AI Matters

Annotated Table of Contents

Upcoming Conferences Sponsored by SIGAI

Registration discount for SIGAI members.

HRI 2017: The 2017 Conference on Human-Robot Interaction, Vienna, Austria.
Conference dates: March 6–9, 2017

IUI 2017: The 22nd Annual Meeting of the Intelligent User Interfaces Community, Limassol, Cyprus.
Conference dates: March 13–16, 2017

ICAIL 2017: The 16th International Conference on Artificial Intelligence and Law, London, UK.
Conference dates: June 12–16, 2017

ICEIS 2017: The 19th International Conference on Enterprise Information Systems, Porto, Portugal.
Conference dates: April 26–29, 2017

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: 10.1145/3054837

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 1

Eric Eaton, Co-Editor (University of Pennsylvania; aimatters@sigai.acm.org)
Amy McGovern, Co-Editor (University of Oklahoma; aimatters@sigai.acm.org)
DOI: 10.1145/3054837.3054838

Wow! It's now the third year of *AI Matters*! Time sure flies! As we start the third volume of *AI Matters*, we are continuing the rollout of new features. Over the past few issues, you may recall that we introduced recurring columns on AI Education, AI Amusements, and Profiles of people in AI. These columns are joined by two new columns in this issue: AI Buzzwords Explained, and AI Policy.

The new column **AI Buzzwords Explained** will focus each issue on introducing readers to a current topic in AI, written by an expert in the field. In this issue, the column focuses on the topic of "scientific workflows."

If you visited the new AI Matters blog (http://sigai.acm.org/ai-matters/) recently, you may have noticed the posts on current policy issues relating to AI by Larry Medsker, our ACM SIGAI Public Policy Officer. Some recent posts include the Stanford *One Hundred Year Study on Artificial Intelligence*, *AI and Future Employment*, and the interview with Jack Clark of OpenAI on *The Public Policy Implications of AI*. We will be re-printing select posts in *AI Matters* in the new **AI Policy** column, but for the most current updates, be sure to see the AI Matters blog.

In the **AI Education** column this issue, Todd Neller describes open-access resources for AI Education. Although you may be familiar with many of these resources, hopefully it will introduce you to some new resources that you can use to further your own education or recommend to students studying AI.

For the second instance of **AI Spotlight**, we interview Jim Kurose, Assistant Director of the National Science Foundation (NSF) for the Directorate for Computer and Information Science and Engineering (CISE). 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.

This issue includes the final set of **abstracts of recent AI doctoral theses**, continuing this feature from our previous issue. Although we've finished the special issues on doctoral abstracts, we invite students to continue to submit abstracts to appear in future issues.

We've made a few other minor changes to bring us into volume 3 of *AI Matters*, including new icons to help you identify the topic of each article. We have also synchronized the volume with the calendar year, and introduced many improvements behind the scenes to help us continue to scale.

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.

AI Buzzwords Explained: Scientific Workflows

Daniel Garijo (Information Sciences Institute and Department of Computer Science, University of Southern California; dgarijo@isi.edu)

DOI: 10.1145/3054837.3054839

The reproducibility of scientific experiments is crucial for corroborating, consolidating and reusing new scientific discoveries. However, the constant pressure for publishing results (Fanelli, 2010) has removed reproducibility from the agenda of many researchers: in a recent survey published in Nature (with more than 1500 scientists) over 70% of the participants recognize to have failed to reproduce the work from another colleague at some point in time (Baker, 2016). Analyses from psychology and cancer biology show reproducibility rates below 40% and 10% respectively (Collaboration, 2015) (Begley & Lee, 2012). As a consequence, retractions of publications have occurred in the last years in several disciplines (Marcus & Oransky, 2014) (Rockoff, 2015), and the general public is now skeptical about scientific studies on topics like pesticides, depression drugs or flu pandemics (American, 2010).

Reproducing the results of a previous study can be a challenge, as even when the original datasets and end results are available, a significant investment in time may be required (Garijo et al., 2013). Fortunately, the community has started to pay attention to initiatives for preserving the data and software used in scientific publications (e.g., Zenodo,[1] Github,[2] etc.). In computational sciences, **scientific workflows** were proposed in the last decade as a means to address reproducibility. A scientific workflow defines the set of computational tasks and dependencies needed to carry out in silico experiments (Taylor, Deelman, Gannon, & Shields, 2006). Typically, scientific workflows are represented as directed graphs, where the nodes represent computational tasks and the edges represent their dependencies. Figure 1 shows an example with two workflows, one for text analytics on the left and another one for neuro-image analysis on the right.

Scientific workflows have been used in many domains, including astronomy [10], brain image analysis (Dinov et al., 2009) and bioinformatics (Wolstencroft et al., 2013). Besides improving reproducibility, scientific workflows have also proved to be helpful in teaching new users to visualize the overall structure of a method, save time when reusing an existing method and debug or inspect and modularize scientific experiments (Goderis, 2008; Garijo et al., 2014).

There are many challenges associated to scientific workflows. During the last decade plenty of systems have been designed to efficiently represent and execute them in both local and distributed environments (e.g., (Wolstencroft et al., 2013; Gil et al., 2011; Deelman et al., 2004; Callahan et al., 2006; Ludscher et al., 2006; Filgueira et al., 2014; Giardine et al., 2005), etc.). Different approaches have focused in optimizing workflow execution (e.g., (Deelman et al., 2004)) and their results (e.g., (Holl, 2014)). Other works have addressed workflow reuse (Garijo et al., 2014), (Goderis, Sattler, Lord, & Goble, 2005), recommendation (Starlinger, Brancotte, Cohen-Boulakia, & Leser, 2014; Bergmann & Gil, 2014) and discovery (Goderis, 2008), (Bergmann & Gil, 2014), as building on previous findings is considered to be critical to push science forward. Here we overview those aspects of workflows related to reproducibility, i.e., workflow preservation, traceability of the results and workflow sharing.

There are two ways in which a workflow may be preserved. The first way is by documenting the method captured by the workflow itself, i.e., providing enough details on each of the tasks of the workflow for anyone to be able to understand their functionality (Garijo & Gil, 2011; Belhajjame et al., 2015). The rationale is simple: given the pace at which software and data evolve, it is difficult to ensure that within five, ten or twenty years the whole workflow will still be reusable. This

[1] https://zenodo.org/

[2] http://github.com/

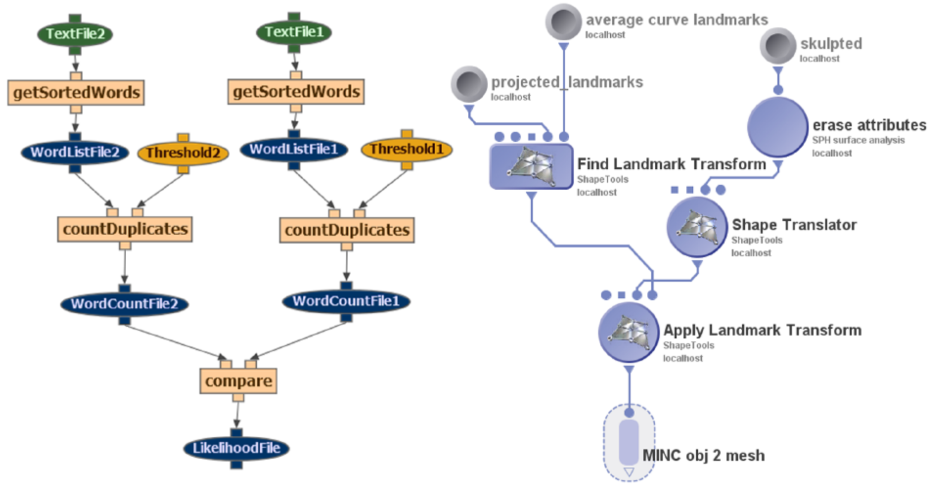

Figure 1: Two scientific workflows from two different workflow systems. The one on the left represents tasks as rectangles and data with ovals, while the one on the right represents task in blue and inputs in grey.

is common in domains where scientific workflows rely on external web services and evolving community-built datasets (e.g., the Protein Data Bank[3] in bioinformatics). New releases of software, changes to the existing APIs or new data discoveries may supersede existing resources, making them outdated and sometimes incompatible with the rest of the tasks in the workflow. Therefore, documentation approaches tend to contextualize, describe and generalize the functionality of every dataset and task used in the workflow. Documentation approaches are usually complemented with sample data, pointing to archived versions of the software to facilitate understanding the original method. Another key feature of these approaches includes documenting the provenance of the results of a workflow. The provenance of a result aims to capture its creation process, i.e., all the steps that contributed to its outcome, including the original datasets and intermediate data. A provenance record also attributes credit to the scientists responsible for producing the result. There is a standard model for provenance publishing on the web (Lebo et al., 2013), and related work has extended it to publish scientific workflow metadata[4] (Garijo & Gil, 2011; Belhajjame et al., 2015; Missier, Dey, Belhajjame, Cuevas-Vicenttn, & Ludscher, 2013). Once a workflow is documented, it may be included as part of a repository (Roure, Goble, & Stevens, 2009; Mates, Santos, Freire, & Silva, 2011; Belhajjame et al., 2013) for others to reuse.

