acm SIGAI

AI Matters

Annotated Table of Contents

Welcome to AI Matters, Volume 3, Issue 2
Eric Eaton& Amy McGovern, Co-Editors
Full article: http://doi.acm.org/10.1145/3098888.3098889
Welcome to AI Matters 3(2)

AI Events
Michael Rovatsos
Full article: http://doi.acm.org/10.1145/3098888.3098890
Summary of upcoming AI Events for 2017

AI Profiles: An Interview with Peter Stone
Amy McGovern & Eric Eaton
Full article: http://doi.acm.org/10.1145/3098888.3098891
Interview with Peter Stone

AI Policy
Larry Medsker
Full article: http://doi.acm.org/10.1145/3098888.3098892
AI Policy Statement

AI Education: Machine Learning Resources
Todd W. Neller
Full article: http://doi.acm.org/10.1145/3098888.3098893
Resources for teaching and learning about ML

AI Amusements: My Favorite Marvin
Tom Williams
Full article: http://doi.acm.org/10.1145/3098888.3098894
AI Crossword Puzzle

AI Practioners: On the Northeast Ohio ACM
Cameron Hughes
Full article: http://doi.acm.org/10.1145/3098888.3098895
Summary of recent events at the Northeast Ohio ACM Chapter

Situated Natural Language Interaction in Uncertain and Open Worlds
Tom Williams
Full article: http://doi.acm.org/10.1145/3098888.3098896
Dissertation abstract

Populating a Linked Data Entity Name System
Mayank Kejriwal
Full article: http://doi.acm.org/10.1145/3098888.3098897
Dissertation abstract

Incremental and Developmental Perspectives for General-Purpose Learning Systems
Fernando Martínez-Plumed
Full article: http://doi.acm.org/10.1145/3098888.3098898
Dissertation abstract

A Persuasive Virtual Chat Agent Based on Sociolinguistic Theories of Influence
Samira Shaikh
Full article: http://doi.acm.org/10.1145/3098888.3098899
Dissertation abstract

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: 10.1145/3098888

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.

Submit to AI Matters!

We're accepting articles and announcements now for the Summer 2017 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, Volume 3, Issue 2

Eric Eaton, Co-Editor (University of Pennsylvania; aimatters@sigai.acm.org)
Amy McGovern, Co-Editor (University of Oklahoma; aimatters@sigai.acm.org)
DOI: 10.1145/3098888.3098889

Welcome to the second issue in our third year of *AI Matters*. This issue features a timely new column: **AI Events**. This column is written by Michael Rovatsos and gives a summary of upcoming AI events for the rest of the year. AI Events will be a regular feature for future issues as well.

The AI Matters blog (http://sigai.acm.org/ai-matters/) contains regular postings on AI and current policy by Larry Medsker, our ACM SIGAI Public Policy Officer. He writes a summary article for each issue on **AI Policy**, but the blog contains the most up-to-date and full information. The blog also contains our latest AI Interview with Peter Stone.

In this issue's **AI Education** column, Todd Neller provides an overview of resources for machine learning, focusing on supervised, unsupervised, and reinforcement learning. His upcoming article will have a special focus on deep learning.

Our latest **AI Interviews** column highlights Peter Stone, Professor at the University of Texas at Austin and the COO and co-founder of Cogitai, Inc. We will continue to interview people involved in all aspects of AI, including academia, industry, and government, and we welcome suggestions for the next person you would like us to interview.

We are also introducing a new column from participating ACM chapters. This issue features a contribution from the Northeast Ohio ACM chapter.

The Call for Dissertation Abstracts remains open, so send us your abstracts! This issue contains four new dissertation summaries. One of our dissertation abstract authors, Tom Williams, also contributed the puzzle for this issue's **AI Amusements** column.

Thanks for reading! Don't forget to send your ideas and future submissions to *AI Matters*!

Eric Eaton is a Co-Editor of AI Matters. He is a faculty member at the University of Pennsylvania in the Department of Computer and Information Science, and in the General Robotics, Automation, Sensing, and Perception (GRASP) lab. His research is in machine learning and AI, with applications to robotics, sustainability, and medicine.

Amy McGovern is a Co-Editor of AI Matters. She is an Associate Professor of computer science at the University of Oklahoma and an adjunct associate 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.

Submit to AI Matters!

We're accepting articles and announcements now for the Summer 2017 issue. Details on the submission process are available at http://sigai.acm.org/aimatters.

AI Events

Michael Rovatsos (University of Edinburgh; mrovatso@inf.ed.ac.uk)
DOI: 10.1145/3098888.3098890

This new column provides information about recent and upcoming events that are relevant to the readers of *AI Matters*, including those supported by SIGAI. In an effort to provide concise summaries, we have abridged the descriptions provided by the organizers at the respective event web sites. We would love to hear from you if you are are organizing an event and would be interested in cooperating with SIGAI, or if you have announcements relevant to SIGAI. You can find out more about conference sponsorship and support at sigai.acm.org/activities/requesting_sponsorship.html.

The 16th International Conference on Artificial Intelligence and Law (ICAIL 2017) *London, UK, June 12-16, 2017*

http://nms.kcl.ac.uk/icail2017

The ICAIL conference is the primary international conference addressing research in Artificial Intelligence and Law, and has been organized biennially since 1987 under the auspices of the International Association for Artificial Intelligence and Law (IAAIL). ICAIL provides a forum for the presentation and discussion of the latest research results and practical applications; it fosters interdisciplinary and international collaboration. The conference will feature a main track for technical papers, a demonstration track, workshops, tutorials, a doctoral consortium and best paper prizes.

International Joint Conference on Rules and Reasoning (RuleML+RR 2017), *London, UK, June 12-16, 2017*

http://2017.ruleml-rr.org

RuleML+RR 2017 is the leading international joint conference in the field of rule-based reasoning, and focuses on theoretical advances, novel technologies, as well as innovative applications concerning knowledge representation and reasoning with rules. Stemming from the synergy between the well-known premier RuleML and RR events, one of the main goals of this conference is to build bridges between academia and industry. RuleML+RR 2017 aims to bring together rigorous researchers and inventive practitioners, interested in the foundations and applications of rules and reasoning in academia, industry, engineering, business, finance, healthcare and other application areas. It will provide a forum for stimulating cooperation and cross-fertilization between the many different communities working on rule-based systems.

The 14th International Conference on Informatics in Control, Automation and Robotics (ICINCO'17)

Madrid, Spain, July 26-28, 2017

http://www.icinco.org

The purpose of the 14th International Conference on Informatics in Control, Automation and Robotics is to bring together researchers, engineers and practitioners interested in the application of informatics to Control, Automation and Robotics. Four simultaneous tracks will be held, covering Intelligent Control Systems, Optimization, Robotics, Automation, Signal Processing, Sensors, Systems Modelling and Control, and Industrial Informatics. Informatics applications are pervasive in many areas of Control, Automation and Robotics. This conference intends to emphasize this connection.

International Conference on the Foundations of Digital Games (FDG '17) *Cape Cod, USA, August 14-17, 2017*

http://fdg2017.org

The International Conference on the Foundations of Digital Games (FDG) is a major international event that seeks to promote the exchange of information concerning the foundations of digital games, technology used to develop digital games, and the study of digital

games and their design, broadly construed. The goal of the conference is the advancement of the study of digital games, including but not limited to new game technologies, critical analysis, innovative designs, theories on play, empirical studies, and data analysis. FDG 2017 will include presentations of peer-reviewed papers, invited talks by high-profile industry and academic leaders, panels, workshops, and posters. The conference will also host a game competition, tech demo session, and a doctoral consortium.

The Taboo Challenge Competition

Melbourne, Australia, August 21, 2017
http://www.essence-network.com/challenge

This challenge competition, presented as a workshop at IJCAI 2017, is based on developing a computer program capable of guessing the name of a city from simple textual hints. The hints provided by the system were gathered from online games successfully played by humans, in order to ensure that hints resemble those produced using human intelligence and that the games are solvable. Guesser agents are evaluated over a predefined set of evaluation games. The competition is open to everyone, researchers, students, developers, and general AI enthusiasts from around the world. Attractive cash awards and an invitation to attend IJCAI in Melbourne await the winners!
Submission deadline: June 10, 2017.

International Conference on Web Intelligence (WI'17)

Leipzig, Germany, August 23-26, 2017
http://webintelligence2017.com

Web Intelligence (WI) aims to achieve a multidisciplinary balance between research advances in theories and methods usually associated with Collective Intelligence, Data Science, Human-Centric Computing, Knowledge Management, and Network Science. It is committed to addressing research that both deepens the understanding of computational, logical, cognitive, physical, and social foundations of the future Web, and enables the development and application of technologies based on Web intelligence. WI17 features high-quality, original research papers and real-world applications in all theoretical and technology areas that make up the field of WI.

