Editorial Introduction to the Special Articles in the Summer Issue

Architectures for Activity Recognition and Context-Aware Computing

Christopher Geib, Vikas Agrawal, Gita Sukthankar, Lokendra Shastri, Hung Bui

■ *This editorial to the summer 2015 AI Magazine introduces the special-issue articles on architectures for activity recognition and context-aware computing.*

The changes brought about by the ubiquity of smartphones and social media are just a small foretaste of changes to come. Soon people will be carrying devices and working in environments that understand not only our personal declarative and demographic facts (information stored in datebooks, calendars, and social media) but also have a deep understanding of the context and intent of our day-to-day activities. The last 10 years have seen the development of novel architectures and technologies for domain-focused, task-specific systems that know many things, such as who (identities, profile, history) they are with (social context) and in what role (responsibility, security, privacy); when and where (event, time, place); why (goals, shared or personal); how are they doing it (methods, applications); and using what resources (device, services, access, and ownership).

Smart spaces and devices will increasingly use such contextual knowledge to help users move seamlessly between devices and applications, without having to explicitly carry, transfer, and exchange activity context. Such systems will qualitatively shift our lives both at work and play and significantly change our interactions both with our physical and virtual worlds.

This dream of seamlessly interacting with our virtual environment has a long history as can be seen in Apple Inc.'s Knowledge Navigator 1987 concept video. However, the combination of dramatic progress in low-power mobile computing devices and sensors, with advances in artificial intelligence and human-computer interaction (HCI) in the last decade, have provided the kind of platforms and algorithms that are enabling context-aware virtual personal assistants that plan activities and recognize intent. This has lead to an increase in work designed to bring these ideas into real world application and address the final technical hurdles that will make such systems a reality.

Example research projects in this area include Patie Maes's work on SixthSense interfaces (Maes 2009; Heun, Hobin, and Maes 2013) and Henry Kautz's work on intelligent assisted cognition systems (Kautz et al. 2002; Kautz et al. 2003; Patterson et al. 2004a; Patterson et al. 2004b; Patterson et al. 2007). The work of Doug Lenat and others (Panton et al. 2006; Lenat et al. 2010) on the Cyc project for codifying commonsense knowledge has proved to be an asset for such work as well, as an understanding of commonsense knowledge is valuable for these systems.

The promise of such interfaces and environments has also been appreciated by industry and government, and has led to much larger research projects. For example, DARPA's Personal Assistant that Learns (PAL) project that developed the Cognitive Assistant that Learns and Organizes (CALO) (Myers et al. 2007; Bui et al. 2008; Tur et al. 2008; Yorke-Smith et al. 2009; Yorke-Smith et al. 2012). It is worth noting that parts of CALO were spun off as Siri[1] and popularized by Apple Inc. as one of the most impressive currently deployed personal assistive systems.

SRI's RADAR project is another example (Faulring et al. 2008; Yoo et al. 2008; Freed et al. 2011), as is Vulcan Inc.'s HALO project (Friedland et al. 2004), which was part of Vulcan's Digital Aristotle project. Even the widely noted recent success of IBM's question-answering system Watson (Ferrucci et al. 2010, Ferrucci 2012) at Jeopardy displays the kind of ability to answer day-to-day questions and some of the basic commonsense reasoning that is an enabling technology for a general-purpose intelligent assistant system.

Along with these research projects we have also seen more work on technologies critical to developing these systems reported in a diverse group of mainstream research conferences including those of the Association for the Advancement of Artificial Intelligence (AAAI), the International Joint Conference on Artificial Intelligence (IJCAI), the ACM Special Interest Group on Computer-Human Interaction (CHI), the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), the User Modeling conference, the Uncertainty in Artificial Intelligence conference (UAI), the International Conference on Automated Planning and Scheduling (ICAPS), and others. However, the interdisciplinary nature of the research required to build these kinds of systems has meant that such work does not neatly fit at any single conference.

As a result two ongoing workshop series have sprung up focused on some of the issues. The first of these is the Plan, Activity, and Intent Recognition workshops (2007–2009, 2011, 2013) that evolved from the Modeling Others from Observations workshops (2004–2006). These workshops were specifically proposed to bring together researchers working on the automated recognition of an agent's activities and plans and the intentions that motivated those actions.

