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: 10.1145/3402562

Join SIGAI

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

Also consider joining ACM.

Our mailing list is open to all.

Notice to Contributing Authors to SIG Newsletters

Submit to AI Matters!

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

AI Matters Editorial Board

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 6(1)

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: 10.1145/3402562.3402563

Issue overview

Welcome to the first issue of the sixth volume of the AI Matters Newsletter.

This issue is smaller than usual as a reflection to the disruptions caused by the on-going COVID-19 pandemic on the scientific community. We open with a call for funding by SIGAI Executive Committe Chair, Sanmay Das, to support AI activities promoting outreach. In our regular articles, we provide an event report by Michael Albert and John P Dickerson on the AAAI/ACM SIGAI Job Fair. In the policy column, Larry Medsker and Farhana Faruqe cover a new series on AI and Bias on the Policy Matters Blog with a commentary on Bias, Fairness, and Discrimination in the context of AI, along with discussions on AI policy issues with respect to work and timeframe for AI impact. Another regular column is our AI crosswords from Adi Botea. We have one contributed article from Cameron Hughes and Tracey Hughes on what constitutes the essential ingredients for AI in a world where AI is increasingly pervading in every walk of life. Enjoy!

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.

ACM SIGAI Activities Fund 2020: Call for Proposals

SIGAI Executive Committee (ACM SIGAI; chair_sigai@acm.org)
DOI: 10.1145/3402562.3402564

Overview

ACM SIGAI invites funding proposals for artificial intelligence (AI) activities that can take place entirely virtually and that have a strong outreach component to students, researchers, practitioners not working on AI technologies, or to the public in general.

We funded several proposals with a similar outreach thrust last year, but those were largely based around physical interactions. With the necessity of responding to the current situation and the Covid-19 pandemic, we are instead focusing this year on activities that can both take place virtually and provide resources, material, or engagement to the broader community through virtual means. Nevertheless, a focus of, and knowledge of, local populations that this could reach, and evidence of the ability to reach such populations, will be positively viewed.

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), virtual 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 virtually. 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 will fund up to five proposals. The funding will be provided either as reimbursement (up to $2000), or as an honorarium (up to $1000), keeping in mind that the honorarium option has tax consequences for the recipient. In the case of honorarium payments, we plan to ensure that the final "deliverable" is in keeping with what was promised in the proposal before sending funding, and that you produce a writeup of the project for submission to AI Matters, the SIGAI newsletter.

We will prioritize the following types of proposals:

a Proposals from ACM affiliated organizations other than conferences (such as ACM SIGAI chapters or ACM chapters).

b Out-of-the-box ideas that can be delivered rapidly,

c New activities (rather than existing and recurring activities),

d Activities with long-term impact, and

e Activities that reach many people. 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 are not appropriate for funding via this call. Likewise, webinars can be hosted in partnership with ACM SIGAI (e.g., `https://learning.acm.org/techtalks/ethicsai`). If you want to participate in programs such as these please reach out to us separately.

Submission Guidelines

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). This could also be a request for an honorarium payment instead;
- 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 reports 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 June 15, 2020 (UTC-12). Proposals should be submitted as pdf documents in any style at https://easychair.org/conferences/?conf=sigaiaaf2020

Committees

The program is organized by the SIGAI Executive Committee, and will be evaluated by them with the help of other SIGAI officers and external reviewers.

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.

Contact

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 Sanmay Das (sanmayd@acm.org) and Nicholas Mattei (nsmattei@gmail.com), the SIGAI Chair and Vice-Chair.

The SIGAI EC consists of the chair, vice-chair, secretary/treasurer, and past chair of SIGAI.