The 11th ACM Conference on Recommender Systems (RecSys 2017)

Como, Italy, 27-31 August 2017
http://recsys.acm.org/recsys17

The ACM Recommender Systems conference (RecSys) is the premier international forum for the presentation of new research results, systems and techniques in the broad field of recommender systems. Recommendation is a particular form of information filtering, that exploits past behaviors and user similarities to generate a list of information items that is personally tailored to an end-users preferences. RecSys 2017, the eleventh conference in this series, will be held in Como, Italy. It will bring together researchers and practitioners from academia and industry to present their latest results and identify new trends and challenges in providing recommendation components. In addition to the main technical track, RecSys 2017 program will feature keynote and invited talks, tutorials covering state-of-the-art in this domain, a workshop program, an industrial track and a doctoral symposium.

Workshop on Hybrid Human-Machine Computing (HHMC 2017)

Guildford, UK, September 20-21, 2017
http://hhmc2017.commando-humans.net

When we talk about "computing" we often mean computers do something (for humans), but due to the more and more blurred boundary between humans and computers, this old paradigm of "computing" has changed drastically toward hybrid human-machine computing world where both humans and machines are working with and for each other. The main goals of the workshop are to bring researchers working in different disciplines but with common research interests together for exchanging research ideas, and to promote interdisciplinary collaborations and experience sharing between different subjects.
Submission deadline: May 31, 2017

The 4th international Workshop on Sensor-based Activity Recognition and Interaction (iWOAR '17) *Rostock, Germany, September 21-22, 2017*
http://iwoar.org

Systems aiming to provide the user with assistance or to monitor their behavior and condition rely heavily on sensors and the activities and interactions that they can recognize. The objective of iWOAR'17 is to discuss these challenges and possible solution approaches. The workshop focuses on: sensors, sensor infrastructures, and sensing technologies needed to detect user behaviors and to provide relevant interactions between systems and users; data and model-driven methods for intelligent monitoring and user assistance that supports users in everyday settings; novel applications and evaluation studies of methods for intelligent monitoring of everyday user behavior and user assistance using sensing technologies; intelligent methods for synthesizing assistance and interaction strategies using sensing technologies.
Submission deadline: June 15, 2017

Workshop on Interaction-Based Knowledge Sharing (WINKS)
Bolzano, Italy, September 21-23, 2017
http://www.iiia.csic.es/winks

This workshop is dedicated to challenges and solutions to knowledge sharing in interaction-based environments, ranging from the Internet of Things to multi-agent systems. Distributed systems increase the need for dynamic interactive knowledge sharing, while at the same time an increasing heterogeneity of resources renders this process more complex. The highly interdisciplinary workshop will involve discussions on requirements and suggestions to endow computational models with knowledge sharing capabilities in interactive scenarios. It will be part of the Third Joint Ontology Workshop (JOWO) (www.iaoa.org/jowo/JOWO-2017).
Submission deadline: July 17, 2017.

The Data Institute San Francisco Conference (DSCO17) *San Francisco, USA, October 15-17, 2017*
http://www.sfdatainstitute.org/conference.html

Deep learning has in recent years become the state of the art technique for a wide variety of computer vision and NLP problems, and produced breakthrough results in area like drug discovery, atomic physics, and dermatology. Topics can cover any development related to deep learning, such as: Training methods such as optimization algorithms and variational inference; regularization techniques; novel and/or highly impactful applications; architectural innovations; generative methods; model interpretability and visualization. As an academic conference, the committee is looking for technical talks, but are also trying to make the conference more accessible than most, so that more people can enjoy a wider range of talks. Therefore, we are asking people to spend some time thinking about how best to present their topic to a technical audience of people who may not necessarily be experts in the specific area of the talk. The goal here is to increase the level of collaboration between academia and industry.

The 9th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K'17) *Funchal, Portugal, November 1-3, 2017*
http://www.ic3k.org

The purpose of the IC3K is to bring together researchers, engineers and practitioners on the areas of Knowledge Discovery, Knowledge Engineering and Knowledge Management. IC3K is composed of three co-located conferences, each specialized in at least one of the aforementioned main knowledge areas: KDIR, the 9th International Conference on Knowledge Discovery and Information Retrieval, KEOD, the 9th International Conference on Knowledge Engineering and Ontology Development, and KMIS, the 9th International Conference on Knowledge Management and Information Sharing.
Submission deadline: June 12, 2017

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 in in multiagent systems, social computation, and human-friendly AI. Contact him at mrovatso@inf.ed.ac.uk .

AI Profiles: An Interview with Peter Stone

Amy McGovern (University of Oklahoma; amcgovern@ou.edu)
Eric Eaton (University of Pennsylvania; eeaton@cis.upenn.edu)
DOI: 10.1145/3098888.3098891

Abstract

This column is the third in our series profiling senior AI researchers. This month we interview Peter Stone.

Introduction

Our third profile for the interview series is Peter Stone, who is a Professor at the University of Texas at Austin and the COO and co-founder of Cogitai, Inc.

Figure 1: Peter Stone

Biography

Dr. Peter Stone is the David Bruton, Jr. Centennial Professor and Associate Chair of Computer Science, as well as Chair of the Robotics Portfolio Program, at the University of Texas at Austin. In 2013 he was awarded the University of Texas System Regents' Outstanding Teaching Award and in 2014 he was inducted into the UT Austin Academy of Distinguished Teachers, earning him the title of University Distinguished Teaching Professor. Professor Stone's research interests in Artificial Intelligence include machine learning (especially reinforcement learning), multiagent systems, robotics, and e-commerce. Professor Stone received his Ph.D in Computer Science in 1998 from Carnegie Mellon University. From 1999 to 2002 he was a Senior Technical Staff Member in the Artificial Intelligence Principles Research Department at AT&T Labs - Research. He is an Alfred P. Sloan Research Fellow, Guggenheim Fellow, AAAI Fellow, Fulbright Scholar, and 2004 ONR Young Investigator. In 2003, he won an NSF CAREER award for his proposed long term research on learning agents in dynamic, collaborative, and adversarial multiagent environments, in 2007 he received the prestigious IJCAI Computers and Thought Award, given biannually to the top AI researcher under the age of 35, and in 2016 he was awarded the ACM/SIGAI Autonomous Agents Research Award.

Getting to Know Peter Stone

How did you become interested in AI?

The first I remember becoming interested in AI was on a field trip to the University of Buffalo when I was in middle school or early high school (I don't remember which). The students rotated through a number of science labs and one of the ones I ended up in was a computer science "lab." The thing that stands out in my mind is the professor showing us pictures of various shapes such as triangles and

squares, pointing out how easy it was for us to distinguish them, but then asserting that nobody knew how to write a computer program to do so (to date myself, this must have been the mid '80s). I had already started programming computers, but this got me interested in the concept of modeling intelligence with computers.

What made you decide the time was right for an AI startup?

Reinforcement learning has been a relatively "niche" area of AI since I became interested in it my first year of graduate school. But with recent advances, I became convinced that now was the time to move to the next level and work on problems that are only possible to attack in a commercial setting.

How did I become convinced? For that, I owe the credit to Mark Ring, one of my co-founders at Cogitai. He and I met at the first NIPS conference I attended back in the mid '90s. We've stayed in touch intermittently. But then in the fall of 2014 he visited Austin and got in touch. He pitched the idea to me of starting a company based on continual learning, and it just made sense.

What professional achievement are you most proud of?

I'm made proud over and over again by the achievements of my students and postdocs. I've been very fortunate to work with a phenomenal group of individuals, both technically and personally. Nothing makes me happier than seeing each succeed in his or her own way, and to think that I played some small role in it.

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

It's cliche, but it's true. There's no better time of life than when you're a Ph.D. student. You have the freedom to pursue one idea that you're passionate about to the greatest possible, with very few other responsibilities. You don't have the status, appreciation, or salary that you deserve and that you'll eventually inevitably get. And yes, there are pressures. But your job is to learn and to change the world in some small way. I didn't appreciate it when I was a student even though my advisor (Manuela Veloso) told me. And I don't expect my students to believe me when I tell them now. But over time I hope they come to appreciate it as I have. I loved my time as a Ph.D. student. But if I had known how many aspects of that time of life would be fleeting, I may have appreciated it even more.

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

I have no idea. When I graduated from the University of Chicago as an undergrad, I applied to four CS Ph.D. programs, the Peace Corps, and Teach for America. CMU was the only Ph.D. program that admitted me. So I probably would have done the Peace Corps or Teach for America. Who knows where that would have led me?

What is a typical day like for you?