The second way to preserve workflows is by capturing their functionality in containers (e.g., Docker[5]) or virtual machines. This way the workflow becomes a black box that performs the experiment functionality, including inputs, software and dependencies for execution. The challenge relies in the creation process of such containers. Approaches like (Chirigati, Shasha, & Freire, 2013) monitor the execution of the experiment to create a virtual machine, while approaches like (Santana-Prez & Prez-Hernndez, 2015) depend on the authors to document the infrastructure details for the workflow. Recent work has proposed a more flexible approach, capturing each of the steps of the workflow as an independent container (Qasha, Cala, & Watson, 2016). Finally notebooks[6] are gaining a lot of momentum as an alternative lightweight method to encapsulate and test script based experiments.

Scientific workflows have demonstrated to be useful to re-execute, reuse and share the methods and tasks commonly used in a community (Garijo et al., 2014). Workflows should be treated as first class citizens in cyberinfrastructure (Gil et al., 2007), since they provide the means of transparent and reproducible

[3]http://www.rcsb.org/pdb/home/home.do
[4]http://vcvcomputing.com/provone/provone.html
[5]https://www.docker.com/
[6]http://jupyter.org/

work. There are still open challenges in workflows, and venues like eScience[7] and Super Computing[8] discuss and publish new research every year.

References

American, S. (2010, October). In Science We Trust: Poll Results on How you Feel about Science. *Scientific American*.

Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, *533*(7604). Retrieved from http://www.nature.com/news/1-500-scientists-lift-the-lid-on-reproducibility-1.19970

Begley, C., & Lee, M. (2012, March). Drug development: Raise standards for preclinical cancer research. , *483*, 531–533. doi: 10.1038/483531a

Belhajjame, K., Zhao, J., Garijo, D., Gamble, M., Hettne, K., Palma, R., ... Goble, C. (2015). Using a suite of ontologies for preserving workflow-centric Research Objects. *Web Semantics: Science, Services and Agents on the World Wide Web*.

Belhajjame, K., Zhao, J., Garijo, D., Garrido, A., Soiland-Reyes, S., Alper, P., & Corcho, O. (2013). A workflow PROV-corpus based on Taverna and Wings. In *Proceedings of the Joint EDBT/ICDT 2013 Workshops* (pp. 331–332). Genoa, Italy: ACM. Retrieved from http://doi.acm.org/10.1145/2457317.2457376 doi: 10.1145/2457317.2457376

Bergmann, R., & Gil, Y. (2014). Similarity assessment and efficient retrieval of semantic workflows. *Information Systems*, *40*, 115 – 127. Retrieved from http://www.sciencedirect.com/science/article/pii/S0306437912001020 doi: http://dx.doi.org/10.1016/j.is.2012.07.005

Callahan, S. P., Freire, J., Santos, E., Scheidegger, C. E., Silva, C. T., & Vo, H. T. (2006). Vistrails: Visualization meets data management. In *In ACM SIGMOD* (pp. 745–747). ACM Press.

Chirigati, F., Shasha, D., & Freire, J. (2013). ReproZip: Using Provenance to Support Computational Reproducibility. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance* (pp. 1:1–1:4). Lombard, Illinois: USENIX Association. Retrieved from http://dl.acm.org/citation.cfm?id=2482949.2482951

Collaboration, O. S. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251). doi: 10.1126/science.aac4716

Deelman, E., Blythe, J., Gil, Y., Kesselman, C., Mehta, G., Patil, S., ... Livny, M. (2004). Pegasus: Mapping Scientific Workflows onto the Grid. In M. Dikaiakos (Ed.), *Grid Computing* (Vol. 3165, pp. 11–20). Springer Berlin / Heidelberg.

Dinov, I. D., Horn, J. D. V., Lozev, K. M., Magsipoc, R., Petrosyan, P., Liu, Z., ... Toga, A. W. (2009). Efficient, Distributed and Interactive Neuroimaging Data Analysis Using the LONI Pipeline. In *Frontiers in Neuroinformatics* (Vol. 3). doi: 10.3389/neuro.11.022.2009

Fanelli, D. (2010). Do Pressures to Publish Increase Scientists Bias? An Empirical Support from US States Data. *PLoS ONE*, *5*(4). doi: 10.1371/journal.pone.0010271.

Filgueira, R., Atkinson, M., Bell, A., Main, I., Boon, S., Kilburn, C., & Meredith, P. (2014). eScience Gateway Stimulating Collaboration in Rock Physics and Volcanology. In *e-Science (e-Science), 2014 IEEE 10th International Conference on* (Vol. 1, pp. 187–195). IEEE.

Garijo, D., Corcho, O., Gil, Y., Braskie, M. N., Hibar, D., Hua, X., ... W.Toga, A. (2014). Workflow Reuse in Practice: A Study of Neuroimaging Pipeline Users. In *10th IEEE International Conference on eScience 2014.*

Garijo, D., & Gil, Y. (2011). A New Approach for Publishing Workflows: Abstractions, Standards, and Linked Data. In *Proceedings of the 6th workshop on Workflows in support of large-scale science* (pp. 47–56). Seattle: ACM. Retrieved from http://dl.acm.org/citation.cfm?id=2110504

Garijo, D., Kinnings, S., Xie, L., Xie, L., Zhang, Y., Bourne, P. E., &

[7]http://escience-2016.idies.jhu.edu/
[8]http://www.supercomp.org/

Gil, Y. (2013). Quantifying Reproducibility in Computational Biology: The Case of the Tuberculosis Drugome. *PLoS ONE*, *8*(11), e80278. Retrieved from http://dx.doi.org/10.1371%2Fjournal.pone.0080278 doi: 10.1371/journal.pone.0080278

Giardine, B., Riemer, C., Hardison, R. C., Burhans, R., Elnitski, L., Shah, P., ... Nekrutenko, A. (2005, October). Galaxy: a platform for interactive large-scale genome analysis. *Genome Research*, *15*(10), 1451–1455.

Gil, Y., Deelman, E., Ellisman, M., Fahringer, T., Fox, G., Gannon, D., ... Myers, J. (2007, December). Examining the Challenges of Scientific Workflows. *Computer*, *40*(12), 24–32. doi: 10.1109/MC.2007.421

Gil, Y., Ratnakar, V., Kim, J., Gonzlez-Calero, P. A., Groth, P. T., Moody, J., & Deelman, E. (2011). Wings: Intelligent Workflow-Based Design of Computational Experiments. *IEEE Intelligent Systems*, *26*(1), 62–72.

Goderis, A. (2008). *Workflow re-use and discovery in Bioinformatics* (Unpublished doctoral dissertation). School of Computer Science, The University of Manchester.

Goderis, A., Sattler, U., Lord, P., & Goble, C. (2005). Seven Bottlenecks to Workflow Reuse and Repurposing. In *The Semantic Web ISWC 2005* (Vol. 3729, pp. 323–337). Springer Berlin Heidelberg. Retrieved from http://dx.doi.org/10.1007/11574620_25

Holl, S. (2014). *Automated Optimization Methods for Scientific Workflows in e-Science Infrastructures* (Doctoral dissertation, University of Bonn). Retrieved from http://hss.ulb.uni-bonn.de/2014/3508/3508.htm

Lebo, T., McGuiness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., ... Zhao, J. (2013, April). *The PROV ontology, W3c Recommendation* (Tech. Rep.). WWW Consortium. Retrieved from http://www.w3.org/TR/2013/REC-prov-o-20130430/

Ludscher, B., Altintas, I., Berkley, C., Higgins, D., Jaeger, E., Jones, M., ... Zhao, Y. (2006). Scientific workflow management and the Kepler system. *Concurrency and Computation: Practice and Experience*, *18*(10), 1039–1065.

Marcus, A., & Oransky, I. (2014, December). Top retractions of 2014. *The Scientist*.

Mates, P., Santos, E., Freire, J., & Silva, C. T. (2011). CrowdLabs: Social Analysis and Visualization for the Sciences. In *23rd International Conference on Scientific and Statistical Database Management (SSDBM)* (pp. 555–564). Springer.

