# Reports on the 2017 AAAI Spring Symposium Series

*Jeannette Bohg, Xavier Boix, Nancy Chang, Vivian Chu, Elizabeth F. Churchill, Fei Fang, Jerome Feldman, Avelino J. González, Takashi Kido, William F. Lawless, José L. Montaña, Santiago Ontañón, Jivko Sinapov, Don Sofge, Luc Steels, Molly Wright Steenson, Keiki Takadama, Amulya Yadav*

■ *The Association for the Advancement of Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2017 Spring Symposium Series, held Monday through Wednesday, March 27–29, 2017, on the campus of Stanford University. The eight symposia held were Artificial Intelligence for the Social Good (SS-17-01); Computational Construction Grammar and Natural Language Understanding (SS-17-02); Computational Context: Why It's Important, What It Means, and Can It Be Computed? (SS-17-03); Designing the User Experience of Machine-Learning Systems (SS-17-04); Interactive Multisensory Object Perception for Embodied Agents (SS-17-05); Learning from Observation of Humans (SS-17-06); Science of Intelligence: Computational Principles of Natural and Artificial Intelligence (SS-17-07); and Well-Being AI: From Machine Learning to Subjectivity-Oriented Computing (SS-17-08). This report, compiled from organizers of the symposia, summarizes the research that took place.*

The AAAI Spring Symposium Series is an annual set of meetings run in parallel at a common site. It is designed to bring colleagues together in an intimate forum while at the same time providing a significant gathering point for the AI community. The two and one half day format of the series allows participants to devote considerably more time to feedback and discussion than typical one-day workshops. It is an ideal venue for bringing together new communities in emerging fields.

The symposia are intended to encourage presentation of speculative work and work in progress, as well as completed work. Ample time is scheduled for discussion. Novel programming, including the use of target problems, open-format panels, working groups, or breakout sessions, is encouraged. Reports of the symposia are published in the AAAI Technical Reports series and distributed to the participants. Most participants of the symposia were selected on the basis of statements of interest or abstracts submitted to the symposia chairs; some open registration is allowed. All symposia are limited in size, and participants are expected to attend a single symposium.

The eight symposia for 2017 were held March 27–29 on the campus of Stanford University.

 ISSN 0738-4602

## Artificial Intelligence for the Social Good

A rise in real-world applications of AI has stimulated significant interest from the public, media, and policy makers. Along with this increasing attention has come a media-fueled concern about purported negative consequences of AI, which often overlooks the societal benefits that AI is delivering and can deliver in the near future. To address these concerns, the symposium on Artificial Intelligence for the Social Good (AISOC-17) highlighted the benefits that AI can bring to society right now. It brought together AI researchers and researchers, practitioners, experts, and policy makers from a wide variety of domains.

The first focus of the symposium was to showcase research work on AI for social good, including work that was motivated by real-world problems in varied domains such as health care, social welfare, urban planning, and computational sustainability. The papers cover many different application domains with significant societal impacts. For example, under the theme of health care, an AI-based system for sanitizing large-scale medical records was discussed; a wandering detection system for Alzheimer suffering patients was presented. Under the topic of social welfare, algorithms for raising awareness about HIV among homeless youth were presented and their real-world deployment discussed; an algorithm for allocating police patrols to prevent crime in cities was shown. For urban planning, algorithms for smart city planning using crowd judgment analysis were presented; techniques for optimal detection of faulty traffic sensors used in route planning were presented. In the field of computational sustainability, techniques for learning the evolution of climate change were presented; animal population censusing techniques using crowdsourcing were discussed; machine-learning algorithms for predicting future poaching attacks from past poaching data in Uganda wildlife parks was illustrated.

The symposium also included two invited talks. A talk given by Carla Gomes (Cornell University) and Henry Kautz (University of Rochester) focused on AI for sustainability and public health. Another talk by Eric Rice (USC School of Social Work) and Sharad Goel (Stanford University) focused on AI for the social sciences.