I live a very full life. Every day I spend as much time with my family as they'll let me (teenagers....) and get some sort of exercise (usually either soccer, swimming, running, or biking). I also play my violin about 3–4 times per week. I schedule those things, and other aspects of my social life, and then work in all my "free" time. That usually means catching up on email in the morning, attending meetings with students and colleagues either in person or by skype, reading articles, and editing students' papers. And I work late at night and on weekends when there's no "fun" scheduled. But really, there's no "typical" day. Some days I'm consumed with reading; others with proposal writing; others with negotiations with prospective employees; others with university politics; others with event organization; others with coming up with new ideas to burning problems.

I do a lot of multitasking, and I'm no better at it than anyone else. But I'm never bored.

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

I don't know. I have many interests and I can't help but pursue them all. And I multitask.

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

Rather than a book, I'll choose an author. As a teenager, I read Isaac Asimov's books voratiously – both his fiction (of course "I, Robot" made an impression, but the Foundation series was always my favorite), and his nonfiction. He influenced my thoughts and imagination greatly.

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.

AI Policy

Larry Medsker (George Washington University; lrm@gwu.edu)
DOI: 10.1145/3098888.3098892

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/).

Introduction

In this installment of the AI Policy column, I review the exciting (unprecedented) events since the Winter 2017 issue of AI Matters. New administration, executive orders, uncertainties – who knew AI Policy would be such a hot area?! This issue covers the following:

- the new SIGAI policy statement,
- suggestions, and request for discussion, of appropriate policy areas for AI, and
- potential issues concerning AI and jobs.

I welcome everyone to make blog comments to enrich our knowledge base of facts and ideas that represent SIGAI members.

Progress on a SIGAI Policy Statement

As a result of concerns about the new administration's executive orders on travel, the ACM and AAAS published statements that included general positions on the nature of scientific work and the need for freedom in research and communications. The following draft statement was discussed by the SIGAI executive committee and in responses to blog posts requesting input. The following statement has been submitted to the USACM for permission to disseminate the SIGAI position:

> The ACM SIGAI executive committee shares the view of its parent organization that "the open exchange of ideas and the freedom of thought and expression are central to the aims and goals of ACM. ACM supports the statute of International Council for Science in that the free and responsible practice of science is fundamental to scientific advancement and human and environmental well-being. Such practice, in all its aspects, requires freedom of movement, association, expression and communication for scientists. All individuals are entitled to participate in any ACM activity." SIGAI is working on policies to support inclusive participation in our AI-related activities. We encourage event organizers to share their efforts and experiences with us through our *AI Matters* newsletter at aimatters@sigai.acm.org and blog postings at https://sigai.acm.org/aimatters/blog/.

AI and Employment

A position seen in the media is that, just like other technological revolutions, new jobs will be created to replace the old ones. But is this a rationalization? Maybe the rate of technological change is of a different order in the AI and Big Data age compared to the industrial revolution. A more optimistic outcome than automation leading to mass unemployment is to see these technologies as a tool that will allow people to achieve more; for example, working together with cognitive assistants. So, which way will it be?

MIT economist Erik Brynjolfsson and co-author Andrew McAfee, in "The Second Machine Age", explore the question of what jobs will be left once software has perfected the art of driving cars, translating speech, and other tasks once considered the domain of humans. Along with the impacts of AI R&D, the rapidly emerging field of data science, spawned by the ubiquitous role of data in our society, is producing tools and methods that surpass human ability to manage and analyze data. Some researchers estimate that 50% of total US employment is in the high-risk category, meaning that associated occupations can potentially be automated. In the first wave, they predict that most workers in transportation and

logistics occupations, together with the bulk of office and administrative support workers and labor in production occupations, are likely to be substituted by computer capital.

Policymaking will no doubt lag behind the technology. Now is the time to discuss and advocate policies that address innovating our education systems, redefining employment, and investigating alternate economic systems. A goal for future AI Matters blog postings is to monitor countries and individuals who are thinking about and experimenting with alternate ways to address the ongoing advances in AI R&D and their impacts on employment that are already being seen.

Policy Areas Relevant to AI

At the time of writing this column, we await information on the current administration's policies on science, and particularly on perspectives on AI. The Obama administration released the reports *Artificial Intelligence, Automation, and the Economy* and *Preparing for the Future of Artificial Intelligence*. Potential implications of AI for society include the speed of change due to advances in technology; loss of control and privacy; job destruction due to automation; and a need for laws and public policy on AI technology's role in the transformation of society. An important point is that, compared to the industrial revolution, AI's impact is happening much faster and at a much larger scale of use than past developments. We see increasing evidence that nongovernmental organizations are recognizing the likelihood of disruption of operations that will happen whether or not change is intentional and planned.

In our current political environment, not much information is available about the new administration's understanding of AI technology and the need for policies, laws, and planning. Appointment to key administrative positions have yet to be made, and the status of the White House Office of Science and Technology Policy is not available on the Website. AI technology and applications will continue to grow rapidly, but whether or not public policy will keep pace is in doubt. The administration may take the position that AI will not cause job losses for many decades, and that perspective could distort assumptions about labor market trends and lead to policy mistakes.

At the recent AAAS Science and Technology Policy Forum, Matt Hourihan, who runs the R&D Budget and Policy Program at AAAS, gave preliminary perspectives on the next federal budget's impact on R&D. He compared the responses by Congress in previous administrations; for example, bipartisan pushback on efforts to reduce NIH budgets. He also discussed the relative emphasis in administrations on applied vs. basic research funding in non-defense spending, and the possibility of reducing applied funding in the next budget. Key slides and details from his presentation are available and links are in the Resources section below. Hourihan says, "In fact, there is a strong argument to be made that the first Trump Administration budget is the toughest of the post-Apollo era for science and technology, even with substantial information gaps still to be filled in." While still awaiting details, "the picture that *does* emerge so far is one of an Administration seeking to *substantially* scale back the size of the federal science and technology enterprise nearly across the board – in some cases, through agency-level cuts not seen in decades."

A goal for this column and blog posts is to monitor the administration's movement toward AI policies, development of budget policies that impact science and AI in particular, and the trends in R&D and its impacts on individuals and society. As always, input from the SIGAI membership is most welcome.

Upcoming

The theme for the SIGAI Public Policy posts for May is "Relevant Policies for AI R&D." We will look at potential policies today that could anticipate impacts of policies, or lack of policies, on progress in research on AI and preparation for the impacts of AI on individuals and society. Policy areas include budget allocation, anticipation of future employment, and the combined impacts of AI and data science. We welcome your input and discussion at the AI Matters blog!

Resources

- AAAS: https://www.aaas.org/program/center-science-policy-and-society-programs

- Artificial Intelligence, Automation, and the Economy: https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF
- Whitehouse Report on the Future of AI: https://obamawhitehouse.archives.gov/blog/2016/10/12/administrations-report-future-artificial-intelligence
- Preparing for the Future of Artificial Intelligence: https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf
- DC Data Science, AI, and Policy: http://www.datacommunitydc.org/data-science-dc/
- Ajay Agrawal, Joshua Gans, and Avi Goldfarb. The Obama Administration's Roadmap for AI Policy: https://hbr.org/2016/12/the-obama-administrations-roadmap-for-ai-policy
- The Second Machine age: http://secondmachineage.com/
- Human Work in the Robotic Future, Policy for the Age of Automation: https://www.foreignaffairs.com/articles/2016-06-13/human-work-robotic-future
- De Lange Conference on Humans, Machines, and the Future of Work: http://delange.rice.edu/index.html
- Artificial Intelligence and Life in 2030: http://ai100.stanford.edu/2016-report
- Matt Hourihan, AAAS, on The Ups and Downs of the U.S. Science Budget: https://www.aaas.org/blog/member-spotlight/matt-hourihan-ups-and-downs-us-science-budget
- AAAS reports and slides on the R&D Budget and Policy: https://www.aaas.org/program/rd-budget-and-policy-program
- AAAS, The Trump Administration's Science Budget: Toughest Since Apollo?: https://www.aaas.org/news/trump-administrations-science-budget-toughest-apollo
- AAAS, First Trump Budget Proposes Massive Cuts to Several Science Agencies: https://www.aaas.org/news/first-trump-budget-proposes-massive-cuts-several-science-agencies

Larry Medsker is a Research Professor of Physics and 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 a Professor in Computer Science and in Physics, where he was a co-founder of the Siena Institute for Artificial Intelligence. His research and teaching continues at GW on the nature of humans and machines and the impacts of AI on society and policy[a]. Professor Medsker's research in AI includes work on artificial neural networks and hybrid intelligent systems. He is the Public Policy Officer for the ACM SIGAI.

[a] http://www.humai.org/humai/ and http://humac-web.org/

AI Education: Machine Learning Resources

Todd W. Neller (Gettysburg College; tneller@gettysburg.edu)
DOI: 10.1145/3098888.3098893

Introduction

In this column, we focus on resources for learning and teaching three broad categories of machine learning (ML): supervised, unsupervised, and reinforcement learning. In our next column, we will focus specifically on deep neural network learning resources, so if you have any resource recommendations, please email them to the address above.

