Reports of the 2013 AAAI Spring Symposium Series

Vita Markman, Georgi Stojanov, Bipin Indurkhya, Takashi Kido, Keiki Takadama, George Konidaris, Eric Eaton, Naohiro Matsumura, Renate Fruchter, Don Sofge, William F. Lawless, Omid Madani, Rahul Sukthankar

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the AAAI 2013 Spring Symposium Series, held Monday through Wednesday, March 25–27, 2013. The titles of the eight symposia were Analyzing Microtext; Creativity and (Early) Cognitive Development; Data-Driven Wellness: From Self-Tracking to Behavior Change; Designing Intelligent Robots: Reintegrating AI II; Lifelong Machine Learning; Shikakeology: Designing Triggers for Behavior Change; Trust and Autonomous Systems; and Weakly Supervised Learning from Multimedia. This report contains summaries of the symposia, written, in most cases, by the cochairs of the symposium.*

Analyzing Microtext

Much progress has been made in recent years in several areas within natural language processing. However, so far there has been less work related to microtext (for example, instant messaging, transcribed speech, and microblogs such as Twitter and Facebook). Microtext is made up of semistructured pieces of text that are distinguished by their brevity, informality, idiosyncratic lexicon, nonstandard grammar, misspelling, use of emoticons, and sometimes simultaneous interwoven conversation. These characteristics make microtext challenging to analyze. Most existing tools are trained on properly spelled and well-punctuated corpora, and therefore have problems correctly tagging and parsing microtext.

The 15 presentations focused on a broad range of microtext data sources: chat from online games, microblogs from Twitter, Facebook posts, and SMS communications. Some of the themes included creating a part of speech tagger for Twitter; sentiment extraction from tweets; gender and author detection in short noisy text; personality trait identification based on language used in social media; clustering of microtext by topic; detection of hedging and its relationship to gender, among many others. In addition to the contributed presentations and posters, the symposium included two invited talks from the leading researchers in microtext and social network analysis. Noah Smith (Carnegie Mellon University) spoke regarding the challenges and novelties of tagging and parsing microtext. His talk highlighted the need to reconsider what we call "noise" in data, for example, numerous abbreviations such as "SMH" (shake my head) and "OMG" that would be considered noise in some types of text are important "parts of speech" in Twitter and even warrant their own tag! Sofus Macskassy (Facebook) spoke about discovering Twitter users' topics of interest by examining the entities they mention in their tweets as well as various types of tweeting behavior (social banter versus event-based tweets). He also discussed an approach that leverages

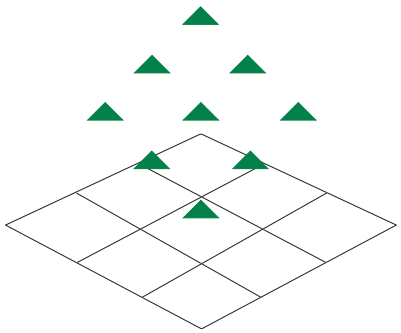

Wikipedia to disambiguate and categorize the entities in the tweets.

The symposium also included an invited panel of prominent researches in the area of microtext that was augmented by lively audience participation. The topic of the panel was the future of microtext. The panelists included Susan Herring (Indiana University), Bernardo Huberman (HP Labs), Rachel Greenstadt (Drexel University), and Alek Kolcz (Twitter). The panel included representatives from both academe and industry to give a fuller, more rounded perspective on the topic. Among the topics discussed were questions regarding the importance of analyzing microtext not just for research, but also for business. The discussion touched upon how improving tools for dealing with microtext can help inform business intelligence technologies of tomorrow. From a research perspective, we asked questions such as what defines microtext? Is social interactivity required (for example, the ability to comment or retweet) or can any news headline be considered mictortext just because it is short? When is microtext too long to be considered microtext? These questions were also echoed in the brilliant and engaging plenary talk by Doug Oard (University of Maryland).

An important question that emerged regarding the future of microtext research is whether microtext should be merged with other domains. This discussion lead to the observation that microtext research presently is fragmented across several research communities. Fostering interaction among this fragmented community is a challenge. There was support for the idea of associating microtext symposium with various conferences, as opposed to aligning with any specific conference. This would allow for maximal cross-pollination of ideas and ensure that the research in this domain is informed from various relevant disciplines.