Such recognition of the plans and goals of another agent is one of the fundamental abilities we associate with intelligence. Performing this task involves making inferences about intelligent entities based on prior knowledge about the world and observations of the agent's behavior, the agent's interaction with the environment and other agents. Correctly performing this task allows people to identify what another person is doing, why the person is doing it, and predict his or her next actions. Inferring other humans' intentions is required for almost all interpersonal interactions. Thus as machines become ever more capable, and we want them to interact with us in ever more complex situations, the ability of machines to recognize, understand, and predict our actions has become critical.

This research is often broken down into three subareas. *Activity recognition* refers to the problem of segmenting and classifying low-level data gathered by cameras or wireless sensors into a description of a single activity (such as walking). *Plan recognition* refers to the mapping of sequences of atomic actions to high-level plans stored in a plan library. *Intention recognition* is the problem of identifying the high-level goals of action; for example, predicting that a software user is looking for a file or that a soccer player will kick left rather than right. Commonly used techniques discussed at the workshops for plan and activity recognition include theoretical frameworks of abduction, grammatical methods, probabilistic graphical models, partially observable Markov decision processes, adapting single-agent frameworks to multiagent cases, and cognitive models of social reasoning.

The second workshop series in this area is the Activity Context-Aware Systems Workshop series (2011, 2012, and twice in 2013). This workshop has focused on other enabling technologies for such systems including system architectures, standardization efforts for representations of context, identification of compelling use cases, domain-specific reasoning algorithms, and proposals for languages, data structures, and algorithms for representing and reasoning about activity context. The workshop also focused on issues like semantic computing, task modeling, context representation, creating fresh methods for capture, transfer, and recall of activity context across multiple platforms, supporting both individuals and groups. The workshop addressed questions about possible barriers to the adoption of such technologies like privacy, scalability, and the proprietary nature of some platforms.

The Activity Context-Aware Systems workshop organizers were also very interested in creating results that were usable by both industry and academia, believing that such collaboration was the only way to create standards for application and device context transfer through peer-to-peer technologies and services in the cloud. They have also encouraged significant cooperative research and development of pervasive computing applications by working toward a common language for interchange of activity context, task, and activity model representations and context representations for assistive cognition devices, the digital workplace, and the consumer play space.

Both of these workshop series have focused on areas of research that are critical if we are to build smart environments. Such environments will help users reason and make better decisions faster and with greater confidence. They will do this by understanding the tasks users are engaged in and giving the user understanding of the provenance, quality, and derivation of recommended information; understanding of context specific actions, dialogs, commands, and queries; and proactive advice based on steps others have previously taken in similar situations.

With these special articles in *AI Magazine,* we bring together extended versions of selected papers from the 2013 Plan, Activity, and Intent Recognition workshops and the Activity Context-Aware Systems workshops to highlight the state of the art in the technologies that enable these new methods of interacting with virtual and cyber physical systems. While there were a great many papers in these workshops that are worthy of inclusion here, we have chosen seven that exemplify four general themes that run throughout the work at both workshops: (1) Activity Modeling, Representation, Plan Recognition, and Intent Prediction; (2) Context-Aware Activity Guidance; (3) Context Exchange, Integration, and Security; and (4) Context Capture, Storage, Transfer, Retrieval, Management, and Presentation Systems.

Activity Modeling, Representation, Plan Recognition, and Intent Prediction

This first theme asks which human activities can be reliably detected. What representation frameworks are suitable for modeling activities and context switching, and enable uniform context recall universally (across devices, platforms, and technologies)? What types of, and how much, context information can be captured and incorporated into activity models?

Oriel Uzan, Reuth Dekel, Or Seri, and Ya'akov Gal's article, Plan Recognition for Exploratory Learning Environments Using Interleaved Temporal Search, presents novel algorithms for inferring users' activities in a class of flexible and open-ended educational software called exploratory learning environments (ELEs) that support interaction styles including exogenous actions and trial and error, providing a rich educational environment for students but challenging teachers to keep track of students' progress and to assess their performance. The authors present techniques for recognizing students' activities in such pedagogical software and visualizing these activities to students. It describes a new plan-recognition algorithm that uses a recursive grammar that takes into account repetition and interleaving of activities. This algorithm was evaluated empirically using two ELEs for teaching chemistry and statistics used by thousands of students in several countries. It was able to perform comparably to, or outperform, the state-of-the-art plan-recognition algorithms for both of these settings when compared to a gold standard that was obtained by a domain expert.