AAAI/ACM SIGAI Job Fair 2020: A Retrospective

Michael Albert (University of Virginia; albertm@darden.virginia.edu)
John P Dickerson (University of Maryland; john@cs.umd.edu)
DOI: 10.1145/3402562.3402565

Introduction

For the sixth year running, AAAI and ACM SIGAI jointly ran the popular AAAI/ACM SIGAI Job Fair. In lockstep with the growth of AAAI and the growth of the greater artificial intelligence and machine learning (AI/ML) community, our once-small job fair also grew. This year, thirty-eight companies and universities formally attended—typically with a booth, team of recruiters, swag, and other representatives—increasing from twenty-six companies during the job fair's previous run in 2019, and twenty-one companies in the year prior to that. Last year, we purchased a dedicated domain—https://aaaijobfair.com/—for the job fair. This year, we provided a link on that site through which job-seekers—students, postdocs, practitioners, and maybe even a few faculty—could upload their resumes or CVs. We then shared that data and contact information for slightly under four hundred job-seekers with participants on the other side: prospective employers. Those employers are listed in the section below.

Participating Employers

- Association for Computing Machinery (ACM)
- AI Singapore
- Alexander von Humboldt Foundation
- Amazon
- Apple
- AppZen
- Arthur AI
- Audatic GMbH
- Baidu
- Beijing Century Tal
- Bloomberg
- Charles River Analytics
- Dataminr
- Elsevier
- Google
- Happy Elements
- Hewlett Packard
- IBM
- Jane Street
- Kitware
- Mayo Clinic
- Microsoft
- NLMatics
- Openstream
- Point 72
- Raytheon BBN Technologies
- SGInnovate
- SigOpt
- Sony
- SuperbAI
- The Take
- Tongdun Technology
- United Technologies Research
- University of Chicago Crime Lab
- University of Southern California
- University of Zurich
- Waymo
- Visa

We kicked off the job fair with a brief motivational speech from the two organizers. Immediately following this—as in the two fairs prior to the present one—firms and universities were given the option to speak for 60 seconds (often a touch longer, but who's counting!) accompanied by a single slide. Most (over 75%) of employers chose to speak. As in previous iterations of the job fair, this introductory session served multiple purposes. First, it coalesced a group of interested job-seekers in a central location. Second, it coalesced recruiters and representatives from employers in that same central location—they wanted to hear about their competition, and also match faces with the names on the resumes and CVs they were provided. Third, importantly, this served to introduce job-seekers to firms that they might not have heard of—start-ups, smaller non-profits, and firms that may have less of a public presence. As usual, potential employers hailed from all over the world (e.g., China, Germany, Norway, Singapore, Switzerland, US) and from industry, academia, and government. Job-seekers were as diverse as the burgeoning AI/ML community is, presumably hailing from all over the world.

Figure 1: The job fair kicked off with a brief intro from organizers. This helped gather job-seekers and representatives from employers.

Figure 2: A representative from each of the participating firms gave a 1–2 minute, single-slide pitch.

We hope that all participants (on both sides!)

Figure 3: A view from the action on the ground.

in this year's fair enjoyed their time and found the experience worthwhile. We'd especially like to thank AAAI—*especially* Monique Abed from AAAI, for her boots-on-the-ground help at the conference and her work, with Carol Hamilton from AAAI, triaging emails in the months leading up to the job fair. If you have any comments regarding the fair itself, or suggested improvements, please get in touch!

Michael Albert is an Assistant Professor at the University of Virginia where he has joint appointments at the Darden School of Business and the School of Engineering and Applied Sciences. His research focuses on combining machine learning and algorithmic techniques to automate the design of markets.

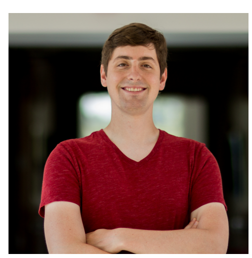

John P Dickerson is an Assistant Professor of Computer Science at the University of Maryland. His research centers on solving practical economic problems using techniques from computer science, stochastic optimization, and machine learning.

AI Policy Matters

Larry Medsker (The George Washington University; lrm@gwu.edu)
Farhana Faruqe (The George Washington University; faruqe@gwmail.gwu.edu)
DOI: 10.1145/3402562.3402566

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.