Missier, P., Dey, S., Belhajjame, K., Cuevas-Vicenttn, V., & Ludscher, B. (2013). D-PROV: Extending the PROV Provenance Model with Workflow Structure. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance* (pp. 9:1–9:7). Lombard, Illinois: USENIX Association. Retrieved from http://dl.acm.org/citation.cfm?id=2482949.2482961

Qasha, W., Cala, J., & Watson, P. (2016). A Framework for Scientific Workflow Reproducibility in the Cloud. In *IEEE 12th International Conference on eScience.*

Rockoff, J. D. (2015, September). Amgen Finds Data Falsified in Obesity-Diabetes Study Featuring Grizzly Bears. *The Wall Street Journal*. Retrieved from http://www.wsj.com/articles/amgen-finds-data-falsified-in-obesity-diabetes-study-featuring-grizzly-bears-1441123200

Roure, D. D., Goble, C. A., & Stevens, R. (2009). The design and realisation of the myExperiment Virtual Research Environment for social sharing of workflows. *Future Generation Comp. Syst.*, *25*(5), 561–567.

Santana-Prez, I., & Prez-Hernndez, M. (2015). Towards Reproducibility in Scientific Workflows: An Infrastructure-Based Approach. *Scientific Programming*, *2015*, 11. doi: 10.1155/2015/243180

Starlinger, J., Brancotte, B., Cohen-Boulakia, S., & Leser, U. (2014). Similarity Search for Scientific Workflows. *Proceedings of the VLDB Endowment*, *7*(12), 1143–1154.

Taylor, I. J., Deelman, E., Gannon, D. B., & Shields, M. (2006). *Workflows for e-Science: Scientific Workflows for Grids*. Secaucus, NJ, USA: Springer-

Verlag New York, Inc.
Wolstencroft, K., Haines, R., Fellows, D., Williams, A., Withers, D., Owen, S., ... Goble, C. (2013). The Taverna workflow suite: designing and executing workflows of Web Services on the desktop, web or in the cloud. *Nucleic Acids Research*.

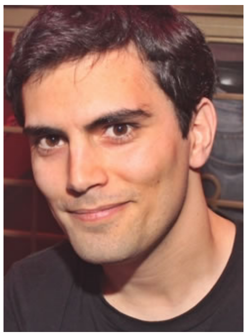

Daniel Garijo is a post-doctoral researcher at the Information Sciences Institute of the University of Southern California. His research activities focus on e-Science and the Semantic web, specifically on how to increase the understandability of scientific workflows using provenance, metadata, intermediate results and Linked Data.

AI Policy: Organizations, Resources, and Recent Symposia

Larry Medsker (George Washington University; lrm@gwu.edu)
DOI: 10.1145/3054837.3054840

Abstract

The AI Matters Policy column is a new column that is appearing twice a month in the AI Matters blog (https://sigai.acm.org/aimatters/blog/). Selected posts will be summarized in each issue of *AI Matters*.

Introduction

As I start my term as ACM SIGAI Public Policy Officer, I have a few initial goals:

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

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

Organizations Related to AI and Policy

The results of my initial inquiries about external groups relevant to AI public policy are in the list at the end of this article. I encourage SIGAI members to contribute additional information on opportunities for us to partner with other groups. An example is work by the American Association for the Advancement of Science, particularly the Center of Science, Policy, and Society. While AAAS policy issues are usually not directly related to AI, a regular look at their Policy Alert notifications is useful for larger policy issues, and helpful to see opportunities for SIGAI to be involved in public policy events. A recent one directly related to AI was the 41st Annual AAAS Forum on Science & Technology Policy. At the panel "Best Friend or Worst Nightmare? Autonomy and AI in the Lab and in Society," AI professionals discussed the role of policy in integrating new technologies into people's lives, particularly for autonomous systems.

AAAI Fall Symposium Series

The FSS-16[1] on November 17-19, 2016 comprised six symposia, all of which are relevant to AI public policy: Accelerating Science: A Grand Challenge for AI, Artificial Intelligence for Human-Robot Interaction, Cognitive Assistance in Government and Public Sector Applications, Cross-Disciplinary Challenges for Autonomous Systems, Privacy and Language Technologies, and Shared Autonomy in Research and Practice.

At the Cognitive Assistance session, two experts on the future of technology and policy spoke about in the future of cognitive assistance in government and public sector applications. Mark Maybury, Chief Security Officer at Mitre, spoke about the unprecedented rapid changes in AI technology applications and the prospects for good and bad impacts on society. Edward Felton, Deputy U.S. CTO in the Office of the Chief Technology Officer, reviewed recent and current initiatives including the impact of AI and cognitive assistants.

Cognitive Assistance is to some extent inspired by the IBM Watson Jeopardy project and now expanding to areas such as healthcare and autonomous systems. Presentations in the different FSS-16 symposia revealed overlap in the areas of public policy and impacts on society, including the proliferation of data, and the accompanying difficulty of managing big data and protecting privacy of individuals and institutions. Most areas of AI research and its applications—importantly, machine learning, neural computing, unsupervised analytics techniques, and

[1] http://www.aaai.org/Symposia/Fall/fss16symposia

deep learning—and connections to brain science have potential impacts on individuals and society, including policy, legal, and privacy issues.

A theme at FSS-16 was human-machine relationships and the need for stakeholders to be in dialogue about legal impacts and potential legislative actions. Should the creators of autonomous systems be responsible for the actions of those systems? Could autonomous systems gain personhood and legal responsibility? Who should decide on and how should we implement ethical frameworks in autonomous systems? How do we monitor and provide input to AI-related policies in the next administration? Public policy must address the encouragement or discouragement of short-term technology development goals, the longer-term implications of autonomous systems including autonomous vehicles, and the increasing influence of AI on human activities.

Upcoming

The theme for the SIGAI Public Policy posts for March is "AI and Future Employment." We will look at potential policies today that could mitigate impacts of AI on future jobs and the economy. Policy areas include innovative education systems, ideas for alternate economic systems, and regulatory changes to promote technological innovation. The results from March will be the topic for a public policy article in the next issue of *AI Matters*.

Resources

- AAAI Symposium series: http://www.aaai.org/Symposia/symposia.php
- AAAI Conferences: http://www.aaai.org/Conferences/
- ACM: http://techpolicy.acm.org
- INNS: http://www.inns.org
- AAAS:https://www.aaas.org/program/center-science-policy-and-society-programs
- White House reports: https://www.whitehouse.gov/blog
- White House Report on the Future of AI: https://obamawhitehouse.archives.gov/blog/2016/10/12/administrations-report-future-artificial-intelligence
- Frontiers Conference: http://www.frontiersconference.org/
- Office of Science and Technology Policy: https://www.whitehouse.gov/ostp
- National Academies of Science, Engineering, and Medicine: http://www.nationalacademies.org
- Online Ethics Center for Engineering and Science: http://www.onlineethics.org
- Brain Initiative: https://www.braininitiative.nih.gov
- DC Data Science, AI, and Policy: http://www.datacommunitydc.org/data-science-dc/
- National Science Foundation: https://www.nsf.gov
- Computing Research Association: http://cra.org
- Computing Community Consortium: http://cra.org/ccc/
- Union of Concerned Scientists: http://www.ucsusa.org
- CMU Cognitive Assistance Lab: http://www.cs.cmu.edu/~NavCog/
- IBM Analytic Solution Center: https://www.ibm.com/connect/federal/us/en/federalinnovationcenters/analytics

- Arkin, R.C. 2013. Lethal autonomous systems and the plight of the non-combatant. *AISB Quarterly* 137: 4-12.
- Balaban, C. 2014. Impacts of autonomous systems. In the *Technical Report on the 2014 AAAI Fall Symposium Series*, Arlington, VA.
- Bonnefon, J., Shariff, A., and Rahwan, I. 2016. The social dilemma of autonomous vehicles. *Science* 352(6293): 1573-1576.
- Marchant, G. and Lindor, R. 2012. The coming collision between autonomous vehicles and the liability system. *Santa Clara Law Review* 52(4): 1321–1340.
- Moor, J. 2006. The nature, importance, and difficulty of machine ethics. *Machine Ethics* 21(4): 13-20.
- Nahamoo, D. 2014. Cognitive computing journey. In PPAA '14 *Proceedings of the First Workshop on Parallel Programming for Analytics Applications*. pp. 63-64.

