI, and NO-FAQS projects that utilize epistemic visualization. Tracey is also a member of the advisory board for the NREF (National Robotics Education Foundation) and the Oak Hill Robotics Makerspace. She is the lead researcher of the technical team for the NEOACM CSI/CLUE Robotics Challenge. Tracey Hughes is the co-author with Cameron Hughes of many books on software development and Artificial Intelligence.

Artificial Intelligence and Jobs of the Future: Adaptability Is Key for Human Evolution

Sudarshan Sreeram (Vidya Mandir Senior Secondary School, Mylapore, Chennai - 04, India; sudarshan.sreeram@gmail.com)

DOI: 10.1145/3299758.3300060

Humans, the most intelligent species on planet Earth, have always found a way to adapt to their dynamic environment and survive in a consistent manner. We have made tremendous progress in evolving from primitive men and women in small, dreary caves to modern men and women in large, bustling cities. Our passionate strive to survive and affinity to adapt is similar to that shown by mosquitoes; regardless of the continuous development and massive deployment of insect repellent products, these undeniably irking insects, which cause a plethora of life-threatening diseases, always make a striking comeback (Stanczyk, Brookfield, Field, & Logan, 2013).

Humans have a natural tendency to be preoccupied with what interests us the most. We lack the ability to not think of anything, even for a second. Boredom is not an option for humans; this is evidenced by a research experiment conducted in 2014 by Timothy Wilson, a social psychologist professor at the University of Virginia, Charlottesville (Wilson et al., 2014). The results of this research found that humans would go to the extent of shocking themselves with small jolts of electricity, when locked alone in a lab, just to keep themselves occupied. Our minds were built to be aware of the world around us and engage in actions and decisions that shape it. Thus, it is not possible to stop continuously flowing thought processes (Whitehead, 2014).

Our rich and vivid history may be described as a continuously flowing river, which does not heed to a specific, predictable pattern. However, the river does take similar looking turns as it flows. Applying this hypothetical case to real world history, we see that major events, be it global revolutions, wars, famines, natural disasters, and even plagues, do recur, though, and at irregular intervals.

The time interval between two successive events is influenced by two factors our mindset and our reaction towards change. For example, consider an article published by Columbia University on the psychology of gullible mindsets (Valhouli, n.d.). According to this publication, a fourteen-year-old boy convinced forty-three of fifty people to support a ban on water. The boy referred to water as dihydrogen monoxide, the chemical name of water, in a petition that was circulated through the internet. An important factor to consider is the way this petition was framed to influence human minds; he claimed that dihydrogen monoxide was a harmful chemical that was found in tumors of cancer patients. Similar convincing sentences use widely known terms such as sweating, vomiting, erosion, and lethal to lure an uninformed reader. The reader takes the authenticity of the provided information for granted without checking facts or running background checks on its source. The information is then dispersed through casual conversations and discussions to a wider network, leading to the general public being falsely informed. Therefore, it is evident that our mindsets are one of the most gullible and volatile entities on our planet; widespread controversies, futuristic sci-fi movies, factually inaccurate data, and popularized opinions of highly influential representatives do a very good job at molding it.

Change-events are stimuli for human societies. When the change is for the better and provides visible benefits, the reaction is most likely positive. For example, the establishment of independence or democracy in a country, the creation of new job opportunities, etc. In a similar way, when the change is disruptive and removes certain key defining aspects of human life, the reaction is most likely negative. For example, the increase in income taxes and fuel prices or the decrease in job satisfaction leading to lack of peaceful life (Rohani & Pahazri, 2018).

Industrial revolutions, dubbed as the perma-

nent markers of change, have had monumental impacts, both positive and negative, on the lives of humans. An industrial revolution can be thought of as a steering wheel of a car that represents the human civilization; this car has made three sharp turns and is on a straight road, looking forward to making a fourth turn.