The second focus of the symposium was to frame the area of AI for social good. With several discussion sessions, the symposium participants concluded that AI for social good is a burgeoning multidisciplinary research field with lots of potential for crosscutting collaboration among researchers in computer science, psychology, social work, criminology, and ecology. In this research area, the focus is on use-inspired research, on solving problems that really matter in the real world. Primarily, the focus is on uplifting fundamental human rights and values, and to focus on entities and domains that could benefit from the introduction of AI-based systems. It is also important to remember that having a very sharp distinction of AI for social good research is not always feasible, and often unnecessary. While there has been significant progress, there still exist many major challenges facing the design of effective AI-based approaches to deal with the difficulties in real-world domains. One of the challenges is interpretability since most algorithms for AI for social good problems need to be used by human end users. Second, the lack of access to valuable data that could be crucial to the development of appropriate algorithms is yet another challenge. Third, the data that we get from the real world is often noisy and biased. Fourth, the validation of AI algorithms in the real world is a big challenge. While it is relatively easy to validate algorithms in computer simulations, the process of moving this validation to the real world is not well known. Finally, it is important that algorithms developed in AI for social good often act as a decision support system for human end users, instead of completely replacing them.

The Artificial Intelligence for the Social Good symposium was organized by Eric Horvitz, Barbara Grosz, Amy Greenwald, David Parkes, Carla Gomes, Stephen Smith, Gregory Hager, Ann W. Drobnis, Nicole Sintov, Milind Tambe, Amulya Yadav (cochair), Fei Fang (cochair), and Bryan Wilder. This report was written by Amulya Yadav and Fei Fang. The papers presented at the symposium were published as AAAI Technical Report SS-17-01 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Computational Construction Grammar and Natural Language Understanding

Construction-based approaches to grammar take the basic unit of language to be a pairing of form and meaning, called a *construction*. Constructions cut across traditional linguistic boundaries, allowing the tight integration of morphosyntactic, semantic, and pragmatic information while ranging widely in size (from lexical items to multiunit expressions) and generality (from fixed idioms to productive grammatical patterns). This theoretical framework flexibly accommodates phenomena and constraints from a variety of perspectives, notably including observations about human language processing and learning from cognitive linguistics and psycholinguistics.

The aforementioned characteristics offer much promise for natural language applications, especially those that rely on deeper understanding of semantic relations, situational grounding, and pragmatic inference — in short, the integration of cues related to other aspects of intelligence. Despite this potential, the impact of construction-based approaches in the AI and computational linguistic communities has been limited thus far by three main factors: computational formalisms, linguistic resources, and research community.

Despite an active international community dedicated to construction grammars, research has been primarily linguistic in orientation and not designed to support computational implementation. In particular, formalizing structures and models that capture aspects of meaning and context raises many thorny challenges that go beyond those involved in syntactic parsing.

In addition, linguistic resources have been limited. Road-coverage constructional inventories and large-scale annotated corpora (analogous to grammars and treebanks) designed with a constructional viewpoint have not been widely available.

Finally, while a handful of teams have long pursued computational construction grammars, a broader interdisciplinary research community has yet to be established.

The last few years, however, have seen a growing number of constructionally inspired computational formalisms and implementations, along with multiple efforts to build larger-scale construction-based resources. This AAAI symposium was designed to nurture this emergent community of computational construction grammar and foster deeper connections with the broader allied fields of AI.

To that end, the first two days of the meeting were designed to showcase advances on various fronts while allowing ample time for richer discussion and practical demonstrations. Invited speakers focusing on linguistic issues included Adele Goldberg (Princeton University) on the foundations of construction grammar, William Croft (University of New Mexico) on event structure, and Hans Boas (University of Texas at Austin) on the development of construction inventories. Other linguistically oriented talks highlighted progress in the closely related field of frame semantics, with reports on FrameNet efforts for English, German, Brazilian Portuguese, and Swedish, as well as the creation and use of data resources specializing in, for example, prepositional semantics and causal relations.