AI and DC

News Items for February, 2020

OECD launched the OECD.AI Observatory, an online platform to shape and share AI policies across the globe.

The White House released the American Artificial Intelligence Initiative: Year One Annual Report and supported the OECD policy

Bias, Ethics, and Policy

The Policy Matters blog has started a series on AI and Bias, with posts on background and context of bias in general and then focused on specific instances of bias in current and emerging areas of AI. The information is intended to inform ideas and discussions on public policy. We look forward to your comments and suggestions. Extensive work such as "A Survey on Bias and Fairness in Machine Learning" by Ninareh Mehrabi *et al*. is one of the background resources for the conversation. Additional resources are provided by Barocas, *et al*. The guest co-author of this column is Farhana Faruqe, doctoral student in the George Washington University Human-Technology Collaboration program.

AI Bias and Discrimination

Discrimination, unfairness, and bias are terms used frequently these days in the context of AI and data science applications that make decisions in the everyday lives of individuals and groups. Machine learning applications depend on data sets that are usually a reflection of our real world in which individuals have intentional and unintentional biases that may cause unfair actions and discrimination. Broadly, fairness is the absence of any prejudice or favoritism towards an individual or a group based on their intrinsic or acquired traits in the context of decision-making.

Direct Discrimination. As described by Ninareh Mehrabi *et al*., "Direct discrimination happens when protected attributes of individuals explicitly result in non-favorable outcomes toward them". Some traits like race, color, national origin, religion, sex, family status, disability, marital status, recipient of public assistance, and age are identified as sensitive attributes or protected attributes in the machine learning world. It is not legal to discriminate against these sensitive attributes, which are listed by the FHA and Equal Credit Opportunity Act (ECOA).

Indirect Discrimination. Even if sensitive or protected attributes are not used against an individual, still indirect discrimination can happen. For example, residential zip code is not categorized as a protected attribute, but from the zip code one may find out about race which is a protected attribute. So, "protected groups or individuals still can get treated unjustly as a result of implicit effects from their protected attributes".

Systemic Discrimination. In the nursing profession, the custom is to expect a nurse to be a woman. So, excluding qualified male nurses for nursing position is an example of systematic discrimination. Systematic discrimination is defined as "policies, customs, or behaviors that are a part of the culture or structure of an organization that may perpetuate discrimination against certain subgroups of the population."

Statistical Discrimination. In law enforcement, racial profiling is an example of statistical discrimination. In this case, minority drivers are pulled over more often than white drivers. The authors define “statistical discrimination is a phenomenon where decision-makers use average group statistics to judge an individual belonging to that group.”

Explainable Discrimination. In some cases, discrimination can be explained using attributes like working hours and education, which is legal and acceptable as well. In a widely used dataset in the fairness domain, males on average have a higher annual income than females because on average females work fewer hours per week than males do. Decisions made without considering working hours could lead to discrimination.

Unexplainable Discrimination. This type of discrimination is not legal as explainable discrimination because “the discrimination toward a group is unjustified”. Some researchers have introduced techniques during data pre-processing and training to remove unexplainable discrimination.

AI Bias and Fairness

In terms of decision-making and policy, fairness can be defined as “the absence of any prejudice or favoritism towards an individual or a group based on their inherent or acquired characteristics”. Six of the most used definitions are equalized odds, equal opportunity, demographic parity, fairness through unawareness or group unaware, treatment equality.

The concept of equalized odds and equal opportunity is that individuals who qualify for a desirable outcome should have an equal chance of being correctly assigned regardless of an individual’s belonging to a protected or unprotected group (e.g., female/male). Along with other concepts like “demographic parity” and “group unaware” are illustrated by the Google visualization research team with nice visualizations using a “simulating loan decisions for different groups”. The focus of equal opportunity is on the outcome of the true positive rate of the group. On the other hand, the focus of the demographic parity is on the positive rate only. Consider a loan approval process for two groups: group A and group B. For demographic parity, the overall number of approved loans should be equal in both group A and group B regardless of a person belonging to a protected group. Since the focus for demographic parity is on overall loan approval rate, the rate should be equal for both groups. Some people in group A who would pay back the loan might be disadvantaged compared to the people in group B who might not pay back the loan; however, the people in group A will not be at a disadvantage in the equal opportunity concept, since this concept focuses on true positive rate. As an example of fairness through unawareness “an algorithm is fair as long as any protected attributes A are not explicitly used in the decision-making process”. All of the fairness concepts or definitions either fall under individual fairness, subgroup fairness or group fairness. For example, demographic parity, equalized odds, and equal opportunity are the group fairness type; fairness through awareness falls under the individual type where the focus is not on the overall group.