The first turn, which had the largest radius, witnessed the urbanization and industrialization of several rural societies as well as the deployment of steam engines, which played a major role in the development of transportation technologies during the 18th and 19th centuries. The second turn saw the growth of industrial centers around the world right before the first world war and the introduction of mass-produced electrical appliances, which included technological advancements that would change the world forever. The third turn, which started during the early 1980s, took the car onto the digital highway; the digital revolution and the boom of the internet made it possible for almost everyone in the world to stay connected. The fourth turn, the most anticipated of them all, deals with significant advancements in the fields of 3D printing (Kyle, 2018), bio-circuits (Prox, Smith, Holl, Chehade, & Guo, 2018), quantum computing (Chong, 2018), robotics, IoT (Agiwal, Saxena, & Roy, 2018), nanotechnology (Alvarez, Chan, Elimelech, Halas, & Villagrán, 2018), renewable energy technologies, self-driving vehicle technologies (Bansal & Kockelman, 2018), and finally, the most controversial and popular of them all Artificial Intelligence (AI) (Lu, Li, Chen, Kim, & Serikawa, 2018).

Ever since its inception in 1956, when John McCarthy held a conference on the subject (McCarthy, Minsky, Rochester, & Shannon, 1955), Artificial Intelligence has always caught the imagination of those who take it upon themselves to predict the outline of future societies (Barrat, 2015). It is undoubtedly an inspiring and captivating field; however, it is met with a great deal of controversies and misconceptions regarding its purpose and intention. Some research suggests controversial ways in which it would influence human life while others see the potential for a future where AI systems complement human societies. An oversimplified and partially misunderstood influence of AI is its effective take over of human-centralized jobs in the future.

The reason behind these misconceptions and controversies is due to a misunderstanding of one word Technology. Technology is a life-altering word that has found its place in human society. Over the past few decades, it has been integrating itself into all aspects of human life, and now, it is integral to our daily lifestyles. Our connection with technology has grown so deep that we cannot even fathom an alternate timeline in which humans live without it. Although widely debated, its intended purpose is to make human life simple and easy. It helps us achieve milestones and immensely speeds up the progress of our society as a whole.

We have witnessed the rapid transformation of our landscapes and interests as technology molds the shape of our future. It functions as a tool of change that teaches us how to adapt and evolve to meet the requirements of future societies. Today, some of us question these requirements and are doubtful of what this future might hold. It is easier and faster for humans to find flaws in an idea or a concept than it is to understand its true value.

Humans have shown varied reactions to the use of automation technology and robotics in several industrial areas (Knight, 2018). Consider the widely discussed transition of the manufacturing industry from recruiting skilled human workers to employing automated robotic technologies. The first flaw humans see with this transition is the massive unemployment of skilled laborers. What we fail to see is that this transition not only prevents us from working in hazardous conditions but also reduces a companys expenditure in manufacturing processes. In addition, the end product is of a much higher quality due to reduction in human errors, and the time taken by the different processes is significantly reduced, thereby increasing production rates to meet the demand. Yet, despite all these advantages, we still consider the loss of jobs as a prominent result.

Humans have always developed a notion of resistance and anxiety over advanced devices and technologies, be it self-driving cars or artificial intelligence. Even in the 1980s, the term computerphobia was widely circulated in magazines, newspapers, and psychology studies (Lafrance, 2015). The term was mainly used

to address a persons fear of being threatened and replaced by a computer. In addition, people felt hostile in the presence of a computer. Eventually the presence of computers in our daily lives grew exponentially. This is quite synonymous to what is going on in the present day with artificial intelligence technologies. The general trend marks a peak period of hype and gradually settles into the books of history.

The Gartner hype cycle (*Hype Cycle Research Methodology*, n.d.) for emerging technologies in 2018 depicts the presence of artificial intelligence on two of the five regions in the hype cycle graph innovation trigger and peak of inflated expectations. In addition, the graph also predicts that these various technologies under AI may reach the plateau of productivity a period of mainstream adaptation in roughly ten or more years (Panetta, 2018).