- Zhao, H., Dimovitz, K., Staveland, B., and Medsker, L. 2016. Responding to challenges in the design of moral autonomous vehicles. In the *Technical Report on the 2016 AAAI Fall Symposium on Cognitive Assistance in Government and Public Sector Applications*, Arlington, VA.

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: Open-Access Educational Resources on AI

Todd W. Neller (Gettysburg College; tneller@gettysburg.edu)
DOI: 10.1145/3054837.3054841

Introduction

Open-access AI educational resources are vital to the quality of the AI education we offer. Avoiding the reinvention of wheels is especially important to us because of the special challenges of AI Education. AI could be said to be "the really interesting miscellaneous pile of Computer Science". While "artificial" is well-understood to encompass engineered artifacts, "intelligence" could be said to encompass any sufficiently difficult problem as would require an intelligent approach and yet does not fall neatly into established Computer Science subdisciplines. Thus AI consists of so many diverse topics that we would be hard-pressed to individually create quality learning experiences for each topic from scratch.

In this column, we focus on a few online resources that we would recommend to AI Educators looking to find good starting points for course development.

AITopics.org

AITopics.org (*AITopics.org*, 2016) is AAAI's portal for sharing quality AI resources and information concerning research and researchers. Covering almost 200 AI topics, it is intended to be a starting reference place for AI instructors and students from high school through first-year graduate school. Its virtual archive of classic publications collects influential works from AI's history. A video collection links to hundreds of AI-related videos. However, the core offering of the site is the topic information categorized as good starting places, general readings, organizations, educational resources, classic articles and books, news, hardware and software, videos, podcasts, etc.

Consider not only browsing this site, but contributing valuable resources to our community via this site as well.

Educational Advances in Artificial Intelligence

Since 2010, the Educational Advances in Artificial Intelligence (EAAI) Symposium (*Educational Advances in Artificial Intelligence*, 2016) has been the premier venue for the sharing of AI educational innovations, research, and teaching resources. Collocated with the AAAI conference, EAAI is a two-day symposium featuring a technical paper track, a Model AI Assignments track (see below) modeled after SIGCSE's Nifty Assignments track, invited speakers, and break-out sessions that have featured mentorship of new faculty, curricular development brainstorming, and educational robotics demonstrations.

This venue provides an excellent opportunity to exchange experiences and resources with many dedicated AI Educators. Consider submitting your best AI educational work to the symposium, and read what others have been doing to advance AI education in the EAAI section of the AAAI proceedings.

Model AI Assignments

The goal of the Model AI Assignments (MAIA) session of EAAI is to gather and disseminate the best assignment designs of the Artificial Intelligence (AI) Education community. Recognizing that assignments form the core of student learning experience, MAIA invites AI educators to submit assignment materials that exemplify an approach to teaching AI topics at all levels. These assignments are then double-blind peer reviewed, and accepted assignments are presented at EAAI and shared via our repository (Neller, 2016) which has become the largest peer-reviewed AI assignment repository to date.

This repository offers excellent ready-made and adaptable assignments to further high-quality AI learning experiences. Consider the current MAIA offerings and what you might share with others through MAIA in the future.

Your Favorite Resources?

These are but a few good starting points for AI instructors to avoid reinventing wheels and instead discover many finely-crafted resources online. If there are other sites you would add to this list of starting points, 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

AITopics.org. (2016). http://aitopics.org. (Accessed: 2016-11-14)

Educational advances in artificial intelligence. (2016). http://eaai.stanford.edu. (Accessed: 2016-11-14)

Neller, T. (2016). *Model AI assignments.* http://modelai.gettysburg.edu. (Accessed: 2016-11-14)

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 Profiles: An Interview with Jim Kurose

Amy McGovern (University of Oklahoma; amcgovern@ou.edu)
Eric Eaton (University of Pennsylvania; eeaton@cis.upenn.edu)
DOI: 10.1145/3054837.3054842

Abstract

This column is the second in a new series profiling senior AI researchers. This month focuses on Jim Kurose.

Introduction

Our second profile for the interview series is Jim Kurose, Assistant Director of the National Science Foundation (NSF) for the Computer and Information Science and Engineering (CISE). Please note that NSF is hiring and would love to have you apply!

Bio

Figure 1: Jim Kurose

Dr. Jim Kurose is an Assistant Director of the National Science Foundation (NSF), where he leads the Directorate for Computer and Information Science and Engineering (CISE) in its mission to support fundamental CISE research, education and transformative advances in cyberinfrastructure across the nation. He is currently a Distinguished Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst, where he has been a faculty member since receiving his PhD in Computer Science from Columbia University. His research area is computer networking, but he did manage to pass a PhD qualifying exam in AI. He is proud to have received a number of research, teaching and service awards over the years, and is particularly proud of the many students with whom he's been so fortunate to work. With Keith Ross, he is the author of the widely adopted textbook *Computer Networking: a Top Down Approach*. Jim is a Fellow of the ACM and IEEE.

Getting to know Jim Kurose

How did you become interested in CS?

My undergraduate degree is in Physics (from Wesleyan University), which didn't have a program in CS at the time. But I took the only two CS courses offered—and loved them both; I worked in the computing center, and had a student job that involved analyzing the various plays run by Wesleyan's football opponents (definitely "small data"!). Probably most importantly, I did some Monte Carlo modeling that complemented the experimental part of my undergrad thesis. I loved physics, but I also had a sense that I'd love computer science, and so I went to grad school expecting to get a MS degree in CS. There, I fell in love with CS research when I met a couple of great faculty who became my PhD advisors.

What was your most difficult professional decision and why?

The hardest decisions are always the ones that affect other people. When there are decisions that run contrary to what a person wants (e.g., passing a PhD qualifying exam), you really need to believe that the decision is in that person's best interests. The people we work with are always so talented that the challenge is really one of helping find the environment in which a given individual will thrive, be happy, and grow.

What professional achievement are you most proud of?

Without a doubt—the students I've taught and mentored; that includes nearly 30 PhD students, and many, many MS and undergrad students. It's really a privilege to have a job that can impact others. There's nothing that makes a day (or a week!) like getting a note from a former student and hearing that you've helped make a difference in that person's life. In second place is the undergraduate textbook (*Computer Networking, a Top-Down Approach*) that I've written with Keith Ross—we wrote that because we both love to write and teach, and have been incredibly pleased (and perhaps a bit shocked!) to see how it has been adopted at so many universities around the world. I am also very proud and honored to be able to serve the CS community in my current position as Assistant Director at the National Science Foundation, where I lead the Directorate for Computer and Information Science and Engineering.

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

Hey—great question! I've given a talk on exactly that topic: "Ten pieces of advice I wished my advisor had told me." I've given this talk at a bunch of student workshops in my research area over the years. Among my favorites in that list are learning how to communicate (write, speak, and tell the narrative of your work), finding role models, and studying broadly.

What would you want for your career if you couldn't do CS?

Impossible to say! I think there's a surprising degree of randomness in where we end up, and how we get there. As the saying goes, "What a long strange trip it's been!" As I mentioned, I didn't go to grad school planning to get a PhD—but my grad school experience turned out to be phenomenal. Nor did I really choose grad school from a particularly career-oriented point-of-view; I just wanted to be where my girlfriend (and now wife) wanted to be. Both turned out great, but the lesson, I think, is to be open to opportunities and to follow your passion. Sounds a bit trite, perhaps, but definitely true.

What is a "typical" day like for you?

No two days are alike in my job at NSF. I spend lots of time working with the amazing CISE staff (program directors, division directors, and administrative team) on both current and future programs; I spend a lot of time interacting with staff from the other directorates at NSF—a real treat as well; and I also spend a good deal of time working with other Federal agencies. Last, I really enjoy spending time *in* the CS community, at meetings and visiting campuses and hearing about the amazing things going on, as well as individual and institutional hopes, aspirations, and concerns.

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

Pretty much all of the aspects of my job at NSF. Let me add that CISE is always looking for smart, dedicated and talented folks from the research community who might be interested in serving a rotation as an NSF/CISE Program Director. I'd encourage anyone interested to contact the relevant CISE division director or me—we'll be happy to tell you more about the opportunities.

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

We all depend on so many other people. As students, we depend on our teachers, staff, mentors and other students; as faculty, we depend on our students, colleagues and collaborators; in academic leadership, we depend on the people with whom we work to help make things happen. For these many activities to be successful we need to rely on other people, and be reliable to those with whom we work; we really do achieve both more and better things by working together. At NSF, it's been great to work with Lynne Parker, NSF/CISE Division Director for Information and Intelligent Systems, and her team, who provide NSF's technical vision, leadership and management of programs in AI and Information and Intelligent Systems more broadly.