Christopher W. Geib and Christopher E. Swetenham's article, Parallelizing Plan Recognition, exploits the opportunity provided by modern multicore computing devices to parallelize plan recognition algorithms to decrease run time. Viewing plan recognition as parsing based on a complete breadth first search, makes their engine for lexicalized intent recognition (ELEXIR) (Geib 2009, Geib and Goldman 2011) particularly suited for parallelization. Geib and Swetenham document the extension of ELEXIR to utilize such modern computing platforms discussing multiple possible algorithms for distributing work between parallel threads and the associated performance wins. They show that the best of these algorithms provides close to linear speedup (up to a maximum number of processors), and that features of the problem domain have an impact on the achieved speedup.

Context-Aware Activity Guidance

The second theme looks at questions of how to model and represent activities, objects, resources, actions, and their semantics in their context during task performance. How do we design activity and context models to enable the searching of repositories of previous activities that have behaviorally and semantically similar components to current activity requirements? What does it take to combine numerical (and subsymbolic) and knowledge-driven (symbolic) approaches for reasoning, together with abductive reasoning, to create meaningful real-time guidance for users?

Knowledge workers perform work on many tasks per day and often switch between tasks. When performing work on a task, a knowledge worker must typically search, navigate, and dig through file systems, documents, and emails, all of which introduce friction into the flow of work. Mik Kersten and Gail C. Murphy show in their article, Reducing Friction for Knowledge Workers with Task Context, how this friction can be reduced, and productivity improved, by capturing and modeling the context of a knowledge worker's task based on how the knowledge worker interacts with an information space. Captured task contexts can be used to facilitate switching between tasks, to focus a user interface on just the information needed by a task, and to recommend potentially other useful information. They report on the use of task contexts and the effect of context on productivity for a particular kind of knowledge worker, software developers, with qualitative findings of the use of task contexts by a more general population of knowledge workers.

Context Exchange, Integration, and Security

The third theme looks at how can we integrate and exploit the growing amount of information available from devices, services, the environment, and the various sources of general background knowledge, in order to support activity context recognition tasks. What common ontologies or data vocabularies will be useful? What communication techniques and formalisms will be most effective in specific domains? How can the externalized cognitive state transfer be properly affected?

For humans and automation to effectively collaborate and perform tasks, all participants need access to a common representation of potentially relevant situational information, or context. In their article, A General Context-Aware Framework for Improved Human System Interactions, Stacy Lovell Pfautz, Gabriel Ganberg, Adam Fouse, and Nathan Schurr describe a general framework for building context-aware interactive intelligent systems that consists of three major functions: (1) capture human system interactions and infer implicit context; (2) analyze and predict user intent and goals; and (3) provide effective augmentation or mitigation strategies to improve performance, such as delivering timely, personalized information and recommendations, adjusting levels of automation, or adapting visualizations. The authors' goal is to develop an approach that enables humans to interact with automation more intuitively and naturally and that is reusable across domains by modeling context and algorithms at a higher level of abstraction. They first provide an operational definition of context and discuss challenges and opportunities for exploiting context; then they work toward a general platform that supports developing context-aware applications in a variety of domains. Pfautz and her colleagues explore an example use case illustrating how this framework can facilitate personalized collaboration within an information-management and decision support tool.

Laura Zavala, Pradeep K. Murukannaiah, Nithyananthan Poosamani, Tim Finin, Anupam Joshi, Injong Rhee, and Munindar P. Singh have been developing a high-level, semantic notion of location called *place* that is described in their article, Platys: From Position to Place-Oriented Mobile Computing. In their view, a place, unlike a geospatial position, derives its meaning from a user's actions and interactions in addition to the physical location where they occur. For this purpose they have considered elements of context that are particularly related to mobile computing. The authors are enabling the construction of a large variety of applications that take advantage of place to render relevant content and functionality and, thus, improve user experience. The main problems they have addressed in this article to realize this place-oriented mobile computing vision are representing places, recognizing places, and engineering place-aware applications. A key element of their work in Platys is the use of collaborative information sharing where users' devices share and integrate knowledge about places. The place ontology described in their article facilitates such collaboration. Declarative privacy policies allow users to specify contextual features under which they prefer to share or not share their information.