Machine Learning

In addition to the much-loved and ubiquitous *Artificial Intelligence: a Modern Approach* textbook (Russell & Norvig, 2009), there are a number of excellent introductory texts and tools specific to ML. Christopher Bishop's *Pattern Recognition and Machine Learning* (Bishop, 2006) is among the clearest introductions to ML with emphasis on Bayesian techniques. Kevin Murphy's *Machine Learning: A Probabilistic Perspective* (Murphy, 2012) is more comprehensive and advanced, bringing together the best of many authors into a massive introduction. David Barber's *Bayesian Reasoning and Machine Learning* (Barber, 2012) is widely recommended for advanced undergraduates and graduate students for self-study without considerable prior background. It is also freely available[1]. Tom Mitchell's *Machine Learning* (Mitchell, 1997) is the classic introduction to the field, offering a solid foundation and broad perspective.

There are many online resources that are very helpful for the study of ML. I highly recommend Andrew Ng's free online Coursera Machine Learning course[2]. Experiential learning is key, so getting datasets to practice ML techniques with is vital to one's learning. The UC Irvine Machine Learning Repository[3] is an excellent source for datasets that one can browse datasets by ML task, attribute type, size, application area, etc. Kaggle, a site that supports ML learning through ML competitions, also offers keyword searchable datasets[4]. A variety of ML assignments are available via the Model AI Assignments repository[5]. Many more recommendations of good resources for learning ML may be found at https://www.quora.com/How-do-I-learn-machine-learning-1.

Supervised and Unsupervised Learning

For those interested in focusing on supervised and unsupervised learning techniques that span from classical statistical methods with high bias, low variance, and better interpretability (e.g. linear regression) to AI methods with low bias, high variance, and better predictive performance, here are resources I would particularly praise and highlight:

Hastie, Tibshirani, and Friedman's *The Elements of Statistical Learning* (Hastie, Tibshirani, & Friedman, 2009) is a particularly well-written introduction spanning the bias-variance trade-off spectrum, and it is freely available[6]. For a gentler undergraduate-friendly introduction featuring labs using R, I highly recommend Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani's *An Introduction to Statistical Learning with Applications in R* (James, Witten, Hastie, & Tibshirani, 2014), also freely available[7] and integrating well with the freely available RStudio[8].

For a free, open-source, Java-based set of Data Mining tools, Weka offers a broad range of classification and regression tools[9].

[1] web4.cs.ucl.ac.uk/staff/D.Barber/textbook/090310.pdf
[2] https://www.coursera.org/learn/machine-learning
[3] http://archive.ics.uci.edu/ml/
[4] https://www.kaggle.com/datasets
[5] https://www.kaggle.com/datasets
[6] https://statweb.stanford.edu/ tibs/ElemStatLearn/
[7] http://www-bcf.usc.edu/ gareth/ISL/
[8] https://www.rstudio.com/products/rstudio/
[9] http://www.cs.waikato.ac.nz/ml/weka/

Ian Witten and Eibe Frank's *Data Mining: Practical Machine Learning Tools and Techniques* (Witten & Frank, 2005) is Weka's companion textbook.

Reinforcement Learning

For those interested in reinforcement learning (RL) in particular, I strongly recommend Richard Sutton and Andrew Barto's *Reinforcement Learning: An Introduction* (Sutton & Barto, 1998). This was the first text that presented a clear, unified view of dynamic programming, Monte Carlo, and temporal-difference learning techniques, and remains the best foundational reading for study of RL. It is also freely available[10]. More algorithms are briefly described in Csaba Szepesvári's *Algorithms for Reinforcement Learning* (Szepesvári, 2010), also freely available[11]. Richard Sutton recommends Marco Wiering and Martijn van Otterlo's *Reinforcement Learning: State-of-the-Art* (Wiering & van Otterlo, 2012) as "a valuable resource for students wanting to go beyond the older textbooks and for researchers wanting to easily catch up with recent developments".

Your Favorite Resources?

These are but a few good starting points for learning about ML. If there are other resources you would recommend, we invite you to register with our wiki and add them to our collection at http://cs.gettysburg.edu/ai-matters/index.php/Resources.

References

Barber, D. (2012). *Bayesian reasoning and machine learning*. New York, NY, USA: Cambridge University Press.

Bishop, C. M. (2006). *Pattern recognition and machine learning (information science and statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: data mining, inference and prediction* (2nd ed.). Springer.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2014). *An introduction to statistical learning: With applications in r*. Springer Publishing Company, Incorporated.

Mitchell, T. M. (1997). *Machine learning* (1st ed.). New York, NY, USA: McGraw-Hill, Inc.

Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. The MIT Press.

Russell, S., & Norvig, P. (2009). *Artificial intelligence: A modern approach* (3rd ed.). Upper Saddle River, NJ, USA: Prentice Hall Press.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: an introduction* (1st ed.). Cambridge, MA, USA: MIT Press.

Szepesvári, C. (2010). *Algorithms for reinforcement learning*. Morgan and Claypool Publishers.

Wiering, M., & van Otterlo, M. (Eds.). (2012). *Reinforcement learning: State-of-the-art*. Springer.

Witten, I. H., & Frank, E. (2005). *Data mining: Practical machine learning tools and techniques, second edition (morgan kaufmann series in data management systems)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

[10] http://incompleteideas.net/sutton/book/the-book.html

[11] https://sites.ualberta.ca/ szepesva/RLBook.html

Todd W. Neller is a Professor of Computer Science at Gettysburg College. A game enthusiast, Neller researches game AI techniques and their uses in undergraduate education.

AI Amusements: My Favorite Marvin

Tom Williams (Tufts University; williams@cs.tufts.edu)

DOI: 10.1145/3098888.3098894

Solve this acrostic puzzle by filling answers to the clues on the left into the dashed spaces to the right. Copying the letter above each numbered space into the same-numbered space in the grid above will yield a quotation from a work whose title and author will be spelled out in the first letter of each of your answers. Words are separated in the quotation grid by black squares, and as such may wrap between grid rows.