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

I can still remember being completely blown away as a kid when I saw *2001: A Space Odyssey*. It was visually stunning, had the HAL 9000 computer (of course, I'd never even seen a computer then), and was wildly inscrutable to a twelve-year-old. For CS/AI-related books, my favorites are anything written by Isaac Asimov, and *Snowcrash* by Neal Stephenson. Beyond science fiction, I've just finished The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies by Erik Brynjolfsson and Andrew McAfee. All of these books speak to the relationship between humans and technology—a topic of increasing importance for everyone.

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.

Learning Peripersonal Space Representation in a Humanoid Robot with Artificial Skin

Alessandro Roncone (Yale University; alessandro.roncone@yale.edu)
DOI: 10.1145/3054837.3054843

The peripersonal space (PPS) is a model of the nearby space that is of special relevance for the life of any complex animal. It is defined as the space surrounding our bodies, and is best described as a multisensory and sensorimotor interface with a fundamental role in the sensory guidance of actions. Owing to it is the ability to perform timely and appropriate actions toward objects located in the nearby space, which is critical for the survival of every animal. Depending on context, these actions may constitute an approaching or an avoidance behavior. In the case of defensive behavior, this creates a "margin of safety" around the body, such as the flight zone of grazing animals or the multimodal attentional space that surrounds the skin in humans (Graziano & Cooke, 2006). Peripersonal space thus deserves special attention and probably justifies the specific neural circuitry devoted to its representation. The brain has to dynamically integrate information coming from several modalities being these motoric, visual, auditory or somatosensory (Graziano & Cooke, 2006; Holmes & Spence, 2004). In animals and humans, PPS representations are gradually formed through physical interaction with the environment and in a complex interplay of body growth and neural maturation processes. *Self-touch* (also called double-touch) is presumably one of the behaviors that impact the formation of multimodal body representations (Rochat, 1998). They occur across different motor and sensory modalities, with the motor/proprioceptive and tactile starting already in prenatal stage. Vision is presumably incorporated later, during the first months after birth hand in hand with the maturation of the visual system. Perhaps even later, contingencies will encompass external objects.

This simplified developmental timeline constitutes the skeleton of our work in the humanoid robot. The humanoid in question is the iCub (Fig. 1), a robot with human-like morphology and a subset of the sensory capacities of the human body (Metta et al., 2010). Lately, it has been equipped with a set of tactile sensors, which provide information about local pressure upon contact. We model peripersonal space on the iCub by focusing on the integration of visual and tactile inputs. Our work shows how this representation can be learned following a developmental progression, starting with the robot exploring its own body by self-touch (Roncone, Hoffmann, Pattacini, & Metta, 2014) and registering correlations in motor and tactile spaces. Subsequently, visual information derived from self-observation kicks in and makes provision for learning multimodal responses. Then, learning is extended to other objects nearing and contacting the robot's body (Fig. 1). Our model keeps in register each "spatial" visual receptive field (RF)

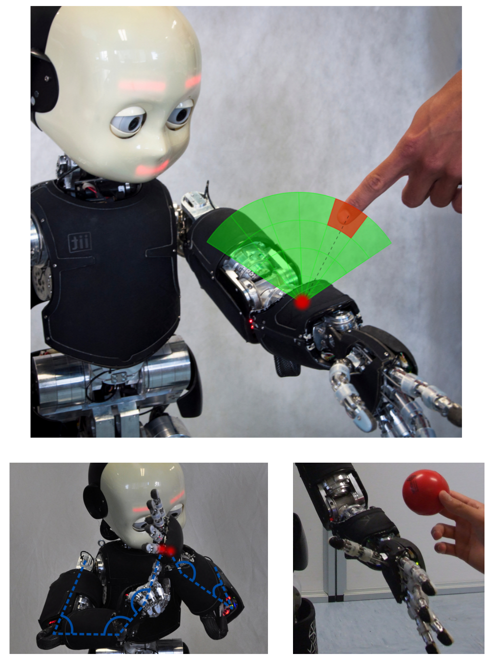

Figure 1: **(top)** The iCub robot with a depiction of a visuo-tactile receptive field above one of the left forearm taxels. **(bottom-left)** double-touch behavior with a schematic of the kinematics and joint angles. **(bottom-right)** An object approaching the left forearm. Tracking of the object employs an optical flow tracker (Roncone et al., 2015).

to a taxel (tactile element) of the robot's skin. RFs are proxies for the neural responses, each of them represented by a probability density function. Probabilities are updated incrementally and carry information about the likelihood of a particular stimulus (e.g. an object approaching the body) eventually contacting the specific taxel at hand. We use the distance to the taxel and its time to contact to compactly identify the stimulus in a bi-dimensional parameter space. An important feature of the proposed model is that learning is fast, proceeds in parallel for the whole body, and is incremental. That is, minutes of experience with objects moving toward a body part produce a reasonable representation in the corresponding taxels that is manifested in prior to contact activations. This representation naturally serves the purpose of predicting contacts with any part of the body of the robot, which is of clear behavioral relevance. Furthermore, we implemented an avoidance controller whose activation is triggered by this representation, thus endowing the iCub with a "margin of safety". Finally, minor modifications to the controller result in a reaching behavior that can use any body part as the end effector.

This work is promising for future applications: with robots leaving the factory and entering less controlled domains, possibly sharing the space with humans, safety is paramount and multimodal awareness of the body surface and the space surrounding it is fundamental. The robotics implementation departs in many respects from the mechanisms that presumably operate in the primate brain. The correspondence between biology and robotics is established at a behavioral level rather than in the details of the implementation. In particular, for mostly practical reasons, we assume that the robot's kinematics and mapping of tactile information into reference frames is given from hand-coded models. The implementation of the double-touch behavior itself is taken as a primitive (Roncone et al., 2014). Conversely, learning/calibration of the spatial receptive fields around individual taxels is primarily addressed here and relates to biology.

References

Graziano, M., & Cooke, D. (2006). Parieto-frontal interactions, personal space and defensive behavior. *Neuropsychologia*, *44*, 845-859.

Holmes, N. P., & Spence, C. (2004, Jun). The body schema and the multisensory representation(s) of peripersonal space. *Cogn Process*, *5*(2), 94–105.

Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., et al. (2010). The icub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks*, *23*(8-9), 1125-1134.

Rochat, P. (1998). Self-perception and action in infancy. *Exp Brain Res*, *123*, 102-109.

Roncone, A., Hoffmann, M., Pattacini, U., & Metta, G. (2014). Automatic kinematic chain calibration using artificial skin: self-touch in the iCub humanoid robot. In *Robotics and Automation (ICRA), 2014 IEEE Int. Conf. on* (pp. 2305–2312).

Roncone, A., Hoffmann, M., Pattacini, U., & Metta, G. (2015). Learning peripersonal space representation through artificial skin for avoidance and reaching with whole body surface. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ Int. Conf. on* (p. 3366-3373).

Alessandro Roncone is a Postdoctoral Associate at the Social Robotics Lab in Yale University. He received his Ph.D. in Robotics, Cognition and Interaction Technologies while working at the Robotics, Brain and Cognitive Sciences department and the iCub Facility in IIT, Genova, under the supervision of prof. Giorgio Metta. The goal of his project was to implement a model of peripersonal space on the iCub humanoid robot. He further elaborated on that line of research by leveraging peripersonal space representations in two directions: i) better, richer body representations ii) distributed motor control via whole-body awareness. He is currently focusing on the exploitation of natural interactions between humans and robots in the context of human-robot collaboration.

Multimodal Concepts for Social Robots

Olivier Mangin (Yale University; olivier.mangin@yale.edu)
DOI: 10.1145/3054837.3054844

Whether they try to mimic human cognition or give us a new glimpse of it as modelling tools, robots have become an essential component of cognitive sciences. My doctoral work embraces this duality of robotic research along two topics that are the study of language acquisition and learning by imitation.

How do concepts such as *red* or *grasp* emerge from our perceptual experience? When S. Harnad formulated the *symbol grounding problem* (Harnad, 1990), he shifted the traditional artificial intelligence's focus away from symbol manipulation problems to the definition, origin, and existence of these symbols. I hence study that question in line with two fundamental perspectives from developmental robotics: intelligence is the result of a developmental process and intelligence is embodied, i.e. it happens in the context of a sensorimotor interaction between the body and the world (Steels, 2008). How do we learn to recognize the words '*red*' and '*grasp*'? How do we relate the action of grasping performed by ourselves and someone else? Despite their apparent similarity to the first one, these two additional questions were originally studied in isolation, by distinct communities. One of the main contribution of my doctoral research is to bring these problems together in an interdisciplinary approach. To do so, I frame them into a common model: *the emergence of concepts from perception*, that covers a definition of motion primitives, the acquisition of words, as well as the discovery of objects in vision.

