

# IUI Workshop on Interactive Machine Learning

**Saleema Amershi**  
 Microsoft Research  
 One Microsoft Way  
 Redmond, WA 98052  
 samershi@microsoft.com

**Maya Cakmak**  
 Willow Garage, Inc.  
 68 Willow Road  
 Menlo Park, CA 94025  
 mcakmak@willowgarage.com

**W. Bradley Knox**  
 Massachusetts Institute of  
 Technology  
 77 Mass. Ave.  
 Cambridge, MA 02142 USA  
 bradknox@mit.edu

**Todd Kulesza**  
 Oregon State University  
 1148 Kelley Engineering Ctr  
 Corvallis, OR, 97331  
 kuleszto@eecs.oregonstate.edu

**Tessa Lau**  
 Willow Garage, Inc.  
 68 Willow Road  
 Menlo Park, CA 94025  
 tlau@willowgarage.com

## ABSTRACT

Many applications of Machine Learning (ML) involve interactions with humans. Humans may provide input to a learning algorithm (in the form of labels, demonstrations, corrections, rankings or evaluations) while observing its outputs (in the form of feedback, predictions or executions). Although humans are an integral part of the learning process, traditional ML systems used in these applications are agnostic to the fact that inputs/outputs are from/for humans.

However, a growing community of researchers at the intersection of ML and human-computer interaction are making interaction with humans a central part of developing ML systems. These efforts include applying interaction design principles to ML systems, using human-subject testing to evaluate ML systems and inspire new methods, and changing the input and output channels of ML systems to better leverage human capabilities. With this Interactive Machine Learning (IML) workshop at IUI 2013 we aim to bring this community together to share ideas, get up-to-date on recent advances, progress towards a common framework and terminology for the field, and discuss the open questions and challenges of IML.

## Author Keywords

End-user programming; Active learning; Reinforcement learning with human feedback or guidance; Interactive clustering; Feature labeling; Programming by demonstration; Transparency and feedback in machine learning; Empirical studies and computational models of human teaching; Democratizing machine learning; Human-in-the-loop intelligent systems

## ACM Classification Keywords

I.2.6 Learning: General; H.1.2 User/Machine Systems: Human information processing; H.5.0 Information Interfaces and Presentation (e.g. HCI): General

## OVERVIEW

Programmability and personalization are key components in the long term success of intelligent systems that assist humans in everyday tasks. Developments in the field of machine learning (ML) have created a range of tools that make these possible. However, the use of ML tools as black-boxes limits their effectiveness and restricts their inputs and outputs. Instead, a growing community of researchers at the intersection of ML and human-machine interaction are making interaction design a central part of building such intelligent systems and changing the way ML tools are used [15, 11, 8]. Similarly, in a series of workshops on agents that learn interactively from human teachers [2, 3, 6, 1], we have seen user interaction issues come up repeatedly, giving rise to a number of novel interaction techniques [13, 12, 9]. Therefore, with this Interactive Machine Learning (IML) workshop at IUI—a progression from the theme of these past four workshops and other IML workshops [7, 4, 5]—we aim to bring these communities together, to focus on interaction and user interface issues in designing intelligent systems that learn from their users.

## WORKSHOP TOPICS

The workshop will cover topics ranging from novel ML techniques that change the role of humans in the process—*e.g.*, answering queries from the learner [14], labeling features [10], or providing evaluative reinforcement [12]—to interface considerations such as transparency about the learning system's internal processes or the amount of control given to the human. We also aim to draw participation from industry representatives by emphasizing real-world problems, such as personalization in e-commerce, desktop or mobile assistants, and end-user programmable robots.

In the workshop's Call For Papers, the following topics were listed as relevant:

- End-user programming
- Active learning
- Reinforcement learning with human feedback or guidance
- Interactive clustering
- Feature labeling
- Programming by demonstration
- Transparency and feedback in ML
- Empirical studies and computational models of human teaching
- Democratizing ML
- Human-in-the-loop intelligent systems

#### WORKSHOP FORMAT

We are planning a full-day workshop with contributions in the form of two-page extended abstracts. The workshop program will involve the following:

- invited speakers focusing on different issues related to IML;
- a discussion session, focusing on open questions related to IML and potentials for collaborations across different communities;
- a poster session involving all contributed abstracts.

#### INVITED SPEAKERS

The following invited speakers will present during IML:

- Rich Caruana - Microsoft Research
- Rebecca Fiebrink - Princeton University
- Chad Jenkins - Brown University
- Ashish Kapoor - Microsoft Research
- Henry Lieberman - MIT
- Michael Littman - Brown University

#### PROGRAM COMMITTEE

The IML workshop's Program Committee is as follows:

- Luis Carlos Cobo - Georgia Institute of Technology
- Nick DePalma - MIT
- Hayley Hung - University of Amsterdam
- Kshitij Judah - Oregon State University
- George Konidaris - MIT
- Brian Lim - Fraunhofer CSE
- Manuel Lopes - INRIA
- Andreas Paepcke - Stanford University
- Stephanie Rosenthal - Bossa Nova Robotics
- Simone Stumpf - City University London
- Bener Suay - WPI
- Kaushik Subramanian - Georgia Institute of Technology
- James E. Young - University of Manitoba

#### ORGANIZERS

**Saleema Amershi** is a Researcher in the Computer Human Interactive Learning (CHIL) group at Microsoft Research. She received her PhD in Computer Science from the University of Washington's Computer Science & Engineering department. Her research interests are at the intersection of human-computer interaction and machine learning. In particular, her work involves designing effective end-user driven machine learning for a variety of real-world applications. Concrete examples include image classification, access control in online social networks, and alarm triage in large-scale computer networks.