A definition of bias can be in the three categories data, algorithm and a user interaction feedback loop: **Data** – behavioral bias, presentation bias, linking bias, and content production bias; **Algorithmic** – historical bias, aggregation bias, temporal bias, and social bias falls; **User Interaction** – popularity bias, ranking bias, evaluation bias, and emergent bias. Bias is a large domain with much to explore and take into consideration. Bias and public policy will be discussed in future blog posts.

AI and Work

The AI and Work session in the recent AAAI FSS-19 Symposia was a good example of exploration and research that should inform public policy making. All topics are related to, or could be expedited by, good public policy and aware policy makers. The deployment of AI technologies in the future will likely require humans to collaborate with AI systems, and this realization highlights the need for more sustained research on how to design such systems. High levels of autonomy and the ability to learn and interact with other systems, including humans redesigning work and rethinking incomes with bold ideas to improve the lives of workers and provide more interesting jobs with more meaning, purpose and dignity.

1. How do we design effective human-AI teaming?

2. What does participatory design look like for AI in the context of work?
3. What training do people need to be able to work successfully with smarter systems?

Time Frame for AI Impact

An interesting IEEE Spectrum article "AI and Economic Productivity: Expect Evolution, Not Revolution" by Jeffrey Funk questions popular claims about the pace of AI's impact on productivity and the economy. He asserts that "Despite the hype, artificial intelligence will take years to significantly boost economic productivity". If correct, this will have serious implications for public policy making. The article raises good points, but many of the examples do not look like real AI, at least as a dominant component. Putting "smart" in the name of a product does not make it AI, and automation does not necessarily use AI.

On a broader note, we should care about the technology language we use and beware of the usual practices in commercialization. As discussed previously, expanding the meanings of terms like AI, machine learning, and algorithms makes rational discourse more difficult. Some of us remember marketing of expert systems and relational databases: companies do a disservice to society by claiming each breakthrough technology actually is in their products. Here we go again today, with anything counting as AI depending on the point you want to make and the products you want to sell.

Another issue raised by the article relates to startups as the leaders of economic impact, as opposed to innovations from established industry and government labs. Any technology has an adoption curve, going from early adopters through the laggards, of about seven years. If you add to that the difficulties of making a startup succeed, a decade or so is probably the minimum timescale. A better perspective on revolution versus evolution could come from longitudinal evaluations looking at trends. In that case, a good endpoint for a hypothesis about dramatic impact on productivity might be the 2030-2035 time frame. Another difficulty of using a vague and broad notion of AI is that policymakers could miss the revolutionary impact of data science, which can, but may not, involve real AI. Data science probably has the best chance of dramatically impacting society and the economy soon and has the advantage of not having to involve designing and manufacturing physical objects, and thus not waiting for consumers to adopt new products. Data Science is already affecting society and employment through obvious, and not so obvious, revolutionary impacts on our work and lives.

Please join our discussions at the SIGAI Policy Blog.

References

1. Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. "A survey on bias and fairness in machine learning". CoRR, abs/1908.09635, 2019.

2. Moritz Hardt, Eric Price, and Nati Srebro. 2016. "Equality of Opportunity in Supervised Learning". In Advances in Neural Information Processing Systems 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 3315–3323. http://papers.nips.cc/paper/ 6374-equality-of-opportunity-in-supervised-learning.pdf.