Computationally oriented sessions featured a smorgasbord of implemented systems and formalisms varying in motivations and areas of emphasis (embodied construction grammar, fluid construction grammar, dynamic construction grammar, and template construction grammar, construction-based tree-adjoining grammar). These disparate systems also demonstrated the breadth of issues and applications that have been explored in a constructional context, tackling aspects of language processing, language learning, scalable language understanding systems, and the modeling of distributional regularities. Additional sessions drew connections between construction-based approaches and related fields, including talks featuring cognitive architectures and autonomous robotic systems and a panel on neurolinguistics organized by Michael Arbib (University of Southern California). Points of connection with more established natural language processing approaches were also explored, with James Allen (University of Rochester) suggesting more interaction with large-scale natural language understanding systems that could be viewed through a constructional lens.

The final morning shifted the focus to the future. A session organized by Nancy Chang (Google) and Tiago Torrent (Federal University of Juiz de Fora) addressed the development of more meaningful benchmarks for construction grammar and natural language understanding. Panel discussions were devoted to surveying candidate language phenomena to include in a community benchmark; metrics suitable for evaluating progress in construction-based understanding; natural language applications capturing relevant phenomena; and the feasibility of developing standards for semantic and construction annotation.

These topics spurred lively discussion and debate on the promise and the pitfalls of seeking measurable benchmarks. Participants differed, for example, in their commitment to modeling human cognition versus building practical applications; as well as in their interest in evaluating intermediate representations versus end-to-end results. Examples of shared tasks and benchmarks successfully driving progress in other areas of AI were balanced against cautionary tales of the unintended effects of prematurely emphasizing quantitative metrics over qualitative markers of progress. While these issues remain far from resolved, participants were in consensus on the value of continued fruitful dialogue in future iterations of the symposium, with the hope of reaching a wider community of researchers in AI, cognitive science, and computational linguistics.

The Computational Construction Grammar and Natural Language Understanding symposium was organized by Luc Steels and Jerome Feldman along with the program committee of Katrien Beuls, Nancy Chang, and Adele Goldberg. This report was written by Luc Steels, Jerome Feldman, and Nancy Chang. The papers presented at the symposium were published as AAAI Technical Report SS-17-02 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Computational Context: Why It's Important, What It Means, and Can It Be Computed?

Logic is context free. But context may explain human perceptions, thinking, and behavior. Context can be clear, uncertain, or illusory. Based on presentations at the symposium on computational context, context is clear when a narrative sufficiently coheres to produce meaning. Military rules of engagement demand clear contexts. Autonomous systems must operate in unambiguous contexts. Clear context for a team is how the team should be structured to perform a mission. Central governance can impose a clear context; however, imposition reduces innovation.

Uncertain contexts can derive from the computational complexity created by multiagent systems; confirmation bias from recommender systems increases uncertainty; and deception, anomalies, and privacy can make context uncertain, as can the interdependence between two complementary states. However, computational assistants can help to reduce uncertainty. Autonomous systems can manage uncertainty, but autonomous systems can also fail. Autonomous systems may manage physically complex contexts better than humans, until humans are introduced. Perturbations against a team may help to clarify context; for example, a competition between teams. But modeling perturbations, especially between multiple autonomous hybrid human-machine-robot systems, is a challenge.

Proactive decision systems are context dependent. However, context may also be an illusion. For example, humans misjudge Edward Adelson's checkerboard illusion, but photometers do not; and prey often overconsume resources without the predators that help to make an environment sustainable. Carlo Rovelli, a physicist, wrote "reality is not as it appears." In 1944, supporting Albert Einstein's theory of relativity, a *New York Times* editorial declared that the physical world was "largely illusory." After searching numerous behavioral and social data sets (for example, polls) for context, Ehtibar Dzhafarov and colleagues found no "evidence for contextuality." These conclusions support Jacob Bekenstein's claim that the holographic principle suggests human awareness of three-dimensional reality is really a two-dimensional illusion.