David C. Uthus (postdoctoral fellow NRC/NRL), Craig Martell (Naval Postgraduate School), Ed Hovy (Carnegie Mellon University/USC ISI), and Vita Markman (Disney Interactive) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-13-01.

Creativity and (Early) Cognitive Development

Cross-domain general creativity is probably a uniquely human faculty. From a child who constructs a new toy using the old and broken ones, to the scientist who works out a theory and makes a profound impact on human civilization, the process invariably evokes the feelings of surprise, astonishment, and wonder. Though we understand what creativity is at an intuitive level, it has turned out to be quite difficult to define it formally and explore it scientifically.

Some researchers of creativity make a distinction between historical creativity (H-creativity) and psychological creativity (P-creativity), which is about small creative deeds, probably new only to the individual performing them. We hypothesize that they share the same basic cognitive mechanisms, and that creative perception (in viewing an artifact) involves the same mechanisms that are responsible for generating creative artifacts. Moreover, these mechanisms can also be observed during cognitive development: a constant reconceptualization of one's understanding of his or her environment in the process of agent-environment interaction, maturation, and education. If this hypothesis is accepted, then it suggests that by exercising and stimulating creative perception, we can also strengthen the ability to generate creative ideas and artifacts in the individual.

Broadly speaking, there were three research strands that were sort of braided in the symposium. The first strand concerns attempts to characterize creativity and connect these characterizations to cognitive development. For example, how does an oft-touted mantra of creativity, "think like a child," relate mechanisms of creativity to the child's cognitive development? How does seeing typical cognitive behavior as a series of small-c creative acts, which, in some individuals, continue in adulthood, lead to big-C creative deeds? How does symbolic play relate to creativity? How can we incorporate creativity in the action-based constructivist model of cognitive development championed by Piaget? How is distinction between strong and weak AI relevant for artificial creativity systems?

The second strand focused on the experimental methods for studying creativity, for example, developing an action-based approach to measure creativity in nonverbal toddlers as young as 19 months. Another issue is to study whether distractions promote divergent thinking thereby aiding creativity. A third research area in this strand is to study the cross-domain aspect of creativity: for example, are people able to recognize which works of art were inspired by which pieces of music?

The third strand of research represented in our symposium contained computational architectures for creativity, for example, attempts to incorporate metaphor-guided pretense play in the interaction between a child and humanoid robot; to model analogy to generate pedagogical explanations and novel compositions; to design a computational developmental agent based on Piaget's theories; to design a computational creativity system that allows researchers and developers to build ad hoc mash-ups of whatever processes and representations are most suited to a given application; to generate perceptually similar image pairs to stimulate the viewer's creativity and imagination, and so on.

Over the two-and-a-half days of intense discussion, we identified a few interdisciplinary research themes related to creativity and cognitive development that we hope to pursue in the near future. These included how to

characterize, measure, and model creativity in nonverbal and cognitively different agents like babies, young children, and autistic people. How to incorporate the effect of social context and norms in creativity? How to incorporate play and playlike activities in education to foster creativity? We all expressed our enthusiasm to pursue these questions in small interdisciplinary groups, and meet again in future symposia or workshops related to this theme.

Georgi Stojanov and Bipin Indurkhya served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-13-02

Data-Driven Wellness: From Self-Tracking to Behavior Change

The goal of the AAAI 2013 symposium on data-driven wellness was to explore approaches of improving personal wellness with AI technologies and understand ourselves with data-driven evidences. The symposium was a follow-up event to the 2012 AAAI spring symposium, Self-Tracking and Collective Intelligence for Personal Wellness. Because of the participants' strong desire to hold further events, we organized the symposium again, this time incorporating the concept of behavior change. The 2013 event focused on data-driven wellness, which derives behavior change of our daily life starting from self-tracking our health. For example, when we know our genome information (for example, the possibility of becoming diabetic) after testing our saliva (that is, self-tracking), we might decide not to eat high-calorie foods, or we might start running to keep or lose weight (that is, behavior change). The important stream of improving our health is promoted by the amount of data about our health (that is, data-driven wellness) acquired by current technologies (for example, calories calculation by smartphone). These streams contribute to creating societal or social activities on health improvement (for example, society or social activities that support diabetes patients). Our symposium explored such AI technologies and discussed possible solutions for our wellness. Our scope of interests included (1) self-tracking technology, (2) behavior change analysis and platform, (3) cognitive and biomedical modeling, (4) data-driven wellness and collective intelligence, and (5) wellness service applications and field study.