3. Martin Wattenberg, Fernanda Viegas, and Moritz Hardt. "Attacking discrimination with smarter machine learning." Accessed at https://research.google.com/bigpicture/attacking-discrimination-in-ml/, 2016.

Larry Medsker is Research Professor and 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. 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.

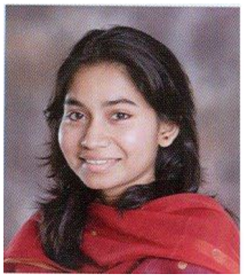

Farhana Faruqe is a doctoral student in the GW Human-Technology Collaboration Lab and Ph.D. program. Her research includes work on impacts of cognitive assistants on human-technology collaboration. She investigates issues in AI and Ethics, human-machine trust, and impacts of AI system bias on individuals and society.

Towards AI Ingredients

Cameron Hughes (Northeast Ohio ACM Chair; cameronhughes@acm.org)
Tracey Hughes (Northeast Ohio ACM Secretary; tracey.hughes@neoacmchapter.org)
DOI: 10.1145/3402562.3402567

We have AI in our cars, in our mobile phones, and AI in our video games. We have AI in medicine, AI in the military applications, and AI in government agencies. It's getting harder to find an aspect of our daily lives that doesn't purport to have some kind of interaction with AI. We are relinquishing more of the personal and professional decision-making process to vestiges of evolving notions of AI. Not only are we starting to defer to AI for the decision-making process, we are subtly transferring the ultimate responsibility for the decisions and the consequences of those decisions to the AI. The public's acceptance and reliance on various aspects of AI is becoming normalized. One major problem with this scenario is that we as a society are unclear about what constitutes AI. Our social position on AI is: we may not be able to concisely or correctly define it, but we all know it when we see it, right? Clearly the integral part that AI has in our society makes this position untenable and we can and should do better with our definition.

Even among AI researchers, educators, and practitioners, there is some consternation and disagreement about what constitutes AI and what doesn't, and the fact that we are currently in an AI hype cycle doesn't help matters. It's no wonder that in the general public the term "AI" is routinely misconstrued and misapplied. In the AI community, we have a responsibility to properly demarcate the tenets of Artificial Intelligence, its mathematics, science, and application. We need to define things clearly for the laymen and the public at large. But a clear, concise definition or presentation of AI for the laymen or the public at large is a tall order. AI research areas and techniques cover a wide range. Consequently, the commercial and government applications that deploy AI techniques from various research areas can have significant differences. Table 1 shows some of the research areas from the Bio-inspired [1] approach to AI and the symbolic approach to AI.

Table 1. Areas of Research from Bio-inspired and Symbolic approaches to Artificial Intelligence.

Some Areas in Bio-Inspired Artificial Intelligence Research

- Artificial Neural Networks,
- Cellular Automata,
- Bio-inspired, nature-inspired algorithms, search and machine learning algorithms,
- Deep Learning,
- Neural Networks,
- Behavior-based modeling,
- Swarm Intelligence,
- Evolutionary Computing,
- Nature-Inspired Metaheuristic Algorithms,

Some Areas in Symbolic Artificial Intelligence Research

- Knowledge Representation using frames, scripts, oav, conceptual graphs, modal logics,
- Expert Systems,
- Logic-based Machine Learning e.g. Inductive Logic,
- Programming and Relational Learning,
- Common Sense Reasoning,
- Situational Calculus, Event Calculus,
- Answer Set Programming,
- Symbolic and Mathematical Logic,
- Agent-Oriented Programming,
- Associative Memory Models

Bio-inspired approaches to AI have different assumptions, goals, vernacular, techniques and tools than what are typically found in symbolic approaches to AI. They both represent two very different schools of thought when it comes to the possibilities of replicating human intelligence and behavior by computer programs or in computer hardware. While there are hybrids of both approaches, most researchers tend to pick sides and consequently the boundaries and the fundamental definition of what constitutes AI can vary dramatically depending on which side has the podium. To make matters worse, vagueness and ambiguity are introduced once these differing techniques, tools, and vernacular are manifested in the form of commercial or government applications. Further, many of these applications carry serious social implications and can have major positive or negative impact on society. Table 2 shows some uses of AI-based systems that are used in the law enforcement, the military, and the legal system. The AI used in each of these areas can irrevocably change the trajectories of the human lives involved.