Whatever the context, individuals act differently whether alone or in a team. Can the computation of context enable AI to adapt to uncertain situations, to change across contexts, and to adapt to contexts with hybrid human-machine-robot teams? The Department of Defense needs hybrid teams to automatically have "a common perception of the surrounding world and able to place it into context." But does it matter whether we understand how deep machine learning solves problems for autonomous cars as long as these cars safely transport humans? Even if the context cannot be specified, if an AI program improves performance, does it matter whether context can be predetermined to be clear, uncertain, or illusory? Integrating systems to work together for autonomous agents and hybrid teams presents computational challenges, but from what we have learned at our symposium, addressing and solving these challenges with theory and in the field will advance the science of computing context for autonomy and autonomous teams.

The Computational Context symposium had 35 attendees, including 8 invited speakers, 12 regular speakers, and 4 impromptu speakers.

The Computational Context: Why It's Important, What It Means, and Can It Be Computed? symposium was organized by Ranjeev Mittu, W. F. Lawless, Don Sofge, and David Aha. This report was written by W. F. Lawless and Don Sofge. The papers presented at the symposium were published as AAAI Technical Report SS-17-03 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Designing the User Experience of Machine-Learning Systems

Machine learning is just one subarea of the broader artificial intelligence toolkit, but it is the one that has caught much public imagination of late. This symposium brought together a multifaceted group to explore what machine learning means for user experience: what challenges do we face in creating desirable, useful, usable, and reliable user experiences that incorporate machine-learning techniques? We invited participants to consider issues that lie at the intersection of machine learning and user experience. Questions included how to face the challenges of machine learning, such as what application and domain-specific challenges exist for experts who work with machine learning and predictive modeling? and observations from those who study the effects of such systems and services on people and their practices, and, ultimately, on social structures.

The symposium participants were as diverse as the topics we covered, hailing from industry and academe, from social and technical sciences, and from design. Attendees hailed from artificial intelligence's many subfields to engineering; computer science; HCI; interaction, UX, and product design; and sociology and anthropology. Participants were at different career stages and in a range of roles and levels in their various organizations.

The first day centered on paper presentations and discussion, the second day on participatory activities: a workshop exploring assumptions, fears, and hopes for machine learning, a panel with industry experts from Nissan, Ford, and Renault on autonomous vehicles, and a demonstration session that allowed participants to gain hands-on experience with some of the systems discussed in presentations. The third half-day was a session for discussion, reflection, and planning for next steps.

Papers presented addressed the impact of machine learning on a range of topics from the philosophical to the critical to the practical, and included new methods from the intersection of technology and design; the transformation of fashion recommendation systems; identifying key needs of experts who use machine learning; explorations into the psychology of end-user experiences with systems that use (or don't use) machine ; and issues of privacy.

Themes and topics included communication and collaboration; automation, agency, and control; and bias, trust, and power.

We addressed how to develop tools that support communication and collaboration between system, interaction, product, and service designers. How might we support a more productive dialog between those who apply machine-learning techniques, and those who understand the implications of the choices that developers and designers make in the design of these systems? We looked at getting beyond the black box — enabling better experiments in model training, and in tending and pruning data.

Reciprocal knowledge sharing will move both areas forward and will enable us to create more trusted and trustworthy user experiences. Bringing relevant and inclusive case studies that reflect a range of diverse use cases is one way for better formulation of design opportunities.

Spurred in part by the panel on autonomous vehicles, we discussed the difficulty of designing for complex ecosystems that are multidevice, multiservice, and interconnected — or sometimes disconnected? They all utilize their own forms of learning and predictive modeling, making for considerable design and user experience complexity, and need to work between technical, physical, and social layers.