To promote discussion on possible solutions for these issues, our symposium brought together an interdisciplinary group of researchers, such as biomedical informatics, natural language processing, intelligent agent systems, human-computer interaction, brain science, cognitive psychology, social science, sport science, and behavior science. The symposium also included three invited talks and three guest talks to give us new perspective on data-driven wellness. For example, Rui Chen (Stanford University) introduced the cutting edge of biomedical research paradigms on integrative personalized omics toward precision medicine. Chirag Patel (Stanford University) presented the bioinformatics challenges for data-driven personalized medicine introducing recent hot topics on comprehensive view of the environment for predictive health. Koichi Hasida (National Institute of Advanced Industrial Science and Technology, Japan) gave a talk on a new platform on personal life repository, which promotes distributed personal data store (PDS) for data-driven improvement of our welfare. Sudheendra Hangal (Stanford University) presented the idea of gamification with personal data such as an individual's email logs. Yotam Heineberg (Stanford University) introduced the Stanford CCARE (the center for compassion and altruism research and education) project by presenting the online tracking of people's values and behaviors for compassionate behavioral activation. Palela Day (Stanford University) gave a talk on the role of mindfulness and compassion in behavior design.

A total of 19 technical papers were presented over the course of the two-and-a-half days. Presentation topics included (1) disease dynamics/detection and gene network, (2) brain science and brain interface (3) zone and kansei, (4) care/disable support system, (5) citizen science, people connection, and participation, (6) health data collection and personal behavior discovery. In each of these topic sessions, we scheduled long interactive discussions, which posed new interesting research questions from different points of view, such as our desire for people connection increases our self awareness, and lifestyle changes by recent social media changes our brain structure.

Our symposium provided participants unique opportunities where completely different background researchers presented new ideas in innovative and constructive discussions. We expect that the furure events will present important interdisciplinary challenges for guiding future advances in AI community.

Designing Intelligent Robots: Reintegrating AI II

Artificial intelligence is a largely fragmented field: research communities have developed and matured based on a focus on specialist areas like learning, planning, language, and vision. While the challenge posed by each of these areas alone is immense, this form of specialization often leads to technologies that are not well suited for integration into complete, intelligent agents.

Recent advances in robotics — most notably the availability of standard hardware platforms and open source software frameworks that result in truly reusable code — have brought issues of integration to the fore again. Rather than having our applications driven by what advances in specific technical subfields allow us to do, AI researchers are now in a position where they can drive progress in each subfield through the grand challenge of intelligent robotics, a challenge that speaks to the original impulse behind AI and presents us with an immensely rich source of research questions with real-world impact.

Last year's successful AAAI symposium, Designing Intelligent Robots: Reintegrating AI, brought together a diverse group of researchers interested in intelligent robotics, with the aim of forming a research community that cut across specialist areas. The attendees were able to share their results

and build cross-disciplinary collaborations for meeting unique challenges posed by building integrated systems. This year's symposium built on that progress, with a lively group of participants from all over the world, including five invited speakers from academe and industry and 14 contributed papers.

Discussions at the symposium led to a decision to join forces with recent European efforts to build a community centered around the ai-robotics Google group. The community has strengthened considerably in the last year, and participants were encouraged to submit papers to the AAAI Robotics track, the upcoming AAAI Workshop on Intelligent Robotic Systems, and an upcoming special issue of *Artificial Intelligence* on AI and robotics.

Byron Boots, Nick Hawes, Todd Hester, George Konidaris, Bhaskara Marthi, Lorenzo Riano, and Benjamin Rosman served as cochairs of this symposium. George Konidaris was the author of this report. The papers in the symposium were published as AAAI Press Technical Report SS-13-04.

Lifelong Machine Learning

Humans learn to solve increasingly complex tasks by continually building upon and refining knowledge over a lifetime of experience. This process of continual learning and transfer allows us to rapidly learn new skills, often with very little training. Over time, it enables us to develop a wide variety of complex abilities across many domains.