Table 2. Areas of Research from Bio-inspired and Symbolic approaches to Artificial Intelligence.

LAW ENFORCEMENT

- **Facial Recognition (FaceFirst, FACES, PoliceOne)**
  Identify criminals and missing persons in public spaces and video footage.
  - **Used by:**
  - Police Departments,
  - US Airports,
  - National Human Genome Research Institute,
  - FBI.
  - **Technology Used:**
  - Principal component analysis using eigenfaces,
  - Linear discriminant analysis,
  - Multilinear subspace learning.

- **Smarter Physical Robots (DroneDeploy in CA)**
  - Bomb detonation,
  - Crime surveillance,
  - Crime scene investigation,
  - Accident Scenes,
  - Search and Rescue,
  - Crowd monitoring.
  - **Used by:**
  - Police Departments,
  - ICE,
  - FBI,
  - Border Patrol.
  - **Technology Used:**
  - Unmanned and Remote controlled,
  - Autonomous,
  - Some equipped with Face Recognition and other technology.

- **Pattern Identification and Predictive Policing**
  - Identification of counterfeit goods,
  - Crime detection/prediction,
  - Forensic analysis,
  - Identification of potential perpetrators, victims and locations at increased risk of crime.
  - **Used by:**
  - Police Departments (California, Washington, South Carolina, Alabama, Arizona, Tennessee, New York and Illinois).
  - **Technology Used:**
  - Big Data Algorithms/ML behavior scripts,
  - Neural networks,
  - Databases,
  - Predictive Analytics.

- **Bias Mitigation Tools**
  - Removes racial biasness from police reports that identifies a suspects race.
  - **Used by:**
  - Police Departments.
  - **Technology Used:**
  - Automatic Translation Information,
  - Identification, retrieval and information extraction.

- **Speech Recognition Interface (Nexgen, Dragon Law Enforcement for CAD/RMS Systems)**
  - Police reports, incident reports and search.
  - **Used by:**
  - Police Departments.
  - **Technology Used:**
  - Customized-Language/Statistical modeling.

- **AFI System Data Collection and Mining(Project Maven/Algorithmic Warfare Cross-Function Team)**

- Used to identify individuals, associations, or relationships that pose a potential law enforcement or security risk.
- **Used by:**
- Homeland Security,
- DoD.
- **Technology Used:**
- Computer vision algorithms,
- TensorFlow APIs assist in object-recognition on unclassified data.

- **COPLINK:**
  Develop information and knowledge management systems technologies from heterogeneous data sources.
  - Captures, accesses, analyze, visualize, and share law enforcement-related information in order to solve cases and develop police reports.
  - **Used by:**
  - Police Departments,
  - ICE.
  - **Technology Used:**
  - Database assessment/integration.

MILITARY

- **Lethal Autonomous Weapon Systems**
  Can independently search for and engage targets based on programmed constraints and descriptions. Current systems (as of 2018) are restricted to a human giving final command to attack.
  - **Offensive (Drones, Unmanned Vehicles):** Autonomously search, identify, and locate enemies but can only engage with a target when authorized by mission command.
  - **Defensive:** Autonomously identify and attack oncoming weapon systems.
  - **Used by:**
  - Military.
  - **Technology Used:**
  - Facial recognition,
  - Decision-making algorithms.

LEGAL SYSTEM

- **Prison Sentencing Recommendations (Compas-Correctional Offender Management Profiling for Alternative Sanctions)**
- **Predicting outcomes of future trials**
  - Combat biasness and provide consistency,
  - Used to assist in the sentencing of defendants by human judges,
  - Weighing contradicting legal evidence, rule on cases in order to help humans make better legal decisions.
  - **Used by:**
  - Court systems (New York, Wisconsin, California, Florida, and other jurisdictions).
  - **Technology Used:**
  - Machine Learning.