We also tackled more philosophical and political issues. On the second day of the symposium, we discussed system transparency, with a call for clear provenance models that make explicit the potential biases in machine-learning data sets, sources, and interactions. The basic call was always to provide multiple points of view to mitigate issues to do with bias, and to make bias an explicit topic of investigation itself.

Trust and power were key issues closing the symposium. What dialogue should systems have with their users, and what does it mean for systems to be personable, to have character, and to be socially responsible? The symposium ended with a pledge to craft a summary monograph that will be published to augment the publications in the AAAI 2017 Spring Symposium Technical Report.

The Designing the User Experience of Machine Learning Systems symposium was organized by Mike Kuniavsky, Elizabeth Churchill, and Molly Wright Steenson. This report was written by Elizabeth Churchill and Molly Wright Steenson. The papers presented at the symposium were published as AAAI Technical Report SS-17-04 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Interactive Multisensory Object Perception for Embodied Agents

Learning to perceive and reason about objects in terms of multiple sensory modalities remains a long standing challenge in robotics. Evidence from the fields of psychology and cognitive science has demonstrated that humans rely on multiple sensory modalities (for example, audio, haptics, tactile) in a broad variety of contexts ranging from language learning to learning manipulation skills. Nevertheless, most object representations used by robots today rely solely on visual input due to the difficulty of robotic interaction. Relying on visual input does not allow robots to learn or reason about nonvisual object properties (weight, texture). The goal of the symposium was to investigate how multisensory object representations can be learned and used by robots through interaction with their environment.

The symposium brought together researchers from a variety of different fields: machine learning, developmental and cognitive robotics, assistive robotics, robotic manipulation and control, and neuroscience. The papers accepted to the symposium spanned a diverse set of problems and domains in which robots interact with the environment and utilize visual and nonvisual object representations. The research showed that multisensory perception can allow robots to learn a variety of skills and tasks and that such perception can complement computer vision techniques in situations where vision alone is insufficient.

Several speakers gave invited talks. Alexander Stoytchev discussed how exploratory behaviors coupled with multisensory perception enable autonomous mental development in robots. Byron Boots presented machine-learning models for state estimation and filtering in high-dimensional spaces. Charlie Kemp highlighted the practical benefits of multisensory perception in the domain of assistive robotics. Oliver Brock proposed a design pattern for using multimodal perception in the context of learning manipulation skills. Katherine Kuchenbecker presented methods that enabled robots to learn haptic properties of objects. Jivko Sinapov highlighted the importance of using multisensory perception when teaching robots language. The symposium also featured two talks, from neuroscientists Allison Yamanashi Leib and Moqian Tian.

The symposium attendees came together and addressed the major question of how do we collect large data sets from robots exploring the world with multisensory inputs and what algorithms can we use to learn and act with this data? This question was broken down into three main themes: (1) representations of multimodal robot knowledge; (2) learning for robot perception; and (3) the benefits of multisensory information and how to collect and share data within the community. Specific challenges within these topics include issues such as different sensors producing data at different sampling rates and different resolutions. Also, data produced by a robot acting in the world is typically not independently and identically distributed (a common assumption of machine-learning algorithms). Furthermore, hosting and sharing hardware and software for robot interactive systems often stopped within the cycle of a single PhD student.

We identified several action items to address these challenges and highlight the importance of multisensory robot learning in the robotics community. First, to unify the work in multisensory robot learning, the attendees agreed on selecting a single keyword *interactive multimodal perception,* which we plan to add to several major robotics venues. Second, we plan on hosting future workshops on interactive multimodal perception to strengthen the community of researchers pursuing this work. Third, we plan on creating a new technical committee within the robotics organization and utilize those resources to build a website and mailing list for researchers to share and disseminate their tools, data, and research. Finally, we are proposing a special journal issue on interactive multimodal perception where we hope extended versions of the symposium submissions will be submitted and published.