Computational management of activities that reflect human intention through activity-based computing (ABC) is described by Jakob E. Bardram, Steven Jeuris, and Steven Houben in their article, Activity-Based Computing: Computational Management of Activities Reflecting Human Intention. ABC applies to traditional desktop computing, ubiquitous or pervasive computing, and even wearable and tangible computing. It has emerged as a response to the traditional application- and file-centered computing paradigm, which is oblivious to a notion of a user's activity context spanning heterogeneous devices, multiple applications, services, and information sources. In their article, the authors present ABC as an approach to contextualize information, which is a common problem addressed in artificial intelligence as well.

Context Capture, Storage, Transfer, Retrieval, Management, and Presentation Systems

The final theme explores how automated context capture can be made and to what extent it will require collaborative metadialogue between people and devices. What might be ways of determining the most relevant elements of context for a given task and for an activity or context switch?

Simon Scerri, Jeremy Debattista, Judie Attard, and Ismael Rivera discuss the possibilities in their article, A Semantic Infrastructure for Personalizable Context-Aware Environments. Although a number of initiatives provide personalized context-aware guidance for niche use cases, a standard framework for context awareness remains lacking. This article explains how semantic technology has been exploited to generate a centralized repository of personal activity context. This data drives advanced features, such as (a) personal situation recognition, and (b) customizable rules for the context-sensitive management of personal devices and data sharing. As a proof of concept, the authors demonstrate how an innovative context-aware system has successfully adopted such an infrastructure.

We hope this issue of *AI Magazine* will help enhance collaborative efforts and joint research at the intersection of artificial intelligence and human-computer interaction to resolve multiple issues in context-aware computing and plan, activity, and intent recognition.

Notes

1. See Gruber, T. 2009. Siri, A Virtual Personal Assistant — Bringing Intelligence to the Interface. tomgruber.org/writing/Siri-SemTech09.pdf

References and Suggestions for Further Reading

Bardram, J. E. 2011. The Activity-Based Computing Project. In Activity Context Representation: Papers from the 2011 AAAI Workshop. Technical Report WS-11-04. Palo Alto, CA: AAAI Press.

Bardram, J. E., and Hansen, T. R. 2010. Context-Based Workplace Awareness. *Computer Supported Cooperative Work* 19(2): 105–138. dx.doi.org/10.1007/s10606-010-9110-2

Bolchini, C.; Curino, C. A.; Quintarelli, E.; Schreiber, F. A.; and Tanca, L. 2007. A Data-Oriented Survey of Context Models. *ACM SIGMOD Record* 36(4): 19–26. dx.doi.org/10.1145/1361348.1361353

Buchholz, S.; Hamann, T.; and Hubsch, G. 2004. Comprehensive Structured Context Profiles: Design and Experiences. In *Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops*. Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Bui, H. H.; Cesari, F.; Elenius, D.; Morley, D. N.; Natarajan, S.; Saadati, S.; Yeh, E.; and Yorke-Smith, N. 2008. A Context-Aware Personal Desktop Assistant. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Demo Papers*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Calvary, G.; Coutaz, J.; and Nigay, L. 1997. From Single-User Architectural Design to PAC*: A Generic Software Architecture Model for CSCW. In *Proceedings of the ACM SIGCHI 1997 Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery.

Cappiello, C.; Missier, P.; Pernici, B.; Plebani, P.; and Batini, C. 2004. QoS in Multichannel IS: The MAIS Approach. Paper presented at the 2004 International Workshop on Web Quality, Munich, Germany, 27 July.

Chen, H.; Finin, T.; and Joshi, A. 2003a. An Intelligent Broker for Context-Aware Systems. Paper presented at the Fifth International Conference on Ubiquitous Computing, Seattle, WA, 12–15 October.

Chen, H.; Finin, T.; and Joshi, A. 2003b. An Ontology for Context-Aware Pervasive Computing Environments. *The Knowledge Engineering Review* 18(03): 197–207. dx.doi.org/10.1017/S0269888904000025

Chen, H.; Perich, F.; Finin, T.; and Joshi, A. 2004. SOUPA: Standard Ontology for Ubiquitous and Pervasive Applications. In *Proceedings of the First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services,* 258–267. Piscataway, NJ: Institute for Electrical and Electronics Engineers.