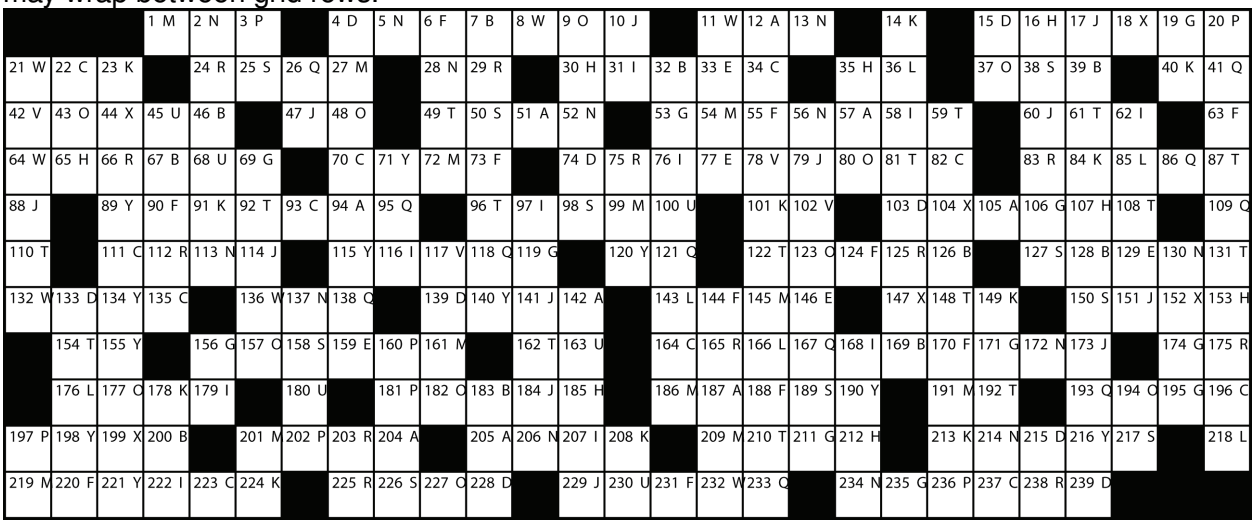

A. Famed transistor growth observation (2 wds.) — 205 187 94 51 105 142 57 12 204
B. George Washington or Brigham Young — 67 46 183 128 39 126 169 32 7 200
C. AAAI 2015 location (vis-a-vis the Colorado) (2 wds.) — 111 223 164 82 34 22 196 237 70 93 135
D. Proof technique — 74 228 103 15 4 139 133 215 239
E. Worked with Java, say — 33 129 146 77 159
F. Kalman filter component (2 wds.) — 90 231 170 124 220 73 55 63 188 144 6
G. Local CompSci society (2 wds.) — 174 235 119 69 53 156 171 106 195 19 211
H. Magazine launched in 1933 — 16 153 30 65 35 107 212 185
I. AAAI past president — 62 222 116 76 179 58 168 97 207 31
J. Annual source of anxiety and opportunity (2 wds.) — 151 60 47 141 88 10 229 173 79 17 184 114
K. Where to find a scanning tunneling microscope (2 wds.) — 84 213 224 91 149 40 178 208 23 14 101
L. Final thesis verb — 85 218 176 36 143 166
M. Where one might show their steps (2 wds.) — 209 72 186 99 27 54 219 145 191 1 161 201
N. Text found on the Great Seal of the United States (3 wds.) — 52 56 130 5 234 28 214 2 13 137 172 206 113
O. Topic in a first AI course (hyph., 2-wds.) — 157 37 43 177 123 48 80 227 9 194 182
P. Commentate a Robocup game, say — 202 236 20 3 197 181 160
Q. Branching technique in a 1st programming course (2 wds.) — 109 121 118 26 193 233 167 41 86 95 138
R. Like some logic (hyph.) — 29 125 75 24 112 165 238 225 203 175 83 66
S. AI paradigm with payoffs (2 wds.) — 98 189 150 127 217 226 38 50 158 25
T. Each of a set of combined distributions (2 wds.) — 96 154 59 162 131 108 87 81 148 49 122 210 61 192 110 92
U. Source of some famous laws (2 wds.) — 180 68 45 230 163 100
V. December air — 117 42 102 78
W. Where one might vie for vouchers — 21 132 232 64 11 8 136
X. Scale with an absolute lower bound — 152 104 199 18 44 147
Y. Frugal conference accommodations (2 wds.) — 190 120 216 221 140 71 134 115 155 89 198

Tom Williams completed a joint PhD in Computer Science and Cognitive Science in May 2017 with Professor Matthias Scheutz at Tufts University, and will be joining Colorado School of Mines in August 2017 as an Assistant Professor of Computer Science. Tom's research focuses on allowing robots to communicate in natural language in uncertain and open worlds, with applications to assistive and search-and-rescue robotics.

AI Practioners: On the Northeast Ohio ACM

Cameron Hughes (Northeast Ohio ACM Chair; cameronhughes@acm.org)
DOI: 10.1145/3098888.3098895

About NEOACM Chapter

I am currently the chair of NEOACM (Northeast Ohio ACM), a professional chapter of the ACM, located in the state of Ohio in the U.S. founded in 2008. NEOACM is quickly coming up on its 10th anniversary. Our chapter, like most, is dedicated to advancing the art, science, engineering, and application of information technology, serving both professional and public interests. We accomplish much of our mission by hosting workshops, speakers, and panel discussions that are open to the public. In my capacity as chair, I have had the opportunity to influence the direction of some of the workshops and panel discussions that we host. I've always had a predilection for most things AI and I'm particularly interested in how the descriptions of embedded AI technologies are worded when those technologies are integral to goods and services targeted toward the average citizen or general public. Sure as insiders we throw around phrases like computational linguistics, particle swarm optimization, machine learning, domain ontologies, agent-oriented architectures, etc. all the time because its normal vernacular. But we know in most instances our techno-speak will need to be translated into terminology that's more consumer-oriented once commercial applications start to be generated. It's the gap that I worry about. The translation from artificial intelligence, computer science, and mathematics specific terminology into descriptions that the average citizen will end up trying to grapple with. For instance:

- What does "smart car" really mean?
- What makes a "smart city", smart?
- What does machine learning have to do with lowering the cost of prescriptions?
- How will AI technologies be described when they are integral to issues that are presented to voters on a ballot?

Will the promoters and politicians come up with descriptions of AI technologies that are accurate, informative, and easy to understand or will they come up with catchy phrases, market speak, slogans and buzzwords that are self-serving? So when I have a chance to suggest a workshop or a panel discussion, there is usually a component that deals with the intersection between the descriptions embedded AI technology and the public interest.

Our Fall Practical AI Panel Discussion

In fall of 2016, NEOACM hosted a panel discussion on artificial intelligence, entitled: We Come In Peace at Kent State University. In that discussion, we covered some of the many hopeful scenarios that AI has to offer and will provide in the future, but we also discussed the challenges that have to be met. The panel discussion was entitled, "We Come In Peace." The phrase: "We Come in Peace" means "having no hostile intent." It is a phrase stereotypically used in science fiction narratives by extraterrestrial visitors upon first meeting the inhabitants of a planet (typically earth). The first science fiction movie that used this phrase in this context was *The Day The Earth Stood Still* by Klaatu after leaving his spaceship:

We have come to visit you in peace – and with good will.

We used this phrase as the title of the discussion because the implementation of AI should "have no hostile intent" to society, the impact should be positive and of great benefit. AI has been successfully applied to a number of societys challenging problems and we covered some of the successful deployments and the potential use of AI in various topics that are essential for social good. This was including but not limited to urban computing, robotics, and public welfare. But we also discussed issues of ethics, liability, safety, and control. But one of the biggest hurdles in addressing our diverse audience comprised of students and professors of computer science and non-technical disciplines, computing professionals, and a large number of average citizens, was making sure everyone had a base

understanding of the meaning of some popular AI terminology and some of the basic techniques used such as data mining, machine learning, and predictive analytics.

There were six panelists:

- **Dr. Pooyan Fazil**, AI and Human-Robot Interaction Researcher from Cleveland State University's Electrical Engineering and Computer Science Department,
- **Dr. Sven Koenig**, fellow of the AAAI (Association for the Advancement of AI),
- **Dr. Mark Vopat**, Applied Ethicists from Youngstown State University,

and three panelists from industry:

- **Paul Carlson**, an Intelligent Community Strategist, who had conducted IT transformation resulting in the City of Columbus being recognized as the 2015 "Most Intelligent Community" in the world by the Intelligent Community Foundation (ICF), achieving international status for Central Ohio,
- **Stuart Johnson**, Vice President of Connected Nation, a leading technology organization committed to assisting states, communities, families, and individuals to adopt and use improved broadband access,
- **Nicholas Wagner**, the cofounder and CTO of AvatarFleet that focuses on bringing innovative software solutions to the trucking industry.

This fall, NEOACM will host "We Come In Peace / We Come in Pieces", our 2nd Annual AI Panel Discussion at Youngstown State University. The new subtitle "We come in Pieces" will address the fact that many AI technologies are 'quietly' embedded within non AI-based technologies. Again the goal for the panel discussion will be to decode and demystify AI terminology, misnomers, and market speak. Our goal will be to help educate the public on some of the AI projects and technologies such as self-driving vehicles, cognitive computing, and smart cities that being implemented in our communities, highlighting some of real social, moral, and economic impact. We will include a discussion of how AI and autonomous systems are being considered locally to deal with water quality and early warning systems within a smart city framework. In this year's event, we will have a demonstration of an AI application developed by a local start-up company based in Cleveland, Ohio.

Our panel this year will include:

- **Doug McCollough**, CIO of Dublin Ohio, Dublin International Institute for the Study of Intelligent Communities and the Smart Mobility Corridor, geared to bring self-driving cars to Ohio,
- **Andrew Konya**, co-founder and CEO of Remesh in Cleveland, developer of Artificial Intelligence that engages and understands large groups of people,
- **Paul Carlson**, Intelligent Community Strategist from Columbus Ohio,
- **Dr. Mark Vopat**, Technology Ethicist at Youngstown State University, and
- **Dr. Jay Ramanathan**, Humanitarian Engineering Center of Ohio State University.

Cameron Hughes is a computer and robot programmer. He holds a post as 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., and a staff Programmer/Analyst at Youngstown State University. Cameron is an advisory board member for the National Robotics Education Foundation and a member of the Oak Hill Robotics Makerspace. He is the project leader of the technical team for the NEOACM CSI/CLUE Robotics Challenge and regularly directs robot programming workshops. Among other books, Cameron is the co-author of *Build Your Own Teams of Robots* (2013) and *Robot Programming: A Guide to Controlling Autonomous Robots* (2016).

Situated Natural Language Interaction in Uncertain and Open Worlds

Tom Williams (Tufts University; williams@cs.tufts.edu)
DOI: 10.1145/3098888.3098896

As intelligent robots become integrated into society, it becomes important for them to be capable of natural, human-like human-robot interaction (HRI). While there has been some progress on enabling natural-language based HRI (Mavridis, 2015), most natural language enabled robots rely on highly scripted interactions, keyword spotting, and shallow natural language processing techniques. For many applications, these methods may be sufficient to achieve the desired behavior, which may be restricted to a small class of tasks. Such methods, however, are not helpful for the development of robots that are generally and flexibly taskable, that can learn about new entities and concepts on the fly, and that are capable of engaging in truly natural *human-like* HRI.