The Interactive Multisensory Object Perception for Embodied Agents symposium was organized by Vivian Chu, Jivko Sinapov, Jeannette Bohg, Sonia Chernova, and Andrea L. Thomaz. This report was written by Vivian Chu, Jivko Sinapov, and Jeannette Bohg. The papers presented at the symposium were published as AAAI Technical Report SS-17-05 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Learning from Observation of Humans

The goal of the Learning from Observation of Humans AAAI symposium was to advance the state of the art in learning from observation and related disciplines by bringing together researchers from a broad set of backgrounds and establishing bridges between the different communities working in these problems.

Learning from observation, also known as learning from demonstration, imitation learning, or behavioral cloning, and related to programming by demonstration and apprenticeship learning, studies how computers can learn to perform complex tasks by observing and thereafter imitating the performance of a (human) actor. Learning from observation offers the promise of allowing machines to learn how to perform complex behaviors that would be difficult to program manually (such as driving vehicles, robotic motion, or playing video games). For example, modern training, education, and entertainment applications make extensive use of virtual agents, which must display complex intelligent behaviors. Driving simulations, for example, require vehicles moving realistically; simulated military training environments require friendly and unfriendly forces to present realistic and intelligent tactical behaviors; and computer games require artificial characters that display believable behaviors to heighten immersiveness of the games. Handcrafting those behaviors requires a significant amount of resources and can be highly error prone, thus being practical only for small and well-defined behaviors. Learning from observation offers a promising alternative. After several decades of research, however, there are still a significant number of open research problems in learning from observation, including algorithmic challenges. One example is designing learning algorithms that do not make the independent and identically distributed assumption that is common in supervised learning, which is violated in most learning from observation settings. Other such examples include how to assess the resulting performance of agents that learned from observation.

The symposium brought together researchers from a variety of subfields of artificial intelligence, including learning algorithms; perception and action (mapping human actions to agent actions, vision, control); demonstration techniques; human-computer and human-robot interaction issues related to learning from observation, or learning from humans; interactive learning from humans; evaluation methodologies (behavior similarity, performance measures); and autonomous driving.

Michael Floyd was the plenary speaker at the symposium. He spoke about a framework to represent the common processes in learning from observation, with the goal of designing a general-purpose learning agent. Floyd also presented a paper, where he discussed a preliminary study about using deep learning to automate feature modeling in learning from observation. His presentation was based on a paper coauthored with J. T. Turner and David W. Aha. Alex Lascarides and Mihai Sorin Dobre described in their talk a supervised approach for learning policies in a highly complex game from small amounts of human data consisting of state-action pairs. Santiago Ontañón then presented a paper that focused on modeling and predicting human driving behavior, with the long-term goal of anticipating the behavior of the driver before dangerous situations occur. He formulated this problem as a learning from demonstration problem, and showed how standard supervised learning methods do not perform well in this task. Brandon Packard's presentation focused on learning from demonstration, and in particular, on the problem of active learning from demonstration in settings where the amount of data that can be acquired from the demonstrator is limited. A paper presented by José L. Montaña focused on cloning recognition problems in learning from observation. He discussed two experimental case studies. The first concerned discrete domains and the second the role of neural networks in autonomous driving. Finally, Amrik Sacha, Elapata Gunaratne, Babak Esfandiari, and Caleb Chan proposed a generic modular framework using multiple learning methods that will enable testing learning from observation in multiple domains. The framework compared an extension on case-based reasoning, called temporal backtracking with probabilistic graphical model–based learning methods, such as Bayesian networks, input-output hidden Markov models, dynamic Bayesian networks, neural networks, and time-windowed Bayesian networks.

At the end of the session the symposium participants held a discussion about several questions related to learning from observation, present and future issues, including the creation of a data set repository for learning from demonstration, which is currently being set up by the symposium organizers, with the goal of helping comparative studies in the field.