How can an animal or a robot acquire and represent skills, so that it can later reuse them and combine them into new more complex skills? My work focuses on the ability for robots to decompose demonstrations of complex motions or skills into a repertoire of primitive gestures, an approach that is inspired on the idea of *motion primitives*, rooted on biological evidences (Mussa-Ivaldi & Bizzi, 2000). It has promising applications for robots and in particular robots programmed by demonstration, an approach inspired from imitation learning in humans, to teach new skills to robots without needing expert programming. Indeed, robot programming by demonstration is currently limited to simple skills and is expected to greatly benefit from the ability to relate complex skills to a growing repertoire of primitive ones (Cangelosi, A. et al., 2010).

My doctoral research brings the following original contributions. First, I focused on the simultaneous combination of motion primitive instead of the more commonly studied sequential one. As a starting point, I built datasets of dance motions, that are short choreographies obtained by the simultaneous mixture of basic gestures. I additionally developed experiments and devised algorithms to learn gestures together with words (symbolic or from continuous speech) that describe the gestures (Mangin & Oudeyer, 2012, 2013). Indeed, social conventions play an important role in the perception of gestures which makes the motion decomposition problem alone ill-posed. Interestingly, in my thesis (Mangin, 2014), I frame the process of learning words, gestures, as well as visual objects, into a multimodal problem that treats all modalities similarly. The fact that one single algorithm enables both the learning of words and gestures establishes a strong symmetry between these problems. As another significant consequence, this experiment shows that, although these three individual problems are ambiguous, rephrasing them together as a multimodal learning problem can overcome their individual ambiguities. In particular, it allows the acoustic concepts of words to both shape and be shaped by the visual model of objects or gestures (Mangin, Filliat, Bosch, & Oudeyer, 2015), thus providing a model of the co-organization of language and concepts (Lupyan, Rakison, & McClelland, 2007).

Finally, these models are based on a simple compression heuristics and operate on unsegmented sentences (raw acoustic signal) in a multimodal sensory flow, without enforcing an *a priori* structure on sentences. Therefore this research enables a bottom-up approach to a

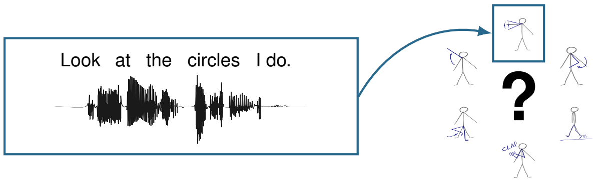

Figure 1: We evaluate the learner's behavior on a multimodal classification task: it hears a new utterance (raw acoustic signal) and choses the best matching gesture among demonstrations.

problem mostly studied from a top-down perspective. Moreover, I designed experiments to elicit *a posteriori* the acquisition of structure; these use behavioral tasks analogous to a child associating words to objects rather than requiring to look for some explicit representation in the robot's brain (see Figure 1).

This work presents an artificial learner that acquires word knowledge from cross-situational information. Future work includes experiments on realistic data to better compare such information in word learning to other cues such as mutual exclusivity or whole object assumption (Markman, 1990). Another direction is the learning of higher levels of structure, as grammar structures for language but also gestures. Finally incremental learning algorithms would enable to study the developmental path of the learning of such structure.

In summary, this research manages to bring together questions coming from several different domains and demonstrates that these questions can shed light on each other. In particular I posit an *algorithmic analogy* between the questions, that includes applications of well known algorithms (nonnegative matrix factorization) to new domains (such as motion decomposition), as well as the development of new algorithms (matrix factorization for inverse reinforcement learning). To conclude, I would like to stress the effort that has been achieved to make data, code[1], and results openly available for dissemination and for reproduction of the experiments.

References

Cangelosi, A. et al. (2010). Integration of Action and Language Knowledge: A Roadmap for Developmental Robotics. *IEEE TAMD*, *2*(3), 167–195.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*(1), 335–346.

Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: redundant labels facilitate learning of novel categories. *Psychological Science*, *18*(12), 1077–1083.

Mangin, O. (2014). *The emergence of multimodal concepts: From perceptual motion primitives to grounded acoustic words* (PhD thesis). Université de Bordeaux.

Mangin, O., Filliat, D., Bosch, L., & Oudeyer, P.-Y. (2015). MCA-NMF: Multimodal concept acquisition with non-negative matrix factorization. *PLoS ONE*, *10*(10).

Mangin, O., & Oudeyer, P.-Y. (2012). Learning to recognize parallel combinations of human motion primitives with linguistic descriptions using non-negative matrix factorization. In *IROS.* Vilamoura, Portugal.

Mangin, O., & Oudeyer, P.-Y. (2013). Learning semantic components from subsymbolic multimodal perception. In *ICDL-EpiRob.*

Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, *14*, 57–77.

Mussa-Ivaldi, F. A., & Bizzi, E. (2000). Motor learning through the combination of primitives. *Philosophical trans., Royal Soc. of London.*, *355*(1404), 1755–69.

Steels, L. (2008). The symbol grounding problem has been solved. so what's next? In *Symbols and embodiment* (chap. 12). Oxford University Press.

Olivier Mangin is a developmental and social roboticist, now a postdoct in Prof. Scassellati's laboratory at Yale University. His interest lies in the apparition of structure in the sensorimotor interaction between an infant or robot and its environment. In particular he studies the development of linguistic skills, and the acquisition and understanding of behaviors during human-robot collaboration. O. Mangin completed his PhD at INRIA in France, under the supervision of P.-Y. Oudeyer.

[1] http://github.com/omangin/multimodal

Learning the State of the World: Object-based World Modeling for Mobile Manipulation Robots

Lawson L.S. Wong (Ph.D. 2016, MIT; now at Brown University; lsw@brown.edu)
DOI: 10.1145/3054837.3054845

Mobile-manipulation robots performing service tasks in human-centric indoor environments have long been a dream for developers of autonomous agents. Tasks such as cooking and cleaning involve interaction with the environment, hence robots need to know about their spatial surroundings. However, service robots operate in environments that are relatively unstructured and dynamic. Mobile-manipulation robots therefore need to continuously perform *state estimation*, using perceptual information to maintain a representation of the state, and its uncertainty, of the world.

By definition, mobile-manipulation robots are capable of moving in and interacting with the world. Hence, at the very least, such robots need to know about the physical occupancy of space and potential targets of interaction (i.e., objects). For the former, there is a long history of representations in the field of navigation and mapping; occupancy grids (Moravec & Elfes, 1985) are a widely-used example. In contrast, object-based representations for robotics are still in their infancy. In my dissertation, I propose a representation based on objects, their 'semantic' attributes (properties such as type and pose), and their geometric realizations in the physical world.

Objects are challenging to keep track of because there is significant *uncertainty* in their states. Object detection and recognition is still far from solved within classical computer vision, and even less so from a robotic vision standpoint. Objects can also be inherently ambiguous because they have the same values for some, or even all, attributes. Besides detection noise, other agents may manipulate objects as well and change object states without informing robots. Compounded over multitudes of objects (thousands or more) and long temporal horizons (days or longer), the above sources of uncertainty give rise to a large and difficult estimation problem.

Figure 1: Mobile-manipulation robots operating in human-centric environments must know about, and be able to model, the world in terms of objects.

Data Association for Semantic World Modeling from Partial Views

A basic world model could simply use an object detector's output on a single image as a representation of the world. However, doing so suffers from errors such as sensor measurement noise, object occlusion, and detection algorithm approximations. Aggregating measurements across different viewpoints can reduce estimation error. The key challenge in this strategy is *identity management*, induced by measurements that often cannot be uniquely mapped to an underlying object.

I proposed a Bayesian nonparametric clustering approach to data association, inspired by the observation that 'objects' are essentially clusters in joint attribute space. Building on top of Dirichlet-process mixture models (Antoniak, 1974), I incorporated crucial domain assumptions, and used the new model to cluster similar attribute measurements in static scenes. Given attribute detections from multiple viewpoints, this algorithm outputs samples from the distribution over hypotheses of object states, where a hypothesis consists of a list of objects and their attribute value distributions (Wong, Kaelbling, & Lozano-Pérez, 2015). In recent work, I extended the model to a dynamic clustering setting to handle objects states that change over time (Wong, Kurutach, Lozano-Pérez, & Kaelbling, 2016). Figure 2 illustrates the full world modeling problem.