Ever since its creation, a computer was tasked with replicating and efficiently performing human tasks. It is just human notion that when a computer is able to do something new, we regain those feelings of anxiety and resistance. People, who are not properly informed of a specific technology, are easily persuaded by the depictions of AI as a threat. So, what is Artificial Intelligence?

Todays AI is a system of complex algorithms that do an exceptional job at processing and analyzing large volumes of data; it is a skill that humans lack due to the volatility and computational limitations of their memory systems. AI systems are purpose built for executing specific tasks, so there are several severe limitations on what it can accomplish; to put it in simple terms, AI is not yet anywhere near the ability to freely think as we do.

Super intelligence or true AI, on the other hand, is a form that current AI technologies could possibly take in the future as we gain a deeper knowledge regarding its fundamental concept the working of the human brain (Hassabis, Kumaran, Summerfield, & Botvinick, 2017). This future is not coming anytime soon; currently, there is no one on our planet who actually understands how super intelligence works (Byttner, 2017). Its progress is most entirely dependent on human brain research, which is currently progressing at a snails pace. Marlene Behrmann, a cognitive neuroscience professor at Carnegie Mellon University, claimed that neuroscientists only have a primitive understanding of human brain functions despite decades of research (Gornall, 2014). Thus, when humans themselves have not yet understood the complete working of their brains, it is improbable to create an intelligent system that can replicate it.

AI systems are excellent at performing repetitive tasks, which is evidenced by their purpose-built nature. It can automate almost any system that is repetitive or predictable. Thus, it is not surprising to see the results of AI technologies from Google and IBM (Hiner, 2018; Russell, 2017; Best, n.d.). Even if AI were to advance to match the human intellect, their only way of physically interacting with this world is through robotic bodies. Even the highest quality of robots are prone to failures, given that their developments should be cost effective. Regardless, robotic technicians will be required to diagnose and address any technical issues. Thus, there will always be space for humans.

Moreover, educational institutions around the world are offering courses on robotic systems to advance the technical skills required to work in this field. In fact, there is pressure on the educational sector to teach students to code from a young age (Barone, 2017). Several technology giants, including Apple, are bringing forward their support to educate children on coding practices (*Education - Teaching Code*, n.d.). Since the world is moving forward with technology as its backbone, the people living in it are expected to be well informed of it in a transparent manner.

As for the role of Artificial Intelligence (AI) in jobs of the future, AI systems would engage in cooperation with humans; these intelligent systems would provide humans with relevant data and suggestions to solve a particular task. There is a clear line that distinguishes jobs and tasks when considering the potential of current AI technologies (Pring, 2018). A job is a simple term for the human state of being preoccupied with multiple tasks while benefiting from a source of income; whereas, a task is a piece of work that acts as a subset of a job. It is a definite possibility that in the future, AI systems would complement the human job

workload rather than completely making it obsolete.

It is evident and inevitable, though, that new fields of interest to humans will rise as the depths of our knowledge in existing fields increase. Our future is influenced by a great deal of factors, AI being just one of them. Contrary to the notion that AI would make us unemployed in the future, this intelligence would, in actuality, boost job growth; though, the variety of jobs offered in the future may not be the same as what is offered in the present day. For example, humans may expect titles such as organ fabricator, commercial space pilot, alternate energy consultant, and trash engineer (*9 Top Jobs in 2030: Future Skills You Need To Learn Now*, n.d.). Prior to advancements in AI, it would have been hard to imagine job titles such as Machine Learning Developer or Artificial Intelligence Specialist. Using the same analogy, we may expect several new job titles in the future as technology evolves.

As jobs evolve over time so do the tools we use. This is evidenced by the history of art; during the 1800s, people had the time and patience to sit for days, months, or even years to complete a single painting, made entirely by hand. Today, we take aid of digital drawing programs and highly precise graphics pen displays or tablets to speed up that process and express our thoughts as they flow. Though, this example does bring into question the value of art and the presence of human touch, which, in my opinion, would fade as we move into the future. However, we do not surrender our creativity and intellectual abilities, but express the same in a different form that consumes less time and effort. Artists of the 19th century had brushes and canvases, artists of the 21st century have graphics tablets and digital canvases, and artists of the future may have technologies that combine the use of AI and Augmented Reality (AR), supported by the human brain, to provide everyone with an artistic toolset. This shows that creativity is a human characteristic that cannot simply disappear in the presence of artificial intelligence.