The Learning from Observation of Humans symposium was organized by Santiago Ontañón, Avelino J. González, and José L. Montaña. This report was written by Santiago Ontañón, Avelino J. González, and José L. Montaña. The papers presented at the symposium were published as AAAI Technical Report SS-17-06 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Science of Intelligence: Computational Principles of Natural and Artificial Intelligence

The motto of the AAAI Symposium on Science of Intelligence: Computational Principles of Natural and Artificial Intelligence was understanding intelligence through algorithms. Most of the discussions were around the most successful pattern recognition algorithms — deep neural networks — which remain indecipherable.

These are the times of record-breaking results of deep neural networks in recognizing objects, playing board and video games, and making unprecedented progress in modeling the sensory cortex. Yet, currently there is no theory that can predict the performance of deep neural networks based on theoretical deduction, before testing them. Deep neural networks are the awakening of Kant's dogmatic dream in the times of big data. "Is this the end of the theory?" asked participant David Donoho (Stanford University) in the symposium.

The symposium brought together about 80 experts and pioneers in the fields of artificial intelligence, cognitive science, and computational neuroscience. There were 17 keynote talks and 29 posters presented, in which two leitmotivs mixed with each other: algorithms (1) to reproduce the computations done in the brain and (2) to create new forms of artificial intelligence. These two are synthesized in the emerging field called science of intelligence, which is dedicated to developing a computational understanding of intelligence — both natural and artificial — and to establishing an engineering practice based on that understanding.

"Intelligence is much more than pattern recognition" pointed out Pat Langley (Institute for the Study of Learning and Expertise), in view of the recent success of deep neural networks. In his talk, Langley depicted the vintage aspirations of symbolic representations and cognitive systems. In fact, this school of thought led to recent spectacular progress on understanding and reproducing high-level human cognitive abilities, without using deep neural networks and neuroscience.

Further, the success of deep neural networks poses a philosophical puzzle. The lack of explanations for the mechanisms of deep neural networks contrasts with the tremendous progress of these algorithms in modeling the brain and advancing artificial intelligence. "We are using a model that we do not understand, to explain something we do not understand, the brain," noted Tomaso Poggio (Massachusetts Institute of Technology and the Center for Brains, Minds, and Machines) during the plenary session. May this be the rebellion of the algorithms?

A promising old idea to resolve the puzzle came to the surface. The idea is to build scientific theories — using hypothesis testing — from empirical observations of artificial intelligence algorithms. This view sees algorithms as an alien intelligence, and shifts the human-centered focus of intelligence to the algorithms — a Copernican turn.

The Science of Intelligence: Computational Principles of Natural and Artificial Intelligence symposium was organized by Gemma Roig and Xavier Boix. This report was written by Xavier Boix. The papers presented at the symposium were published as AAAI Technical Report SS-17-07 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

## Well-Being AI: From Machine Learning to Subjectivity-Oriented Computing

Well-being AI is an information technology that aims to promote psychological well-being (that is, happiness) and maximize human potential. Well-being AI provides a way to understand how our digital experience affects our emotions and our quality of life and how to design a better well-being system that puts humans at the center.

Recently, deep learning and other advanced machine-learning technologies have revolutionized computer vision, speech recognition, and natural language processing. Despite these advances, applying these AI revolutions to human health and wellness problems remains challenging. One of the largest challenges is to understand human subjective knowledge and design better health and well-being systems. We define *subjectivity-oriented computing* as an approach to designing and understanding computational systems by understanding human subjective knowledge. This symposium also discussed subjective intelligence by learning from the human self-awareness process.

More specifically, we explored the methods or methodologies for (1) representation of subjective knowledge; (2) deep learning and other quantitative methods for health and wellness; (3) models, reasoning, and inference; and (4) better well-being systems design.