Despite recent advances in transfer learning and representation discovery, lifelong machine learning remains a largely unsolved problem. Lifelong machine learning has the huge potential to enable versatile systems that can learn continually over an unbounded stream of experience and rapidly acquire new skills by building upon previous knowledge. As Rich Sutton (University of Alberta) stated in his invited talk, "lifelong machine learning is the ultimate big data problem."

Learning over a lifetime of experience involves several procedures that must be performed continually, including (1) discovering representations from streaming data that capture higher-level abstractions, (2) transferring knowledge to accelerate learning, (3) accumulating and maintaining knowledge while avoiding catastrophic forgetting, and (4) incorporating feedback from the environment or other agents. The goal of this symposium was to bring together practitioners in each of these areas to investigate recent progress toward the development of lifelong machine learning, and to identify key challenges and future directions.

The symposium featured four invited talks, each of which advocated a different perspective on lifelong learning. Rich Sutton described a vision of lifelong learning that involves continual interaction with the world, in which the learner gradually accumulates predictive models of sensorimotor data using off-policy reinforcement learning. This view of lifelong learning, implemented in the Horde architecture, has no concept of "tasks," but instead supports learning useful predictive models from streaming data. Efficiency and scalability for continual learning ad infinitum are paramount, given the emphasis on robot learning in Sutton's work. In contrast, Jeff Dean (Google) presented a data and computation-intensive view of lifelong learning through his work on deep learning for image classification, speech recognition, and natural language processing. Dean emphasized the importance of learning layered representations from large amounts of unlabeled data and sharing these representations between learning tasks, distributing computation across many machines for speed and scalability.

Paul Ruvolo (Bryn Mawr College) presented a multitask learning perspective on lifelong learning, discussing his work with Eric Eaton on the Efficient Lifelong Learning Algorithm. This approach learns consecutive supervised or semisupervised tasks and supports continual improvement by sharing knowledge between tasks. Ruvolo also described a mechanism for incorporating self-direction into lifelong learning, in which the agent can actively choose the next task to learn in order to maximize performance. Matthew Taylor (Washington State University) described his work on teaching agents by demonstration, in which human teachers or even other agents share their knowledge through instruction. Taylor's vision would enable heterogenous agents to learn interactively, building upon the collective knowledge and advice of other lifelong learning agents.

Twelve peer-reviewed papers were presented during the symposium, exploring a range of complementary issues. Danny Silver (Acadia University) provided an overview of research toward lifelong learning, from early work on inductive transfer to continual reinforcement learning to current deep learning methods. Several presentations discussed potential objective functions for lifelong learning, including mechanisms for balancing between competing objectives. Many papers focused on reinforcement learning formulations of lifelong learning, exploring such issues as continual learning, scalability through structural policy transfer, spatiotemporal organization of knowledge, fast episodic recall, and autonomous selection of intertask mappings for knowledge transfer. To motivate research, Terran Lane (Google) described a variety of challenging problems in large-scale entity resolution that are suitable for lifelong learning.

One highlight of the symposium was a set of small working sessions for focused brainstorming on specific topics in lifelong learning, followed by large group conversations. Several key issues arose during these discussions. The first issue is scalability to support learning many diverse tasks from unbounded data streams. The need for scalability through shared representations, transfer learning, and knowledge maintenance is motivated by the learner having bounded resources. Another key issue is nonstationarity, both in the distribution of tasks or environments over time, as well as drift within each of these. Consequently, lifelong learners need the ability to discard obsolete knowledge while maintaining theoretical performance guarantees and avoiding catastrophic forgetting. Depending upon the application domain, there may or

may not be explicit learning tasks; in some cases, the agent may need to identify tasks autonomously. We also discussed representation discovery, distributed and parallel learning, the balance between transfer and simultaneous/multitask learning, intrinsic motivation, curriculum design, and collaborative learning.

All participants agreed that there is a current lack of benchmark data sets and standard metrics for lifelong learning. The group discussion identified several applications that are particularly compelling for lifelong learning, including robotic control, vision, natural language processing, machine translation, computational sustainability, and health care. In an effort to motivate research on lifelong learning through applications, the working session led by Danny Silver focused on organizing a sequential learning challenge; we would welcome participation from the broader AI community in organizing such an event. The symposium concluded with a discussion of open challenges critical to the development of lifelong machine learning.