A research paper from Oxford university (Frey & Osborne, 2017), which claims that AI would take over human jobs of the future, bases its results on assuming that all jobs are repetitive and predictable (Surowiecki, 2017). This is not true in the real world as its dynamics hold innumerable factors. Consider the case where a robotic machine breaks down in a large scale factory. A human technician is required at this point for analyzing the situation and factoring multiple variables and carrying out an in-depth root cause analysis. Humans are better at decision making as they have a lot more input and background knowledge to work with and consider, which are not just derived from text book sources (*Why humans learn faster than AI-for now.*, 2018). Critical situations such as fire emergencies and natural disaster management, where even robots arent effective, do require human presence (*Why humans learn faster than AI-for now.*, 2018). Unless super intelligence is reached, AI is nothing all by itself apart from being a follower of specific instructions given by humans. To reach a future where AI would support job growth and cooperate with humans, three fundamental principles must be established: transparency in actions and decisions, restriction on information collection and access of data sources, and foundation for long-term reliability and trust.

In the far future, we may reach a state when artificial intelligence will prioritize and reorganize our skillsets; humans will think but not do. The creation of objects and physical completion of tasks will be taken care by humanoids powered by intelligent AI, as shown in film adaptations of popular sci-fi novels. Contrary to what is shown in these films, humanoids are less likely to turn against humanity as the three principles, coupled with safety protocols and other measures, would be enforced upon them. Humans would also have to make sure that the AI or AI based systems cannot find a way to bypass these security measures and gain access to unauthorized sections of data and services. In the same way, the general public should also be restricted from gaining privileges to access the core of the AI system. This would keep intelligence away from people who would utilize it for social misconduct, disinformation, and to promote corruption.

With these security measures in mind, humans can look forward to a future filled with cooperation and responsibility. In addition, with heavy investments in nuclear fusion and space mining research (James, 2018), humans may soon develop their first zero pollu-

tion, sustainable energy resource. This, along with further advancements in AI, would heavily influence the industries including the manufacturing and agriculture sectors. This would eventually reduce global resource prices and reach a stage wherein everyone would have access to resources. The same holds true for global inflation, where we may observe stable and controlled prices.

It is a social responsibility of humans to create a balance with ecology by which robots backed up by AI would contribute to the human community and its wellbeing; food generation, disease prevention, and medical care should be given the first priority. This would help us ensure a balance and mutual trust between AI based technology and humans.

For now, humans should focus on utilizing these intelligent platforms to solve global challenges such and global warming, climate change, and resource management. Furthermore, humans would need to be educated on the potential of such technological developments and refrain from exaggerating future outcomes to a point where it alters the expectations of general public on our future with AI. As change is inevitable, humans would learn to adapt and live with AI powered humanoids and robots in a harmonious manner. This would be accomplished through shedding our highly opinionated self and embracing AI based technologies as our way forward into the future.

In summary, humans will continue to see further advancements in the field of technology and experience its impacts on their lives. In these situations, humans are required to change their mindsets, be informed of the case, and use their ability to adapt to the environment. Artificial Intelligence would play a prime role in future societies. However, despite widespread controversies, AI systems of the future would be cooperative and collaborative with humans, providing them with the necessary data and suggestions to speed up the completion of job related tasks. This may hold true only if strict security measures are taken to ensure that these intelligent systems are not utilized for the wrong purposes. In conclusion, humans may look forward to an interesting future where adaptability is key for human evolution.