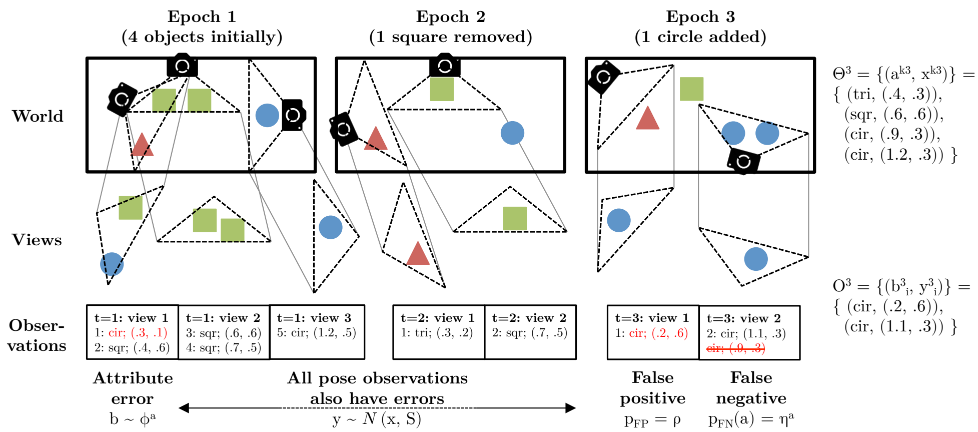

Figure 2: An illustration of the world modeling problem. An unknown number of objects exist in the world (top row), and change in pose and number over discrete epochs. In each epoch, partial views of the world are captured, as depicted by the triangular viewcones. Objects within the viewcones have detectable attributes; in this example, the attributes are shape type (discrete) and 2-D location. The observations are noisy, as depicted by the perturbed versions of viewcones in the middle row. Uncertainty exists both in the attribute values and the existence of objects, as detections may include false positives and negatives (e.g., $t = 3$). The actual attribute detection values obtained from the views are shown in the bottom row; this is the format of input data. Given these noisy measurements as input, the goal is to determine which objects were in existence at each epoch, their attribute values (e.g., Θ^3 in top right), and their progression over time.

Combining Object and Metric Information

One concept lacking in the above work is the notion that objects occupy physical regions of space. The concept of free space, regions that no object overlaps, was also only implicitly represented. It is therefore difficult, in the object-attribute representation, to incorporate absence/'negative' observations, most prominently that observing a region of free space should suggest that no object overlaps that region. This information is handled very naturally in conventional occupancy grids, but grids cannot handle objects elegantly.

The complementary advantages of these two representations inspired a search for a way to maintain estimates of both object and metric information. Because filtering in the joint state is often intractable, I instead adopted the strategy of filtering *separately* in the object and metric spaces by using the previous section's model and occupancy grids. To compensate for the lost dependencies, I then developed a way to *fuse* the filters on demand as queries about either posterior distribution are made (Wong, Kaelbling, & Lozano-Pérez, 2014). Our results suggest that maintaining simple, disparate, and aggressively-factored estimators is potentially superior to keeping a complex joint estimate.

References

Antoniak, C. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics*, *2*(6), 1152-1174.

Moravec, H., & Elfes, A. E. (1985). High resolution maps from wide angle sonar. In *ICRA.*

Wong, L., Kaelbling, L., & Lozano-Pérez, T. (2014). Not seeing is also believing: Combining object and metric spatial information. In *ICRA.*

Wong, L., Kaelbling, L., & Lozano-Pérez, T. (2015). Data association for semantic world modeling from partial views. *The International Journal of Robotics Research*, *34*(7), 1064–1082.

Wong, L., Kurutach, T., Lozano-Pérez, T., & Kaelbling, L. (2016). Object-based world modeling in semi-static environments with dependent Dirichlet-process mixtures. In *IJCAI.*

Lawson L.S. Wong is a postdoctoral fellow at Brown University, working with Stefanie Tellex. He completed his Ph.D. in Jan. 2016, advised by Leslie Pack Kaelbling and Tomás Lozano-Pérez. He is currently supported by a Croucher Foundation Fellowship for Postdoctoral Research.

Evolutionary Online Learning in Multirobot Systems

Fernando Silva (Instituto de Telecomunicações; fsilva@di.fc.ul.pt)
Luís Correia (BioISI; luis.correia@ciencias.ulisboa.pt)
Anders Lyhne Christensen (Instituto de Telecomunicações; anders.christensen@iscte.pt)
DOI: 10.1145/3054837.3054846

Introduction

Robots have the potential to replace manned machines and to carry out tasks in environments that are either remote or hazardous, such as space, deep sea, or underground. However, to create intelligent, reliable, mobile robots, capable of operating effectively in a wide variety of environments, the *limited learning ability* of robots needs to be addressed (Brooks & Matarić, 1993). Mobile robots are typically brittle: they are unable to adapt to changing environmental conditions (e.g., changes in terrain or lighting) or internal conditions (e.g., drift or permanent failure in sensors and/or actuators), and to learn new tasks as they execute.

We study how learning can be achieved in robotic systems through evolutionary algorithms (EAs), a nature-inspired approach that mimics Darwinian evolution. Instead of manually programming the robots to carry out a mission, an EA is executed onboard robots *during task execution* in order to synthesize and continuously optimize the artificial brain or controller of each robot. This approach is known as *online evolution* and can automatically generate the artificial intelligence that controls each robot. However, despite the potential for automatic robot learning, online evolution methods have not yet been able to solve any but the simplest of tasks, and the approach typically requires a prohibitively long time to evolve controllers on real robots (several hours or days) (Silva, Duarte, Correia, Oliveira, & Christensen, 2016).

Contributions of Research

We have introduced a novel online EA called odNEAT (Silva, Urbano, Correia, & Christensen, 2015) for decentralized online evolution of artificial neural network (ANN) controllers in groups of robots that evolve in parallel and exchange candidate controllers to the task. Contrarily to previous approaches, in which the controller structure is defined by the human experimenter, odNEAT evolves both topology and weighting parameters of ANNs. The algorithm starts with minimal networks and effectively complexifies them by adding new neurons and new connections through mutation. In this way, odNEAT can automatically find an appropriate degree of complexity to the current task (Silva, Urbano, et al., 2015).

We have extensively assessed the performance of odNEAT in a number of simulation-based studies, in which the algorithm was shown to enable: (i) *scalability* (Silva, Correia, & Christensen, 2015), as groups of different size can leverage their multiplicity to achieve superior task performance and speed up evolution, (ii) *robustness*, as the controllers evolved can often adapt to changes in environmental conditions without further evolution (Silva, Urbano, et al., 2015), and (iii) *fault tolerance*, as robots executing odNEAT are able to adapt their behavior and to learn new behaviors in the presence of sensor faults (Silva, Urbano, et al., 2015). We have additionally developed different approaches to speed up online evolution, including a technique in which sub-behaviors can be prespecified in the neural architecture (Silva, Correia, & Christensen, 2014), and racing techniques to cut short the evaluation of poor controllers (Silva, Correia, & Christensen, 2016).

For the real-robot experiments, we have implemented odNEAT in groups of Thymio II robots, each of which extended with a Raspberry Pi 2 single-board computer (see Fig. 1). The robots form an ad-hoc IEEE 802.11g wireless network, and communicate with one another by broadcasting UDP datagrams. We have successfully evolved controllers for canonical tasks, including navigation and obstacle avoidance, homing, and aggregation. The controllers were evolved completely on-

Figure 1: Three of our Thymio II robots (two in the front, one in the back), each of which extended with a Raspberry Pi 2 single-board computer.

line, in less than one hour, thereby showing the potential of our approach and a path towards online evolution in a timely manner. As the final topic of the doctoral research, we are experimenting with evolutionary multi-level composition of evolved control. Our final goal is to enable scalable and efficient synthesis of solutions for complex, multi-competence tasks beyond the state of the art in the field, and to help realize the full potential of online evolution.

Acknowledgments. This research was partly supported by FCT under grants SFRH/BD/89573/2012, UID/EEA/50008/2013, and UID/Multi/04046/2013.

References

Brooks, R. A., & Matarić, M. J. (1993). Real robots, real learning problems. In J. H. Connell & S. Mahadevan (Eds.), *Robot Learning* (pp. 193–213). Springer US, Boston, MA.

Silva, F., Correia, L., & Christensen, A. L. (2014). Speeding up online evolution of robotic controllers with macro-neurons. In *17th European Conference on the Applications of Evolutionary Computation* (pp. 765–776). Springer, Germany.