Eric Eaton served as the chair of the symposium. The organizing committee included Terran Lane, Honglak Lee, Michael Littman, Fei Sha, and Thomas Walsh. The papers of the symposium were published as AAAI Press Technical Report SS-13-05.

Shikakeology: Designing Triggers for Behavior Change

How do you trigger learning by seeing? How do you encourage eco-conscious behaviors? How do you trigger health awareness? How do you encourage crime prevention? Our answer is "a shikake." *A shikake* is a Japanese word with various meanings related to triggers for behavior change, but in shikakeology it is defined as having the following three elements or features to clarify its meaning: (1) an embodied trigger for behavior change, (2) the trigger is designed to induce a specific behavior, and (3) the behavior deals with a social or personal issue. The embodied trigger is designed so that it is expected to be perceived and desirable. However, a shikake should be designed so as to not entice or trick anyone but explicitly or implicitly to encourage people to change their behavior by presenting possible alternative behaviors. The alternative behavior needs to be carefully designed so as to be acceptable and even desirable as an approach to the issue being dealt with.

The symposium invited two keynote speakers and 21 technical presentations. The presenters shared the knowledge, methods, experiments, and findings that demonstrate triggers to motivate people and lead to behavior changes. The first keynote presentation given by Jeremy Bailenson (Stanford University) discussed a line of research that leverages embodied experiences — immersive, multisensory, perceptually rich forays into virtual worlds. The second keynote presentation given by Hiroshi Ishiguro (Osaka University) introduced a series of androids and discussed philosophical questions related to the impact on interactions between human-operated androids and collocated people.

The symposium participants discussed concepts and implications of shikake from various points of view such as psychological or social mechanisms (human-environment interaction process, persuasive mechanism, game mechanism, marketing), implications for design (affordance, trigger categories, physical and virtual collaboration spaces, landscape ostranenie), human roles (curation, collaborative decision making), digital technologies (avatar-mediated interaction, human-computer interaction), and theoretical approaches (answer set prolog, abduction). The diversity of topics are indicators that shikakeology addresses a new direction of AI that bridges the gap among independent disciplinaries and the fact that AI is embedded in shikake. One of the takeaways from the workshop is the opportunity for harmony among human, object, environment, and AI, which is worth studying as shikakeology.

Naohiro Matsumura and Renate Fruchter organized this symposium. The papers of the symposium were published as AAAI Press Technical Report SSS-13-06.

Trust and Autonomous Systems

Trust is a key user issue in the development and implementation of autonomous systems, particularly for autonomous systems working with humans. Humans must be able to trust the actions of the machines to want to work with them in the first place, and machines must develop or establish trust in the actions of their human coworkers to ensure effective collaboration. Trust between autonomous robots (and other autonomous systems) and humans is also important. The goal of this symposium was to investigate trust relationships between humans and autonomous robotic systems.

The symposium brought together researchers from a variety of subfields of AI and robotics such as autonomous systems, human-robot interaction, social robotics, cognition, multiagent systems, and planning. Practitioners from other fields including social psychology, mathematics, and argumentation contributed to the discussions. One major focus was human interpersonal relationships and the meaning of trust between humans in various contexts, exploring how these may be (and sometimes are) extended to nonhumans, including pets and machines. A finding from several presenters was that intentionally anthropomorphizing autonomous robotic systems engendered trust. One speaker showed that humans are hormonally predisposed to establish trust relationships with other humans, and that hormones regulate behaviors in humans that are important in establishing trust. Further work was presented on the development of computational cognitive models of trust consistent with neuroimaging (MRI) studies of the human brain performing various trust-related tasks. Another major theme was the focus on humans working with autonomous systems in a heterarchical framework, with the conclusion that the most effective combinations include both human and machine intelligence. There was common agreement that standardized metrics for trust between humans and machines are needed.

The symposium participants exchanged many ideas about the meanings of trust and strategies for engendering trust in autonomous systems and concluded that this meeting was the beginning of a new discipline in robotics and autonomous systems. Several expressed that they would like to attend future symposia focusing on trust and autonomous systems.

Donald Sofge, William Lawless, and Geert-Jan Kruijff served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-13-07.