References

9 top jobs in 2030: Future skills you need to learn now. (n.d.). Retrieved November 5, 2018, from https://www.crimsoneducation.org/ie/blog/jobs-of-the-future

Agiwal, M., Saxena, N., & Roy, A. (2018, Mar). Towards connected living: 5g enabled internet of things (iot). *IETE Technical Review*, 1–13. Retrieved from http://dx.doi.org/10.1080/02564602.2018.1444516 doi: 10.1080/02564602.2018.1444516

Alvarez, P. J. J., Chan, C. K., Elimelech, M., Halas, N. J., & Villagrán, D. (2018, Aug). Emerging opportunities for nanotechnology to enhance water security. *Nature Nanotechnology*, *13*(8), 634–641. Retrieved from http://dx.doi.org/10.1038/s41565-018-0203-2 doi: 10.1038/s41565-018-0203-2

Bansal, P., & Kockelman, K. M. (2018, Mar 01). Are we ready to embrace connected and self-driving vehicles? a case study of texans. *Transportation*, *45*(2), 641–675. Retrieved from https://doi.org/10.1007/s11116-016-9745-z doi: 10.1007/s11116-016-9745-z

Barone, R. (2017, December). *9 reasons your child should learn to code (and one word of caution).* Retrieved November 5, 2018, from www.idtech.com/blog/5-reasons-your-child-should-learn-to-code

Barrat, J. (2015). *Our final invention: Artificial intelligence and the end of the human era*. St. Martin's Griffin.

Best, J. (n.d.). *Ibm watson: The inside story of how the jeopardy-winning supercomputer was born, and what it wants to do next.* Retrieved November 5, 2018, from https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/

Byttner, K. J. (2017, October). *There's only one job ai can't replace.* Retrieved November 4, 2018, from https://nordic.businessinsider.com/theres-only-one-job-ai-cant

-take-over-2017-10

Chong, F. T. (2018). Quantum computing is getting real. *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '18*. Retrieved from http://dx.doi.org/10.1145/3173162.3177152 doi: 10.1145/3173162.3177152

Education - teaching code. (n.d.). Retrieved November 4, 2018, from https://www.apple.com/education/teaching-code/

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254 - 280. Retrieved from http://www.sciencedirect.com/science/article/pii/S0040162516302244 doi: https://doi.org/10.1016/j.techfore.2016.08.019

Gornall, J. (2014, September). *Will we ever understand the human brain?* Retrieved November 4, 2018, from https://www.weforum.org/agenda/2014/09/understanding-human-brain/

Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017, Jul). Neuroscience-inspired artificial intelligence. *Neuron*, *95*(2), 245–258. Retrieved from http://dx.doi.org/10.1016/j.neuron.2017.06.011 doi: 10.1016/j.neuron.2017.06.011

Hiner, J. (2018, January). *Ai and jobs: Where humans are better than algorithms, and vice versa.* Retrieved November 5, 2018, from https://www.zdnet.com/article/ai-and-jobs-where-humans-are-better-than-algorithms-and-vise-versa/

Hype cycle research methodology. (n.d.). Retrieved November 4, 2018, from https://www.gartner.com/en/research/methodologies/gartner-hype-cycle

James, T. (2018). Scouting for resources. In T. James (Ed.), *Deep space commodities: Exploration, production and trading* (pp. 69–79). Cham: Springer International Publishing. Retrieved from https://doi.org/10.1007/978-3-319-90303-3_5 doi: 10.1007/978-3-319-90303-3_5

Knight, H. (2018, May). *How humans respond to robots: Building public policy through good design.* Retrieved November 5, 2018, from https://www.brookings.edu/research/how-humans-respond-to-robots-building-public-policy-through-good-design/

Kyle, S. (2018, Apr). 3d printing of bacteria: The next frontier in biofabrication. *Trends in Biotechnology*, *36*(4), 340–341. Retrieved from http://dx.doi.org/10.1016/j.tibtech.2018.01.010 doi: 10.1016/j.tibtech.2018.01.010

Lafrance, A. (2015, Mar). *The 1980s: An era of computerphobia.* Retrieved November 4, 2018, from https://www.theatlantic.com/technology/archive/2015/03/when-people-feared-computers/388919/