De Virgilio, R., and Torlone, R. 2005. An Extensible Technique for Content Adaptation in Web-Based Information Systems. In *Mobile Information Systems II: IFIP International Working Conference on Mobile Information Systems,* 107–121. Berlin: Springer:

De Virgilio, R.; Torlone, R.; and Houben, G.-J. 2007. Rule-Based Adaptation of Web Information Systems. *World Wide Web* 10(4): 443–470. dx.doi.org/10.1007/s11280-007-0020-2

Dey, A. K.; Hamid, R.; Beckmann, C.; Li, I.; and Hsu, D. 2004. A CAPpella: Programming by Demonstration of Context-Aware Applications. In *Proceedings of the ACM SIGCHI 2004 Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery.

Faulring, A.; Myers, B. A.; Mohnkern, K.; and Freed, M. 2008. A Demonstration of the RADAR Personal Assistant. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence,* 1876–1877. Menlo Park, CA: AAAI Press.

Ferrucci, D. 2012. Introduction to "This Is Watson." *IBM Journal of Research and Development* 56(3,4).

Ferrucci, D.; Brown, E.; Chu-Carroll, J.; Fan, J.; Gondek, D.; Kalyanpur, A. A.; Lally, A.; Murdock, J. W.; Nyberg, E.; and Prager, J. 2010. Building Watson: An Overview of the DeepQA Project. *AI Magazine* 31(3): 59–79.

Finin, T.; Fritzson, R.; McKay, D.; and McEntire, R. 1994. KQML as an Agent Communication Language. In *Proceedings of the Third International Conference on Information and Knowledge Management*. New York: Association for Computing Machinery.

Freed, M.; Bryce, D.; Shen, J.; and O'Reilly, C. 2011. Interactive Bootstrapped Learning for End-User Programming. In Artificial Intelligence and Smarter Living: The Conquest of Complexity: Papers from the 2011 AAAI Workshop. Technical Report WS-11-07. Palo Alto, CA: AAAI Press.

Friedland, N. S.; Allen, P. G.; Matthews, G.; Witbrock, M.; Baxter, D.; Curtis, J.; Shepard, B.; Miraglia, P.; Angele, J.; and Staab, S. 2004. Project Halo: Towards a Digital Aristotle. *AI Magazine* 25(4): 29.

Gal, A., and Segev, A. 2006. Putting Things in Context: Dynamic eGovernment Re-Engineering Using Ontologies and Context. Paper presented at the 2006 WWW Workshop on E-Government: Barriers and Opportunities, Edinburgh, UK, 23 May.

Geib, C. W. 2009. Delaying Commitment in Plan Recognition Using Combinatory Categorial Grammars. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence,* 1702–1707. Palo Alto, CA: AAAI Press.

Geib, C. W., and Goldman, R. P. 2011. Recognizing Plans with Loops Represented in a Lexicalized Grammar. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence,* 958–963. Palo Alto: CA: AAAI Press.

Hang, C.-W.; Murukannaiah, P. K.; and Singh, M. P. 2013. Platys: User-Centric Place Recognition. In Activity Context-Aware System Architectures: Papers from the AAAI 2013 Workshop. AAAI Technical Report WS-13-05. Palo Alto, CA: AAAI Press.

Heun, V.; Hobin, J.; and Maes, P. 2013. Reality eDitor: Programming Smarter Objects. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing*. New York: Association for Computing Machinery.

Houben, S.; Bardram, J. E.; Vermeulen, J.; Luyten, K.; and Coninx, K. 2013. Activity-Centric Support for Ad Hoc Knowledge Work: A Case Study of Co-Activity Manager. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery. dx.doi.org/10.1145/2470654.2481312

Kaenampornpan, M., and O'Neill, E. 2005. Integrating History and Activity Theory in Context Aware System Design. Paper presented at the 1st International workshop on Exploiting Context Histories in Smart Environments, Munich, Germany, May 8.

Kaenampornpan, M., and O'Neill, E. 2004. Modelling Context: An Activity Theory Approach. *Ambient Intelligence: Second European Symposium,* Lecture Notes in Computer Science, 367–374. Berlin: Springer.

Kaenampornpan, M.; O'Neill, E.; and Ay, B. B. 2004. An Integrated Context Model: Bringing Activity to Context. Paper presented at the Workshop on Advanced Context Modelling, Reasoning and Management, Nottingham, UK, 7–10 September.

Kautz, H.; Arnstein, L.; Borriello, G.; Etzioni, O.; and Fox, D. 2002. An Overview of the Assisted Cognition Project. In Automation as Caregiver: The Role of Intelligent Technology in Elder Care: Papers from the AAAI 2002 Workshop. AAAI Technical Report WS-02-02. Palo Alto, CA: AAAI Press.