Silva, F., Correia, L., & Christensen, A. L. (2015). A case study on the scalability of online evolution of robotic controllers. In *17th Portuguese Conference on Artificial Intelligence* (pp. 189–200). Springer, Germany.

Silva, F., Correia, L., & Christensen, A. L. (2016). Leveraging online racing and population cloning in evolutionary multirobot systems. In *19th European Conference on the Applications of Evolutionary Computation* (pp. 165–180). Springer International Publishing, Switzerland.

Silva, F., Duarte, M., Correia, L., Oliveira, S. M., & Christensen, A. L. (2016). Open issues in evolutionary robotics. *Evolutionary Computation*, *24*(2), 205–236.

Silva, F., Urbano, P., Correia, L., & Christensen, A. L. (2015). odNEAT: An algorithm for decentralised online evolution of robotic controllers. *Evolutionary Computation*, *23*(3), 421–449.

Fernando Silva is a doctoral student and researcher at Instituto de Telecomunicações, Portugal. He studies how to enable the synthesis of adaptive and robust robots by taking inspiration from biological principles at multiple levels, and has published over 25 papers on evolutionary computation, machine learning, and robotics.

Luís M. P. Correia is associate professor with habilitation at Department of Informatics of Faculty of Sciences of University of Lisbon, Portugal. He lead the Laboratory of Agent Modelling (LabMAg), 2004-2014, and leads the MAS group of BioISI from 2015. His research interests are artificial life, autonomous robots, and self-organised systems.

Anders Lyhne Christensen is a tenured assistant professor at University Institute of Lisbon, Portugal, where he founded the Bio-inspired Computation and Intelligent Machines Lab. Dr. Christensen has published more than 100 papers on artificial intelligence and robots, and he is on the editorial board of three international journals.

Interactions Between Learning and Decision Making

Theja Tulabandhula (Massachusetts Institute of Technology; theja@alum.mit.edu)
DOI: 10.1145/3054837.3054847

This thesis revolves around learning and decision making and how these two processes interact with each other in surprising ways. To this end, we have studied the effects of machine learning on decision making and vice versa in a variety of theoretical contexts. On the way, we have also touched application areas such as power grid maintenance scheduling and professional car racing.

In the setting studied in this thesis, a learning problem provides estimated probabilities, and those probabilities all feed into a single optimization problem for decision-making. This is how people typically set up facility location problems, traveling repairman and traveling salesman problems, knapsack problems, and so on. The interesting part is how the uncertainty in predictions translates into uncertainty in the decision problem. We used tools from statistical learning theory and robust optimization to provide important insights into this question.

Machine learning with operational costs: Quantifying influence using statistical learning theory (Tulabandhula & Rudin, 2013, 2014b) - We proposed a way to align statistical modeling with decision making, which we called the "Machine Learning with Operational Costs (MLOC)" framework. We focused on the class of problems where we, (a) construct a data dependent predictive model (a classification or regression function), and (b) solve a decision making optimization problem whose parameters depend on the predictive model. We proposed a new method that propagates the uncertainty in predictive modeling to the uncertainty in operational cost, where operational cost is the amount spent by the practitioner in solving the problem. The method allows us to explore the range of operational costs associated with the set of reasonable statistical models as shown in Figure 1, so as to provide a useful way for practitioners to understand uncertainty. To do this computationally, we cast the operational cost as a regularization term in a learning algorithm's objective function, allowing either an optimistic or pessimistic view of possible costs, depending on the regularization parameter.

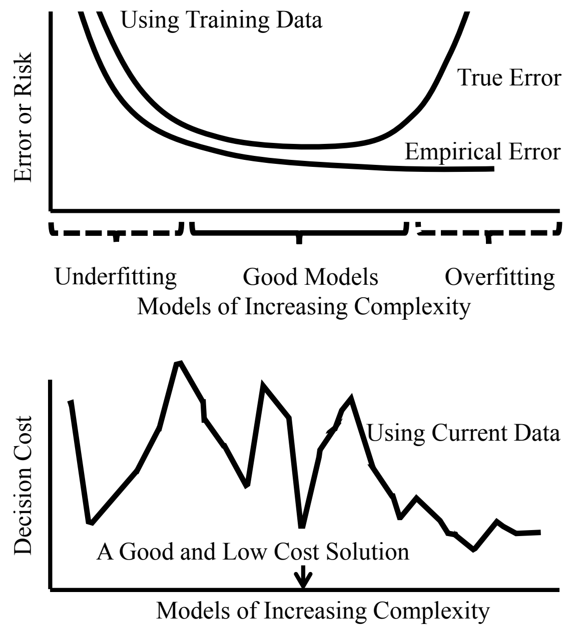

Figure 1: Operational/Decision costs change depending on the predictive model chosen.

From another perspective, if we have prior knowledge about the operational cost, for instance that it should be low, this knowledge can help to restrict the hypothesis space, and can help with generalization. In fact, it may be much more natural for a manager to have a prior belief on the cost to solve a problem than, for instance, a belief on the ℓ_2 norm of the coefficients of a predictive model, or the number of nonzero coefficients. Based on this view, we developed several statistical learning theory generalization bounds that take the operational cost knowledge into account. These bounds are novel, and require bounds on the complexity of a set of functions (covering numbers, VC dimension, Rademacher complexity, etc.). We also showed that learning with operational costs is related to the robust optimization framework.

We explored power grid maintenance using recent data provided by Con Edison for New York City (Tulabandhula & Rudin, 2014b), call center staffing, portfolio optimization and healthcare logistics problems under this framework. Our major findings were that it is very important to understand how the uncertainty in the predictive modeling translates into real-world uncertainty. For instance, for routing utility trucks on the NYC power grid, there can be quite a lot of uncertainty in what route is optimal for inspecting power grid equipment. This is something extremely important for power grid operators to know about; yet, they would not have been able to consider this uncertainty without our methods. Along similar lines and with similar challenges, we also performed knowledge discovery and designed a novel decision support tool for professional car racing (Tulabandhula & Rudin, 2014d).

Decision making backed by machine learning: Learning uncertainty sets for robust optimization (Tulabandhula & Rudin, 2014c) - We want our decision to best handle the the worst possible situation that could arise, out of an *uncertainty set* of possible situations. Classically, the uncertainty set is simply chosen by the user, or it might be estimated in overly simplistic ways with strong distributional assumptions; whereas in this work, we show how to *learn the uncertainty set from data collected in the past*. This method is principled, and backed by statistical learning theoretic bounds. And it will allow practitioners to construct uncertainty sets that are just right – not too conservative, and not too small.

Machine learning beyond decision making priors: An analysis using convex duality (Tulabandhula & Rudin, 2014a) - We considered a supervised learning setting where side knowledge is provided about the labels of an additional set of unlabeled examples. One of the ways such side knowledge can arise is through knowledge about an associated decision making problem (the MLOC setting). We considered many other sources of side knowledge than what had been studied rigorously in the past that lead to linear, polygonal, quadratic or conic constraints constraints on the hypothesis space. We proved bounds on the complexity measures of these constrained hypothesis spaces. These are some of the first results that enable a practitioner to understand the impact of different types of side knowledge on their learning problem.

Acknowledgments. Funding for this thesis came from a Fulbright Science and Technology Fellowship, the Solomon Buchsbaum Research Fund, and NSF grant IIS-1053407.

References

Tulabandhula, T., & Rudin, C. (2013). Machine Learning with Operational Costs. *Journal of Machine Learning Research*, *14*, 1989-2028.

Tulabandhula, T., & Rudin, C. (2014a). Generalization Bounds for Learning with Linear, Polygonal, Quadratic and Conic Side Knowledge. *Machine Learning Journal*, 1-34.

Tulabandhula, T., & Rudin, C. (2014b). On Combining Machine Learning with Decision Making. *Machine Learning Journal*, *97*(1-2), 33-64.

Tulabandhula, T., & Rudin, C. (2014c). Robust Optimization using Machine Learning for Uncertainty Sets. In *Proceedings of the International Symposium on Artificial Intelligence and Mathematics.*

Tulabandhula, T., & Rudin, C. (2014d). Tire Changes, Fresh Air, and Yellow Flags: Challenges in Predictive Analytics for Professional Racing. *Big Data*, *2*(2).

Theja Tulabandhula received his bachelors and masters degrees in electrical engineering from Indian Institute of Technology, Kharagpur, in 2009, where he was awarded the Prime Minister of India Gold medal. He received his Ph.D. degree in 2014 at the Massachusetts Institute of Technology (MIT), Cambridge. There, he was a Xerox-MIT fellow and a Science and Technology Fulbright scholar. He is currently a research scientist in the Analytics Lab at the Xerox Research Centre India.