Lu, H., Li, Y., Chen, M., Kim, H., & Serikawa, S. (2018, Apr 01). Brain intelligence: Go beyond artificial intelligence. *Mobile Networks and Applications*, *23*(2), 368–375. Retrieved from https://doi.org/10.1007/s11036-017-0932-8 doi: 10.1007/s11036-017-0932-8

McCarthy, J., Minsky, M., Rochester, N., & Shannon, C. (1955). *A proposal for the dartmouth summer research project on artificial intelligence.* Retrieved November 4, 2018, from http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

Panetta, K. (2018, August). *5 trends emerge in the gartner hype cycle for emerging technologies, 2018.* Retrieved November 4, 2018, from https://www.gartner.com/smarterwithgartner/5-trends-emerge-in-gartner-hype-cycle-for-emerging-technologies-2018/

Pring, B. (2018, Feb). *Why ai won't take your job.* Retrieved November 4, 2018, from https://www.cio.com.au/article/633045/why-ai-won-t-take-your-job/

Prox, J., Smith, T., Holl, C., Chehade, N., & Guo, L. (2018, Feb). Integrated biocircuits: engineering functional multicellular

circuits and devices. *Journal of Neural Engineering*, *15*(2), 023001. Retrieved from http://dx.doi.org/10.1088/1741-2552/aaa906 doi: 10.1088/1741-2552/aaa906

Rohani, M. M., & Pahazri, N. (2018, Apr). Survey on how fluctuating petrol prices are affecting malaysian large city dwellers in changing their trip patterns. *IOP Conference Series: Earth and Environmental Science*, *140*, 012085. Retrieved from http://dx.doi.org/10.1088/1755-1315/140/1/012085 doi: 10.1088/1755-1315/140/1/012085

Russell, J. (2017, May). *Google's alphago ai wins three-match series against the world's best go player.* Retrieved November 5, 2018, from https://techcrunch.com/2017/05/24/alphago-beats-planets-best-human-go-player-ke-jie/

Stanczyk, N. M., Brookfield, J. F. Y., Field, L. M., & Logan, J. G. (2013, Feb). Aedes aegypti mosquitoes exhibit decreased repellency by deet following previous exposure. *PLoS ONE*, *8*(2), e54438. Retrieved from http://dx.doi.org/10.1371/journal.pone.0054438 doi: 10.1371/journal.pone.0054438

Surowiecki, J. (2017, August). *The great tech panic: Robots won't take all our jobs.* Retrieved November 6, 2018, from https://www.wired.com/2017/08/robots-will-not-take-your-job/

Valhouli, C. (n.d.). *Cutting down the dissonance: The psychology of gullibility.* Retrieved November 5, 2018, from http://www.columbia.edu/cu/21stC/issue-3.4/valhouli.html

Whitehead, N. (2014, July). *People would rather be electrically shocked than left alone with their thoughts.* Retrieved November 4, 2018, from https://www.sciencemag.org/news/2014/07/people-would-rather-be-electrically-shocked-left-alone-their-thoughts

Why humans learn faster than ai-for now. (2018, March). Retrieved November 5, 2018, from https://www.technologyreview.com/s/610434/why-humans-learn-faster-than-ai-for-now/

Wilson, T., Reinhard, D., Westgate, E., Gilbert, D., Ellerbeck, N., Hahn, C., ... Shaked, A. (2014). Just think: The challenges of the disengaged mind. *Science*, *345*(6192), 75–77.

Sudarshan Sreeram is a grade XII student at Vidya Mandir Senior Secondary School, Mylapore, Chennai, India. He is a dedicated student researcher and has published two technical papers at international conferences: 18th IEEE EEEIC, Palermo, Italy and 2018 IEEE SMC Junior, Miyazaki, Japan. Sudarshan is an Apple WWDC18 Scholar and the TUNZA Eco-Generation (Samsung-UNEP) Regional Ambassador to India. His wish is to become a mature researcher in the future and contribute to the field of computing. In addition, he is determined to advance his education to a doctoral level.