The applications and the domain areas shown in Table 2, express the seriousness and growing deference to AI in many aspects of society. We in the AI community have a solemn obligation to define what we mean by Artificial Intelligence, its limits, its applicable scope, failure rates, risks, potential benefits and costs. We have to find a way to clearly and effectively educate and inform the public about this technology. How else does the domain expert, community advocate as well as the laymen or public at large navigate the potential morass of notions that can be attributed to AI?

Considering the severity of the consequences of using AI in these applications and domains, what precautions and communications have been employed regarding the fallibility of AI? Humans are fallible. Our technology is fallible. Data models can be incorrect and incomplete. They can be correct and complete but transient in nature because of structure, geographical or culture changes in the underlying sample sets. Rules learned by machine learning can expire as a result of dramatic changes in the environment from which the data culled. Valid decisions made, or conclusions drawn by the AI today, might be invalid 10 years from now when made in the same environment but different circumstances. We have to effectively communicate the risks involved. But where to start? We've yet to produce a clear concise, correct definition of AI suitable. How do we communicate limits, risks, safety issues?

Could labels on the proverbial container help? That is, should we have labeling in AI suitable for public consumption? In the same way that we are now starting to demand labels for fruits and vegetables that have been genetically modified, or meats that are synthetic. Should we develop labels that describe the ingredients of the AI in these applications or services? Should we apply the notion of expiration dates and safe use to our AI applications and services in the same way that we place expiration dates and directions for safe use on foods and other consumables? Data models that are used as the basis of decision-support systems for Prison Sentencing Recommendation systems or criminal profiling can quickly become outdated as the result of population, cultural, social, or geographic change. So that a model that may be appropriate now, may not be applicable at some future date. Those data models are historic and may not take into account the biasness that exist in the data collection. The cost and effort required to produce data models, or rule-based systems that support the AI might introduce reluctance to routinely update those models or rules. Once these systems are put in place, there will be inertia to prevent change.

What Are the AI Ingredients?

Therefore, should we label the AI application, data model, or rule set with an expiration date or other temporal restrictions? Being able to describe AI ingredients such as:

- **Epistemic metrics,**
- **Reliability indices,**
- **Expiration dates, and other temporal restrictions,**
- **Applicability scope,**
- **Failure rates,**
- **Safety considerations,**
- **Federal Regulations/Law Alignment,**

require that we have shared concise and correct definitions for the various AI technologies for which these ”ingredients” will define and affect. Perhaps the simplest path to educating the public is to provide labels that contain the AI ingredients of an application, device or service. Using labels would allow the public to know what AI is being purported, under which situations it can be reliably and safely used, and when it expires. Table 2 shows some of the domain areas where AI-based systems are in use that have serious consequences for the public. If we had labels that detailed the AI ingredients of these applications then the public would be in a better position to ascertain the value and legitimate uses of such applications.

In AI Matters, Volume 5 Issue 2 entitled: What Metrics Should We Use To Measure Commercial AI?, we discussed the need to not only define AI, but to also be able to measure the AI that is in any given application, device, or service. Clearly and concisely definable AI, measurable AI, and labels that contain AI ingredients are steps in the right direction of educating and informing the public with respect to the proper viability, applicability, and utility of applications, devices, or services that claim the use of AI. The notion of AI ingredients is related to transparent AI and explainable AI. In the next issue of AI Matters, we want to take a closer look at the notions of labeling AI, measuring AI, and transparent AI.

References

[1] Floreano, Dario and Mattiussi, Claudio (2008). *Bio-Inspired Artificial Intelligence. Theories, Methods, and Technologies*. Cambridge, The MIT Press.