Our symposium included eight invited talks that provided new perspectives on well-being computing. Alex Ranter (Stanford University) spoke on ameliorating the labeling bottleneck with weak supervision. Avanti Shrikumar (Stanford University) discussed the issues on interpretable deep learning for genomics. Michael Nova (Pathway Genomics, Inc.) discussed future perspectives on cognitive healthcare using AI. Hirokazu Shirado (Yale University) introduced research topics on human coordination in experimental social networks. Kenji Suzuki (University of Tsukuba) discussed the concept of subjectivity-oriented computing to understand and empower individuals. Atsushi Nakazawa (University of Kyoto) introduced his research on evaluation of care skills using eye-tracking technologies. Steve Cole (University of California, Los Angeles) spoke on social regulation of human genome expression. Finally, Guido Pusiol (Stanford University) introduced his research on systems that care for the elderly using motion detection technologies.

The technical presentations comprised 16 papers and 3 posters or demonstrations. Presentation topics included sensor-based well-being; interactive support systems for elderly persons; interactive support systems for visual impairment; body motion for well-being; meditation detection for well-being; sleep stage estimation for well-being; machine learning for well-being; feeling analysis for well-being; and visualization for well-being.

Takashi Kido (Preferred Networks) discussed the challenges for machine learning and subjective computing in well-being AI. Nicola Bellotto (University of Lincoln) introduced a system named ENRICHME as an ambient intelligence integration for elderly care robots. Tomoyuki Hiroyasu (Doshisha University) introduced the brain functional state analysis of mindfulness using graph theory and functional connectivity. Tomohiro Harada (Ritsumeikan University) proposed methods for improving the accuracy of real-time sleep stage estimation. Keiki Takadama discussed guidelines for applying machine learning to care support systems; and Camille Marie Ruiz (Nara Institute of Science and Technology) introduced an analysis of online activity and expressions on real-life relationships of lonely users.

The Well-Being AI symposium provided participants unique opportunities where researchers with completely different backgrounds were able to come up with new ideas through innovative and constructive discussions. The symposium presented important interdisciplinary challenges for guiding future advances in the AI community.

The Well-Being AI: From Machine Learning to Subjectivity-Oriented Computing symposium was organized by Takashi Kido and Keiki Takadama. This report was written by Takashi Kido and Keiki Takadama. The papers presented at the symposium were published as AAAI Technical Report SS-17-08 in the AAAI Digital Library and included in *The 2017 AAAI Spring Symposium Series: Technical Reports SS-17-01 – SS-16-08* compilation.

**Jeannette Bohg** is a senior research scientist at the Max Planck Institute for Intelligent Systems.

**Xavier Boix** is a postdoctoral fellow at the Massachusetts Institute of Technology.

**Nancy Chang** is a research scientist at Google.

**Vivian Chu** is a PhD candidate at the Institute of Robotics and Intelligent Machines, Georgia Institute of Technology.

**Elizabeth F. Churchill** is a director of user experience at Google, based in San Francisco and Mountain View, USA.

**Fei Fang** is a postdoctoral fellow in the Center for Research on Computation and Society (CRCS) at Harvard University.

**Jerome Feldman** is professor emeritus of the University of California at Berkeley.

**Avelino J. González** is a professor in the Department of Computer Science at the University of Central Florida.

**Takashi Kido** is a research manager of Preferred Networks in Japan.

**William F. Lawless** is a professor of mathematics and psychology at Paine College in Augusta, Georgia.

**José L. Montaña** is an associate professor in the Department of Mathematics, Statistics, and Computer Science at the University of Cantabria.

**Santiago Ontañón** is an assistant professor in the Department of Computer Science at Drexel University.

**Jivko Sinapov** is a postdoctoral associate at the Department of Computer Science, University of Texas at Austin.

**Don Sofge** is a roboticist at the Naval Research Laboratory.

**Luc Steels** is an ICREA research professor at the Institut de Biologia Evolutiva (UPF/CSIC) in Barcelona.

**Molly Wright Steenson** is an associate professor at Carnegie Mellon University in the School of Design.

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

**Amulya Yadav** is a PhD candidate in the Computer Science Department at the University of Southern California.