Kautz, H.; Etzioni, O.; Fox, D.; Weld, D.; and Shastri, L. 2003. Foundations of Assisted Cognition Systems. Technical Report CSE-03-AC-01, University of Washington, Computer Science Department, Seattle, WA.

Labrou, Y., and Finin, T. 1997. A Proposal for a New KQML Specification, Technical Report TR-CS-97-03, University of Maryland Baltimore County, Computer Science and Electrical Engineering Department. Catonsville, MD.

Lenat, D.; Witbrock, M.; Baxter, D.; Blackstone, E.; Deaton, C.; Schneider, D.; Scott J.; and Shepard, B. 2010. Harnessing Cyc to Answer Clinical Researchers' Ad Hoc Queries. *AI Magazine* 31(3): 13–32.

Maes, P. 1994. Agents That Reduce Work and Information Overload. *Communications of the ACM* 37(7): 30–40. dx.doi.org/10.1145/176789.176792

Maes, P. 2009. SixthSense: Integrating Information and the Real World. In *Proceedings, 8th IEEE International Symposium on Mixed and Augmented Reality.* Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Marchetti, C.; Pernici, B.; and Plebani, P. 2004. A Quality Model for Multichannel Adaptive Information. In *Proceedings of the 13th International World Wide Web Conference:* Alternate Track Papers and Posters. New York: Association for Computing Machinery.

Myers, K.; Berry, P.; Blythe, J.; Conley, K.; Gervasio, M.; McGuinness, D. L.; Morley, D.; Pfeffer, A.; Pollack, M.; and Tambe, M. 2007. An Intelligent Personal Assistant for Task and Time Management. *AI Magazine* 28(2): 47.

Panton, K.; Matuszek, C.; Lenat, D.; Schneider, D.; Witbrock, M.; Siegel, N.; and Shepard, B. 2006. Common Sense Reasoning — From Cyc to Intelligent Assistant. *Ambient Intelligence in Everyday Life,* 1–31. Berlin: Springer.

Patterson, D. J.; Fishkin, K.; Fox, D.; Kautz, H.; Perkowitz, M.; and Philipose, M. 2004. Contextual Computer Support for Human Activity. *Human Factors* 32(5): 527–536.

Patterson, D. J.; Kautz, H.; Fox, D.; and Liao, L. 2007. Pervasive Computing in the Home and Community. In *Pervasive Computing In Healthcare,* 79–103. Boca Raton, FL: CRC Press.

Patterson, D. J.; Liao, L.; Gajos, K.; Collier, M.; Livic, N.; Olson, K.; Wang, S.; Fox, D.; and Kautz, H. 2004. Opportunity Knocks: A System to Provide Cognitive Assistance with Transportation Services. *UbiComp 2004: Ubiquitous Computing,* Lecture Notes in Computer Science Volume 3205, 433–450. Berlin: Springer.

Preuveneers, D.; Van den Bergh, J.; Wagelaar, D.; Georges, A.; Rigole, P.; Clerckx, T.; Berbers, Y.; Coninx, K.; Jonckers, V.; and De Bosschere, K. 2004. Towards an Extensible Context Ontology for Ambient Intelligence. *Ambient Intelligence: Second European Symposium,* Lecture Notes in Computer Science, 148–159. Berlin: Springer.

Robinson, R.; Henricksen, K.; and Indulska, J. 2007. XCML: A Runtime Representation for the Context Modelling Language. In *Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops.* IEEE. Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Roussaki, I.; Strimpakou, M.; Kalatzis, N.; Anagnostou, M.; and Pils, C. 2006. Hybrid Context Modeling: A Location-Based Scheme Using Ontologies. In *Proceedings of the Fourth IEEE Annual Conference on Pervasive Computing and Communications Workshops.* Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Sohn, T., and Dey, A. 2003. iCAP: An Informal Tool for Interactive Prototyping of Context-Aware Applications, 974–975. In *Proceedings of the ACM SIGCHI 2003 Conference on Human Factors in Computing Systems.* New York: Association for Computing Machinery.

Strimpakou, M. A.; Roussaki, I. G.; and Anagnostou, M. E. 2006. A Context Ontology for Pervasive Service Provision. In *Proceedings of the 20th International IEEE Conference on Advanced Information Networking and Applications.* Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Sukthankar, G.; Goldman, R.; Geib, C.; Pynadath, D.; and Bui, H. 2014. *Plan, Activity, and Intent Recognition: Theory and Practice.* Amsterdam: Morgan Kaufman.