**Maya Cakmak** is a Post-doctoral Research Fellow at Willow Garage. She received her PhD degree in Robotics from the Georgia Institute of Technology in 2012. She has a B.Sc. degree in Electrical and Electronics Engineering and an M.Sc. degree in Computer Engineering from the Middle East Technical University, Turkey. Her research interests include learning from demonstration, human-robot interaction and interactive machine learning.

**W. Bradley Knox** is a Postdoctoral Associate at the MIT Media Lab. He received a PhD in Computer Science at the University of Texas at Austin. His research interests span machine learning, robotics, psychology, and educational technology, especially machine learning algorithms that learn through human interaction. Dr. Knox has a B.S. degree in Psychology, completed pre-medical coursework as an undergraduate, and was an NSF Graduate Research Fellow from 2008–2011. His research won the Pragnesh Jay Modi Best Student Paper Award at AAMAS in 2010 and was a finalist for the CoTeSys Cognitive Robotics Best Paper Award at Ro-Man in 2012.

**Todd Kulesza** is a Computer Science PhD candidate at Oregon State University, where he is advised by Dr. Margaret Burnett. His research interests cluster around human interactions with intelligent systems, with a particular focus on how end users want to control such systems and the knowledge they need to do so efficiently. His work in this area was awarded a CHI 2012 Honorable Mention by SIGCHI.

**Tessa Lau** is a Research Scientist at Willow Garage, where she is leading an effort to develop simple interfaces for personal robots based on end user programming. More generally, Dr. Lau's research is in the area of intelligent user interfaces: combining techniques from artificial intelligence and human-computer interaction to create systems that enhance human productivity and creativity. She has served on organizing and program committees for major AI and HCI conferences and journals. She also serves on the board of CRA-W, the CRA committee on the status of women in computing research. Dr. Lau holds a PhD in Computer Science from the University of Washington.

**REFERENCES**

1. AAAI Fall Symposium 2012: Robots Learning Interactively from Human Teachers. <http://www.cs.utexas.edu/bradknox/AAAI-FSS-RLIHT12/>.
2. AAAI Spring Symposium 2009: Agents that Learn from Human Teachers. <http://www.cc.gatech.edu/AAAI-SS09-LFH/>.
3. AAMAS 2010 Workshop: Agents Learning Interactively from Human Teachers. <http://www.cs.utexas.edu/bradknox/AAMAS-ALIHT>.
4. CHI 2012 Workshop: End-user interactions with intelligent and autonomous systems. <http://www.city.ac.uk/informatics/school-organisation/centre-for-human-computer-interaction-design/enduserias>.
5. ECAI 2012 Workshop: Machine Learning for Interactive Systems. <http://www.sfbtr8.spatial-cognition.de/mlis-2012/Overview.html>.
6. IJCAI 2011 Workshop: Agents Learning Interactively from Human Teachers. <http://www.cs.utexas.edu/bradknox/IJCAI-ALIHT11/>.
7. NIPS 2009 Workshop: Analysis and Design of Algorithms for Interactive Machine Learning. <http://nips.cc/Conferences/2009/Program/event.php?ID=1503>.
8. Amershi, S., Lee, B., Kapoor, A., Mahajan, R., and Christian, B. Cuet: Human-guided fast and accurate network alarm triage. In *In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (2011), 157–166.
9. Cakmak, M., Chao, C., and Thomaz, A. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development* 2, 2 (2010), 108–118.
10. Druck, G., Settles, B., and McCallum, A. Active learning by labeling features. In *In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2009), 81–90.
11. Kapoor, A., Lee, B., Tan, D., and Horvitz, E. Interactive optimization for steering machine classification. In *Proceedings of the 28th international conference on Human factors in computing systems (CHI)* (2010), 1343–1352.
12. Knox, W., and Stone, P. Reinforcement learning from human reward: Discounting in episodic tasks. In *Proceedings of the 21st IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN09)* (2012).
13. Rosenthal, S., and Dey, A. Towards maximizing the accuracy of human-labeled sensor data. In *Proceeding of the 14th international conference on Intelligent user interfaces (IUI)* (2010), 259–268.
14. Settles, B. *Active Learning*. Morgan & Claypool, 2012.
15. Stumpf, S., Sullivan, E., Fitzhenry, E., Oberst, I., Wong, W., and Burnett, M. Integrating rich user feedback into intelligent user interfaces. In *Proceedings of the 13th international conference on Intelligent user interfaces (IUI)* (2008), 50–59.