What is more, even natural-language enabled robots designed to handle more natural, flexible dialogue typically operate under a set of assumptions that severely restrict the types of language they are prepared to handle. Specifically, many language-enabled robots assume that (1) their knowledge is certain, (2) they operate in a closed world, (3) only entities from a single domain will be referred to, (4) knowledge is centralized, and homogeneous in representation, (5) humans' utterances should be understood as commands or requests, (6) humans' utterances will be expressed directly, and/or that (7) the meaning of humans' utterances will not vary with context.

To advance the state of the art of natural language based HRI, we must develop natural language enabled robots that challenge these assumptions, that is, robots which are able to (1) handle uncertain and open worlds; (2) make use of distributed knowledge that is heterogeneous in domain and in representation; (3) process a wide variety of utterance forms and referring expression forms; and (4) process such utterances in a context sensitive manner.

In this dissertation, I describe algorithms I have developed in service of these goals, and the experimental and theoretical work I have performed which informs those algorithms and mechanisms.

Figure 1: The *Vulcan* Intelligent Wheelchair: one of the robot platforms used in the presented work.

I first present a set of algorithms that provide **reference resolution** and **referring expression generation** capabilities: *SPEX*, the *Spatial Expert*, an architectural component responsible for performing spatial reference resolution in open worlds (Williams, Cantrell, Briggs, Schermerhorn, & Scheutz, 2013); *REX*, the *Referential Executive*, an architectural component responsible for a broader class of referential activities, including domain-independent reference resolution of definite noun phrases in uncertain and open worlds (Williams & Scheutz, 2015a,b, 2016a); *GH-POWER*, an algorithm which incorporates *REX* into a broader *Givenness Hierarchy*-theoretic (Gundel, Hedberg, & Zacharski, 1993) framework in order to additionally resolve anaphoric and deictic expressions in a context sensitive manner (Williams, Acharya, Schreitter, & Scheutz, 2016; Williams & Scheutz, 2017); and *PIA*, an algorithm which uses *REX* for the purposes of referring expression *generation*.

Next, I move on to discuss **pragmatic reasoning**. I begin by presenting experimental evidence demonstrating the extent of *indirect speech act* use in HRI (Briggs, Williams, & Scheutz, 2017), and then present a Dempster-Shafer theoretic framework for both understanding and generating indirect speech acts in a context sensitive manner under uncertainty and ignorance (Williams, Briggs, Oosterveld, & Scheutz, 2015). Next, I demonstrate how this framework can be used to generate clarification requests to resolve pragmatic and referential ambiguity (Williams & Scheutz, 2016b). Finally I move beyond the pragmatics of *human*-robot communication, and discuss the pragmatics of *robot*-robot communication (Williams, Briggs, & Scheutz, 2015).

Finally, I discuss the application of the presented algorithms to assistive robotics, by providing a comprehensive survey of natural language enabled wheelchairs, and then demonstrating how the use of the presented algorithms on the University of Michigan's Vulcan intelligent wheelchair (Figure 1) (Murarka, Gulati, Beeson, & Kuipers, 2009) advances the state of the art of such wheelchairs (Williams, Johnson, Scheutz, & Kuipers, 2017).

Acknowledgments

This doctoral work was funded in part by ONR grants #N00014-11-1-0289, #N00014-11-1-0493, #N00014-10-1-0140 #N00014-14-1-0144, #N00014-14-1-0149, #N00014-14-1-0751, and NSF grants #1111323, #1038257.

References

Briggs, G., Williams, T., & Scheutz, M. (2017). Enabling robots to understand indirect speech acts in task-based interactions. *Journal of HRI*.

Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 274–307.

Mavridis, N. (2015). A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems*, *63*, 22–35.

Murarka, A., Gulati, S., Beeson, P., & Kuipers, B. (2009). Towards a safe, low-cost, intelligent wheelchair. In *PPNIV.*

Williams, T., Acharya, S., Schreitter, S., & Scheutz, M. (2016). Situated open world reference resolution for human-robot dialogue. In *ACM/IEEE HRI.*

Williams, T., Briggs, G., Oosterveld, B., & Scheutz, M. (2015). Going beyond command-based instructions: Extending robotic natural language interaction capabilities. In *AAAI.*

Williams, T., Briggs, P., & Scheutz, M. (2015). Covert robot-robot communication: Human perceptions and implications for human-robot interaction. *Journal of HRI*.

Williams, T., Cantrell, R., Briggs, G., Schermerhorn, P., & Scheutz, M. (2013). Grounding natural language references to unvisited and hypothetical locations. In *AAAI.*

Williams, T., Johnson, C., Scheutz, M., & Kuipers, B. (2017). A tale of two architectures: A dual-citizenship integration of natural language and the cognitive map. In *AAMAS.*

Williams, T., & Scheutz, M. (2015a). A domain-independent model of open-world reference resolution. In *COGSCI.*

Williams, T., & Scheutz, M. (2015b). POWER: A domain-independent algorithm for probabilistic, open-world entity resolution. In *IEEE/RSJ IROS.*

Williams, T., & Scheutz, M. (2016a). A framework for resolving open-world referential expressions in distributed heterogeneous knowledge bases. In *AAAI.*

Williams, T., & Scheutz, M. (2016b). Resolution of referential ambiguity using Dempster-Shafer theoretic pragmatics. In *AI-HRI.*

Williams, T., & Scheutz, M. (2017). Reference resolution in robotics: A givenness hierarchy theoretic approach. In J. Gundel & B. Abbott (Eds.), *The Oxford handbook of reference.*

Tom Williams completed a joint PhD in Computer Science and Cognitive Science in May 2017 with Professor Matthias Scheutz at Tufts University, and will be joining Colorado School of Mines in August 2017 as an Assistant Professor of Computer Science. Tom's research focuses on allowing robots to communicate in natural language in uncertain and open worlds, with applications to assistive and search-and-rescue robotics.

Populating a Linked Data Entity Name System

Mayank Kejriwal (Information Sciences Institute; kejriwal@isi.edu)

DOI: 10.1145/3098888.3098897

Resource Description Framework (RDF) is a graph-based data model used to publish data as a *Web of Linked Data* (Bizer *et al.* 2009). RDF is an emergent foundation for large-scale *data integration*, the problem of providing a unified view over multiple data sources. The structure in RDF data can be conveniently visualized using *directed labeled graphs*, as illustrated in the real-world graph fragments in Figure 1. Nodes in the graph represent entities (e.g. the node with label *dbpedia:Allen_Paul* represents the entity Paul Allen in the DBpedia knowledge graph) and edges represent either attributes of an entity (e.g. '01/21/1953' is the birthdate of Paul Allen) or relationships between two entities (e.g. Paul Allen is the co-founder of the company entity, Microsoft). Facts in the knowledge base are formally represented as a set of *triples*, with a triple comprising a labeled edge (denoted as a *property*) in the RDF graph along with its incoming and outgoing nodes.

An Entity Name System (ENS) is a thesaurus for entities, and is a crucial component in a data integration architecture (Kejriwal 2014). For example, consider an application that queries multiple knowledge graphs. Since entities like Microsoft and Paul Allen are represented in different ways in different graphs, a robust system would need to infer that the different *mentions* of a single entity should be linked to a *canonical* thesaurus entry. Automatically populating an ENS is equivalent to solving an AI problem called *Entity Resolution* (ER), which concerns identifying pairs of entities referring to the same underlying entity (Getoor and Machanavajjhala 2012).

Due to its Web origins, Linked Data exhibits properties that make ER a challenging problem including *scale* (super-linear growth on average since inception), *heterogeneity* (data is published using a wide range of RDF types and properties) and *diversity* (data spans multiple domains, e.g. social media and bioinformatics). Diversity, in particular, is an important concern for state-of-the-art *supervised*

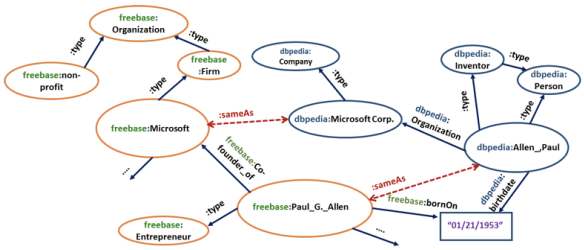

Figure 1: Fragments of real-world RDF graphs, DBpedia and Freebase, with two examples of *Entity Resolution* using a special *sameAs* property

machine learning-based ER systems, since manually labeled examples for each domain are typically not available. Acquiring such examples is also hard since ER exhibits *class skew*: the number of synonym node pairs in a graph is much smaller than the number of non-synonym node pairs.