Related Articles

- Greengard, Samuel. *Policing the Future.* Communications of the ACM. VOL. 55 No. 3. March 2012. pages 19-21.
- Kugler, Logan. *AI Judges and Juries: Artificial intelligence is changing the legal industry*. Communications of the ACM. VOL. 61 No. 12. December 2018. pages 19-21.
- Goose, Stephen., Arkin, Ronald. *The Case for Banning Killer Robots*. Communications of the ACM. VOL. 58 No. 12. December 2015. pages 43-45.
- Pasquale, Frank. *Law and Technology: When Machine Learning Is Facially Invalid*. Communications of the ACM. VOL. 61 No. 9. September 2018. pages 25-27.

- Chen, Hsinchun., Zeng, Daniel., et. al. *COPLINK: Managing Law Enforcement Data and Knowledge*. Communications of the ACM. VOL. 46 No. 1. 2003. pages 28-34.

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 NOFAQS 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.

AI Fun Matters

Adi Botea (Eaton; adibotea@eaton.com)
DOI: 10.1145/3402562.3402568

1	L	2	3	4	I	N	5	■	6	7	A	8
L	■	9			■	■	10	11			■	I
12	13				14	■	15				16	
17				■	18	19			■	20		
21			■	22 G				■	23			
24			25		A		■	26				
■	■	■	27			M	28			■	■	■
29	30	31			■	32	E			33	34	35
36				■	37			S	■	38		
39			■	40				■	41			
42			43			■	44	45				
C	■	46				■	■	47			■	N
48				■	49	E	A				N	

Across: **1)** AI research area. **6)** Individual among mondern-era Veneti. **9)** A word to emphasize. **10)** Woodwind instrument. **12)** Flowers whose purple popular variety symbolizes wisdom. **15)** Helpful accessories that are good to see. **17)** Fish-eating mammal. **18)** Young Flanders. **20)** AI research center in California. **21)** Apple computer brand. **22)** Hair-shaping accessories. **23)** West African country. **24)** Smoking accessory. **26)** Stupid person. **27)** City in New Jersey, USA. **29)** Ligament. **32)** Necklace medallion. **36)** Irish musician. **37)** Sings a tune softly. **38)** ___*, an optimal adversarial-search algorithm. **39)** Pale ___, a type of beer. **40)** El ___, one of the University of Texas locations. **41)** Personal cover. **42)** Convert points into a reward. **44)** Himalayas language. **46)** Girl from a famous picture in image processing. **47)** Garlic in a coq-au-vin dish. **48)** Achilles' weak spot. **49)** AI research area.

Down: **1)** Flat screen technology. **2)** Add a file to an email. **3)** Christmas in Paris. **4)** Logic gate type. **5)** Cooks food into a light brown color. **6)** Direct descendant. **7)** Landlord. **8)** AI research area. **11)** Test___, a dataset or platform for scientific evaluation. **13)** Sch. of Eng. and Appl. Sci. at Harvard. **14)** Grab work from another thread in parallel computing. **16)** Cogito ___ sum, the basis of Descartes' philosophy. **19)** Gods' mountain in Greece. **22)** Successfully seek an extension. **23)** Communicated orally. **25)** Keyword in an IF statement. **26)** Fashion sector. **28)** Evil spirit. **29)** AI research area. **30)** CPU status when waiting for jobs. **31)** Acupuncture utensil. **33)** Water-soluble base in chemistry. **34)** Mini canvas at one's fingertips. **35)** Muscle-shaping activity. **37)** Bright star in the Aries constellation. **40)** ___ pals in a remote relation. **41)** Be active in the space of hard drive disks. **43)** Snake-like fish. **45)** File type based on JAR.

Previous solution: RAILED - BELIE - ENTIRE - ARRANT - ADHERE - SEARCH - LEAN - PAID - SHE - MAC - HERA - LEER - SNARING - BINDS - ISSUERS - PESTS - ELAPSES - RATE - BRIG - CAP - ORR - FAST - BASE - UNIBUS - IBERIA - SEVENS - SOCCER - TREND - TAKERS

Acknowledgment: I thank Karen McKenna for her feedback. The grid is created with the AI system Combus (Botea, 2007).

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.