Tur, G.; Stolcke, A.; Voss, L.; Dowding, J.; Favre, B.; Fernandez, R.; Frampton, M.; Frandsen, M.; Frederickson, C.; and Graciarena, M. 2008. The CALO Meeting Speech Recognition and Understanding System. Paper presented at the IEEE Spoken Language Technology Workshop, 15–19 December.

Yoo, S.; Gates, D.; Levin, L.; Fung, S.; Agarwal, S.; and Freed, M. 2008. Using Semantic Features to Improve Task Identification in Email Messages. In *Natural Language and Information Systems,* Lecture Notes in Computer Science Volume 5039, 355–357. Berlin: Springer. dx.doi.org/10.1007/978-3-540-69858-6_43

Yorke-Smith, N.; Saadati, S.; Myers, K. L.; and Morley, D. N. 2009. Like an Intuitive and Courteous Butler: A Proactive Personal Agent for Task Management. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems,* volume 1. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yorke-Smith, N.; Saadati, S.; Myers, K. L.; and Morley, D. N. 2012. The Design of a Proactive Personal Agent for Task Management. *International Journal on Artificial Intelligence Tools* 21(01): 90–119. dx.doi.org/10.1142/S0218213012500042

Christopher Geib is an associate professor in the College of Computing and Informatics at Drexel University. He received his MS and Ph.D. in computer science from the University of Pennsylvania and a BS in computer science from Union College. His principal research interests are in artificial intelligence methods for intent recognition and planning more generally probabilistic reasoning about actions using formal grammars, and the interface between continuous control systems and logic-based reasoning systems. His work on intent recognition has been applied in multiple application areas including assistive care for the elderly, human robot interaction, and computer network security.

Vikas Agrawal is a principal data scientist and research analyst for cloud and big data at Infosys Limited. He received a BTech in electrical engineering from the Indian Institute of Technology, New Delhi (1997) and an MS (computer science) and a Ph.D. (computational systems biology) from the University of Delaware, and he conducted postdoctoral research at California Institute of Technology. His research interests lie at the intersection of AI and human-computer interaction in activity context-aware systems, virtual personal assistants, mathematical modeling, knowledge-driven data analytics and large simulations in manufacturing, utilities, life sciences, and finance.

Gita Sukthankar is an associate professor and Charles N. Millican Faculty Fellow in the Department of Electrical Engineering and Computer Science at the University of Central Florida, and an affiliate faculty member at the UCF Institute for Simulation and Training. She received her Ph.D. from the Robotics Institute at Carnegie Mellon, an M.S. in robotics, and an A.B. in psychology from Princeton University. Her research focuses on multiagent systems and computational social models. She is the recipient of the Air Force Young Investigator award, the DARPA Computer Science Study Panel, and an NSF CAREER award.

Lokendra Shastri is director of the Advanced Technology Labs at Samsung Research Institute – Bangalore. He has made significant contributions to massively parallel models of inference, knowledge representation, semantic networks, episodic memory, and computational neuroscience, and his research interests include machine learning, natural language processing, speech and visual pattern recognition, and decision making. He received his Ph.D. in computer science from the University of Rochester, an MS in computer science from the Indian Institute of Technology, Madras, and a B.E. (honors) in electronics (with distinction) from the Birla Institute of Technology and Science, Pilani. The author of a book and more than 80 invited chapters and peer-reviewed articles, he holds multiple patents. He has consulted for SRI International (as part of DARPA's Calo Project), the National Science Foundation, Cognitive Technologies Inc., and the United Nations Development Programme.

Hung Bui is a principal research scientist at the Laboratory for Natural Language Understanding, Nuance Communications in Sunnyvale, California. His technical expertise and interests include probabilistic graphical models, Bayesian inference, and machine learning, especially models for sequential and relational data, with applications in human activity and intent recognition, video understanding, and natural language processing. These interests center on the challenging problem of how to make the machine better understand the humans — a problem that plays a central part in intelligent assistive technologies, personal assistance, and dialogue management. He earned his Ph.D. from Curtin University (Australia) and was a senior computer scientist at SRI International, where he led the multiinstitution research team in developing probabilistic inference technologies for understanding user activities during the CALO project (the largest AI project in history, also known as the project that spun off Siri).