Contributions

We propose an Entity Resolution (ER) approach that includes methods to accommodate the competing requirements of diversity, heterogeneity and scale that are necessary for enabling successful population of Linked Data Entity Name Systems. In the rest of this section, we present a high-level overview of the primary aspects of the approach, followed by a brief empirical summary.

Given the success of supervised machine learning approaches in the ER literature, we address diversity by developing a *self-training* machine learning algorithm that executes in two stages (Kejriwal and Miranker 2013). In the first stage, the algorithm generates its own training set by applying a set of inexpensive, non-adaptive similarity heuristics that are known to be robust in a variety of text domains. Because these heuristics are non-adaptive, they do not require user supervision other than the setting of a conservative threshold, which we also automated in later developments (Kejriwal and Miranker 2015b).

Although such heuristics cannot compete with supervised algorithms in 'difficult' real-world domains, we showed that they can be used to obtain 'easy' matches between nodes that have many attribute values in common (Kejriwal and Miranker 2013). Even with conservative thresholds, some of the matches are noisy and the obtained matches are quite sparse. In the second stage, therefore, the algorithm uses a number of strategies to robustly train a supervised machine learning classifier like an SVM. Strategies that were shown to achieve high success rates in our work include the use of an expressive set of features (spanning numeric, string and phonetic features), boosting and feature discretization (Kejriwal and Miranker 2015a).

To accommodate heterogeneity, we developed a *schema-matching algorithm* that not only matches *edge labels* between two graphs, but also determines *compatibly typed* nodes across the graphs. In Figure 1, for example, the algorithm would determine that *freebase:Microsoft* and *dbpedia:Microsoft Corp.* are compatibly typed (i.e. they are both companies) and should be classified, either as a match or non-match, by the self-trained SVM, whereas *freebase:Microsoft* and *dbpedia:Allen_,Paul* are incompatibly typed.

Finally, we accommodate scale *directly* in the design of all the algorithms described above. We show that an approximate version of the algorithms, particularly the training set generator, can be implemented in a shared-nothing parallel paradigm like MapReduce, and successfully handles *data skew*, an important concern in many data-intensive problems.

Empirically, we evaluate the efficacy of our methods by actively avoiding prior assumptions about input domains, and through evaluations on ten RDF test cases spanning multiple domains, including movies, books, people and restaurants. On all datasets, the approach outperforms a popular unsupervised baseline (derived from Locality Sensitive Hashing) by a large margin, and achieves performance (within 10% on a popular accuracy measure) competitive with supervised SVM-based baselines (Kejriwal and Miranker 2015b).

We test the scalability of our approach by implementing our algorithms in both serial and MapReduce architectures. Evaluations in public cloud infrastructure (Hadoop clusters on Microsoft Azure) show that the system scales near-linearly, and is able to effectively resolve entities across encyclopedic graphs (e.g. DBpedia and Freebase) with millions of nodes and edges using relatively small clusters ($<$40 compute nodes).

References

Christian Bizer, Tom Heath, and Tim Berners-Lee. Linked data-the story so far. *Semantic Services, Interoperability and Web Applications: Emerging Concepts*, pages 205–227, 2009.

Lise Getoor and Ashwin Machanavajjhala. Entity resolution: theory, practice & open challenges. *Proceedings of the VLDB Endowment*, 5(12):2018–2019, 2012.

Mayank Kejriwal and Daniel P Miranker. An unsupervised algorithm for learning blocking schemes. In *2013 IEEE 13th International Conference on Data Mining*, pages 340–349. IEEE, 2013.

Mayank Kejriwal and Daniel P Miranker. Semi-supervised instance matching using boosted classifiers. In *European Semantic Web Conference*, pages 388–402. Springer, 2015.

Mayank Kejriwal and Daniel P Miranker. An unsupervised instance matcher for schema-free rdf data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 35:102–123, 2015.

Mayank Kejriwal. Populating entity name systems for big data integration. In *International Semantic Web Conference*, pages 521–528. Springer, 2014.

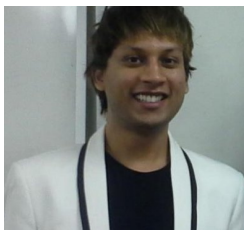

Mayank Kejriwal completed his PhD at the University of Texas at Austin under Daniel P. Miranker. Currently, he works as a researcher and Computer Scientist in the Information Integration group at the Information Sciences Institute, part of the Viterbi School of Engineering at the University of Southern California in Los Angeles, California.

Incremental and Developmental Perspectives for General-Purpose Learning Systems

Fernando Martínez-Plumed (Universitat Politècnica de València; fmartinez@dsic.upv.es)
DOI: 10.1145/3098888.3098898

The stupefying success of Artificial Intelligence (AI) for *specific* problems, from recommender systems to self-driving cars, has not yet been matched with a similar progress in *general* AI systems, coping with a variety of (different) problems. This dissertation deals with **the long-standing problem of creating more general AI systems, through the analysis of their development and the evaluation of their cognitive abilities**.

Given the above challenge, in the presented dissertation we characterise a series of human intelligence attributes (incremental, developmental and lifelong learning) and cognitive-oriented procedures (memory and forgetting) that, combined with the use of symbolic AI and symbolic learning, have helped us to develop both a general-purpose learning approach as well as a knowledge handling tool. This ambitious issue should, furthermore, pervade the evaluation procedures in AI where systems are usually evaluated in terms of task performance, not really in terms of intelligence (ability-oriented evaluation).

Particularly, and regarding the construction of more general AI approaches, this thesis contributes with a pair of settings for learning and knowledge acquisition. Firstly we present a general-purpose declarative learning system (**gErl**) (Martínez-Plumed, Ferri, Hernández-Orallo, & Ramírez-Quintana, 2013, 2017) that meets several desirable characteristics in terms of expressiveness, comprehensibility and versatility. **gErl** (Fig. 1) relies on two compatible mechanisms. The former is the definition of customised learning operators, depending on the data structures and problem at hand, done by the user, using a functional language. The latter mechanism is the use of generalised heuristics, since the use of different operators precludes the system from using specialised heuristics for each of them. The choice of the right pair of operator and rule has been reframed as a decision process (using a reinforcement learning approach). Therefore, not only is this a novel approach, but also allows us to better understand the role of operators and heuristics in machine learning. By performing a series of illustrative experiments we show where the flexibility stands out, since **gErl** is able to solve a wide range of problems (from recursive ones to several IQ tests).

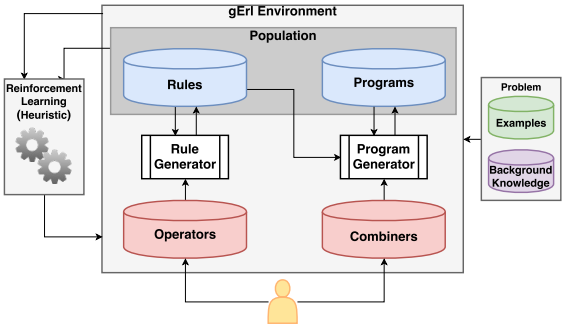

Figure 1: **gErl** takes examples and learning operators as input and returns functional programs.

Secondly, the learning process is also overhauled with a new developmental and lifelong approach for knowledge acquisition, consolidation and forgetting, which is necessary when bounded resources (memory and time) are considered. In this sense we present a parametrisable (hierarchical) approach (Martínez-Plumed, Ferri, Hernández-Orallo, & Ramírez-Quintana, 2015) for structuring knowledge (based on coverage) which is able to check whether the new learnt knowledge can be considered redundant, irrelevant or inconsistent with the old one, and whether it may be built upon previously acquired knowledge. We show that the use of complex knowledge assessment structures jointly with information theory-based principles to characterise knowledge (Fig. 2) allows for a straightforward and principled approach to knowledge handling.

Thirdly, and moving towards AI evaluation, this thesis analyzes whether the use of more ability-oriented evaluation techniques for AI (such as intelligence tests) is a much better al-

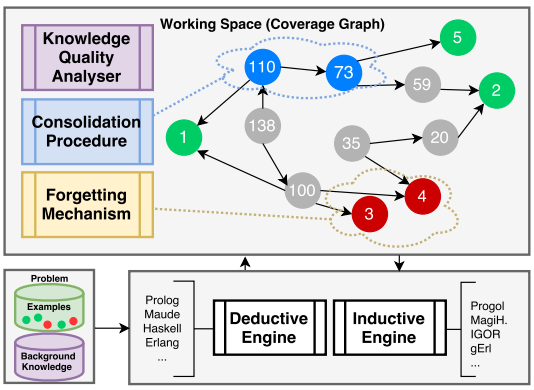

Figure 2: Organisation of complex knowledge structures in terms of coverage and information theory principles.

ternative to most task-oriented evaluation approaches in AI. Accordingly, we make a review of what has been done when AI systems have been confronted against tasks taken from intelligence tests (Hernández-Orallo, Martínez-Plumed, Schmid, Siebers, & Dowe, 2016). In this regard, we scrutinise what intelligence tests measure in machines, whether they are useful to evaluate AI systems, whether they are really challenging problems, and whether they are useful to understand (human) intelligence by analysing over 30 papers featuring AI systems addressing intelligence test problems. Our aim here is to contribute to a more widespread realisation that more general classes of problems are needed when constructing benchmarks for AI evaluation.