Weakly Supervised Learning from Multimedia

Weakly supervised learning (WSL) refers to situations where human feedback in training the system can be described as indirect, uncertain, and limited. Consider going from image- or video-level annotations to localization of objects or concepts in general (actions, events, and so on). For example, given videos tagged with a concept of interest, such as "boat," the task would be to learn where (and when) the instances of the concept occur within the video. Thus, weakly supervised algorithms could output video frames, bounding boxes that intersect the concept, or at an extreme, pixel-level annotations. Similarly, weakly supervised systems that process audio could temporally localize sounds of interest. The great attraction of such approaches is that weak or unreliable feedback is substantially cheaper to obtain than accurate labels at the detailed level of interest.

A wide variety of problems and solution strategies were discussed at the symposium. There were however a few major themes that united much of the work presented: that of using extra knowledge or constraints of various form (for example, originating from the specific domain or task) to make up for the lack of explicit supervision. For instance, for the task of assigning names to faces in a personal picture collection, two faces in the same picture cannot (usually) be of the same person, and pictures taken in close temporal proximity typically contain the same set of people. Domain knowledge of events (for example, birthday, wedding), can help provide constraints too. In work on retrieving actions and events of interest (such as "changing a tire," or "a wedding proposal") from a large video collection, it was shown how classifiers trained on different feature families (multiple views) can be late fused to provide adequate retrieval (and impressive) performance, even though such systems relied solely on relatively low-level audiovisual features. Furthermore, classifiers trained on a large vocabulary of mid- to lower-level concepts, not directly related to events of interest, were shown to help, in particular as the number of positive examples of the events of interest, that is, the labeled data, was reduced to fewer than a handful.

Other instances of constraints or extra knowledge included the use of: the horizon and related geometric constraints for scene understanding, "objectness" for better locating foreground objects, gaze prediction for improved action recognition in egocentric vision, constraining classes or clusters and taxonomic relations among classes (for example in imagenet), the hierarchical nature of patterns for unsupervised learning (in audio), and domain constraints together with depth and appearance cues in robotics for robust object discovery. A variety of techniques were presented including versions of multiple instance learning, online learning under a budget, and low-rank matrix factorization techniques for weakly supervised learning.

The symposium included two panel discussions. In the first, panelists Vitto Ferrari, Matthias Grundmann, Gang Hua, Kevin Murphy, and Harpreet Sawhney proposed data sets and tools to help researchers develop and evaluate approaches for weakly supervised learning. The second panel discussion was on the WSL problem and solution space, where Rita Singh and Bhiksha Raj presented challenges in speech and audio domains with some promising directions, and Irfan Essa and Sidd Srinivasa did the same in vision and robotics.

The final half-day of the symposium was hosted at Google Research, where participants had the opportunity to see relevant large-scale vision and audio research at Google, and to engage in unstructured discussions on WSL.

Learning and classification in the real world does not occur in a vacuum. While direct but costly manual feedback is often very sparse, various constraints in terms of implicit and explicit world knowledge can go a long way toward significantly improving learning performance.

Rahul Sukthankar was the chair and Omid Madani and James Rehg served as the cochairs for this symposium. The sessions were organized so as to provide ample opportunities for unstructured discussion. No symposium proceedings were published so as to enable authors to present their latest results prior to publication.

Vita Markman is affiliated with Disney Interactive Studios.

Georgi Stojanov is an associate professor at the American University of Paris, France.

Bipin Indurkhya is a professor at International Institute of Information Technology, Hyderabad, India.

Takashi Kido is a research manager of Rikengenesis in Japan.

Keiki Takadama is a professor of the University of Electro-Communications in Japan.

George Konidaris is a postdoctoral associate at MIT CSAIL.

Eric Eaton is a visiting assistant professor in the Computer Science Department at Bryn Mawr College.

Naohiro Matsumura is an associate professor at Graduate School of Economics, Osaka University.

Renate Fruchter is the founding director of the Project Based Learning Laboratory at Stanford University.

Donald Sofge is computer scientist at the Naval Research Laboratory in Washington, D.C.

William Lawless is a professor in the Department of Mathematics and Psychology at Paine College in Augusta, Georgia.

Omid Madani is a senior software engineer at Google.

Rahul Sukthankar is a staff scientist at Google.