By the same token, as a final contribution, we show that intelligence tests can also be useful to examine concept dependencies (mental operational constructs) in the cognitive development of artificial systems (although a superficial score comparison is misleading), therefore supporting the assumption that, even for fluid intelligence tests, the difficult items require a more advanced cognitive development than the simpler ones. In this sense, we show (Martínez-Plumed et al., 2017) how several fluid intelligence test problems (odd-one-out problems, Raven's Progressive Matrices and Thurstone's letter series) are addressed by our general-purpose learning system **gErl**, which, although it is not particularly designed on purpose to solve intelligence tests, is able to perform relatively well for this kind of tests. **gErl** makes it explicitly how complex each pattern is and what operators are used for each problem (due its symbolic and declarative nature), thus providing useful insight into the characteristics and usefulness of these tests when assessing the abilities and cognitive development of AI systems.

Summing up, this dissertation represents one step forward in the hard and long pursuit of making more general AI systems and fostering less customary (and challenging) ability-oriented evaluation approach.

References

Hernández-Orallo, J., Martínez-Plumed, F., Schmid, U., Siebers, M., & Dowe, D. L. (2016). Computer models solving intelligence test problems: Progress and implications. *Artificial Intelligence*, *230*, 74-107.

Martínez-Plumed, F., Ferri, C., Hernández-Orallo, J., & Ramírez-Quintana, M. J. (2013). Learning with configurable operators and rl-based heuristics. In A. Appice (Ed.), *New frontiers in mining complex patterns* (Vol. 7765, p. 1-16). Springer.

Martínez-Plumed, F., Ferri, C., Hernández-Orallo, J., & Ramírez-Quintana, M. J. (2015). Knowledge acquisition with forgetting: an incremental and developmental setting. *Adaptive Behavior*, *23*(5), 283-299.

Martínez-Plumed, F., Ferri, C., Hernández-Orallo, J., & Ramírez-Quintana, M. J. (2017). A computational analysis of general intelligence tests for evaluating cognitive development. *submitted (second revision)*.

Fernando Martínez-Plumed is currently a post-doctoral researcher in Artificial Intelligence at the Technical University of Valencia (UPV), Spain, where he obtained his Ph.D. His main research interest encompasses several areas of Artificial Intelligence, machine learning, inductive programming, psychometrics and cognitive systems.

A Persuasive Virtual Chat Agent Based on Sociolinguistic Theories of Influence

Samira Shaikh (UNC Charlotte; sshaikh2@uncc.edu)
DOI: 10.1145/3098888.3098899

Problem Statement

Machine intelligence breaks new ground with each passing day, achieving feats like winning against humans in the game of Go and Jeopardy!. However, the goal of implementing social competence in machines so that they may achieve specific social goals remains elusive. Persuasion, the ability to tailor messaging to the target audience is closely related to social competence (Gass and Seiter, 2015). In using technology to persuade, we delegate a central function of human communication to a computer. This central function is intentional and goal-directed communication - deliberate, purposeful messaging with a clear outcome in mind. Our goal in this thesis is to automate the very process of persuasive communication, by designing a system which can purposefully communicate, without any restrictions on domain or genre or task, and which has the clear intention of persuading the recipients of its messaging.

We investigate how models of social phenomena - specifically persuasion strategies - may be automated in an artificial autonomous agent, in the form of two overarching research questions:

- Can specific persuasive strategies be automated in a virtual chat agent?
- Can active persuasion by individuals during conversation be detected and counteracted by such an agent?.

Our goal is not to create an artificial agent capable of passing the Turing test or the Loebner prize. Rather, our goal is to define specific human persuasive strategies that can be programmed into an agent who can then persuade participants to its own view.

Methodology

The foundation of our persuasive strategies comes from the summative model of attitude (a well-established model of attitude in social psychology), where belief change leads to attitude change, and, ultimately, behavior change (Fishbein and Ajzen, 2011).

We organized the work in three phases. First, we conducted a belief elicitation study to obtain salient beliefs on a variety of social issues (for example, *Should the minimum legal drinking age be lowered from 21 to 18?*). We had two distinct reasons to undertake such a study. The first motivation was to use the beliefs so elicited to design survey instruments. These surveys are crucial to our work, because the responses on the surveys provide the ground truth from which to measure the agent's efficacy in being persuasive in controlled experiments. Our second motivation was to use the corpus of beliefs elicited as a database of natural language statements for the agent to use. The goal of this thesis was not to undertake the task of natural language generation for the agent. Our solution is to use the responses made by humans on the topics as a database of arguments made in favor of and against the topic of discussion. We pre-programmed these arguments in the agent, with the intention of using them at *opportune* moments in the conversation (as explained below).

Next, we programmed behaviors and strategies in the agent that were aimed at persuading individuals through online conversation as well as counteracting persuasion by the participants. The behaviors programmed in the agents are triggered, in part, by a variety of linguistic cues emerging from the conversation, such as dialogue acts, topic, polarity and communication acts. The annotated context of conversation is used to inform the agent's models by updating the underlying beliefs of participants in real time. It is necessary for the agent to create and maintain a represen-

tation of the mental states of the participants with respect to the topic so as to understand their viewpoints. In this work, the mental state or point of view is the overall attitude towards the topic, and its essential pieces are the belief strength (how strongly a belief is held) and belief evaluation (is the evaluation towards the belief positive or negative) on the beliefs related to the topic of discussion. The agent uses the belief models and annotated utterances and selects appropriate behaviors to perform from a list of pre-determined behaviors (operationalized as dialogue acts e.g. assertions, agreements, disagreements, etc).

In the third phase, we ran controlled experiments. The aim of the these experiments was to deploy the agent and validate our persuasion and counter-persuasion strategies in online synchronous conversation environments. We selected the majority-minority influence setting in our experiment design. In social influence research, it has been shown that minorities influence people's thinking, attitudes, and behavior by being consistent in their views and flexible in their negotiation with majority members' (Gardikiotis, 2011). Each chat session in our experiments consisted of four participants, two majority opinion holders, one minority opinion holder and the agent (either wizard or computer agent, who advocated the extreme minority opinion). Procedures for both the Woz and autonomous agent experiments were nearly identical. The only difference was the presence of the wizard or the algorithm acting autonomously. Our main research hypothesis is that the interventions made by the agent would result in attitude change in the participants. Pre- and post-discussion surveys allow us to measure changes in participants belief models, and thus, the shifts in their overall attitude towards the topic of discussion. The system goal in our experiments is multifaceted, attempting to change the belief strength and belief evaluation, and furthermore, the overall attitude of the participants. Accordingly, we used the Wilcoxon signed-rank test (Wilcoxon, 1945) to measure persuasion effectiveness of the algorithm. We observed statistically significant changes in overall attitude from pre- and post-discussion surveys across 10 discussion sessions with 30 participants in total using the Wilcoxon signed rank test *(Wilcoxon T=12, p=0.03)*.

Contributions of this Thesis

Upon placing an agent in the midst of conversations, it is able to discern beliefs that are expressed by the participants in the group, and use them to ascertain participant's opinions on topics of discussion. Using this information and drawing upon theories of influence and persuasion from social psychology, cognitive science and communication, the agent aligns participants towards or against a particular issue. We validated that the agent achieved statistically significant changes in the participant's attitudes, thus demonstrating its effectiveness in being persuasive. In doing so, this work makes contributions to the field of AI and human-computer dialogue.

Acknowledgments

This research was supported by a grant from Army Research Laboratory.

References

Fishbein, M. and Ajzen, I. (2011). *Predicting and changing behavior: The reasoned action approach*. Taylor & Francis.

Gardikiotis, A. (2011). Minority influence. *Social and Personality Psychology Compass*, 5(9):679–693.

Gass, R. H. and Seiter, J. S. (2015). *Persuasion: Social influence and compliance gaining*. Routledge.

Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83.

Samira Shaikh works on computational sociolinguistics, artificial intelligence and natural language processing. She obtained her PhD from the State University of New York - University at Albany in Computer Science in July 2016. She is currently an Assistant Professor of Cognitive Science in the Department of Computer Science department at University of North Carolina - Charlotte.