Reports of the AAAI 2016 Spring Symposium Series

Christopher Amato, Ofra Amir, Joanna Bryson, Barbara Grosz, Bipin Indurkhya, Emre Kiciman, Takashi Kido, W. F. Lawless, Miao Liu, Braden McDorman, Ross Mead, Frans A. Oliehoek, Andrew Specian, Georgi Stojanov, Keiki Takadama

■ *The Association for the Advancement of Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2016 Spring Symposium Series on Monday through Wednesday, March 21–23, 2016, at Stanford University. The titles of the seven symposia were (1) AI and the Mitigation of Human Error: Anomalies, Team Metrics, and Thermodynamics; (2) Challenges and Opportunities in Multiagent Learning for the Real World (3) Enabling Computing Research in Socially Intelligent Human-Robot Interaction: A Community-Driven Modular Research Platform; (4) Ethical and Moral Considerations in Nonhuman Agents; (5) Intelligent Systems for Supporting Distributed Human Teamwork; (6) Observational Studies through Social Media and Other Human-Generated Content, and (7) Well-Being Computing: AI Meets Health and Happiness Science.*

AI and the Mitigation of Human Error: Anomalies, Team Metrics, and Thermodynamics

Today, robots are building other robots that drive humans. Watching this rapid progress with artificial intelligence, many scientists, engineers, and citizens have become alarmed by what they see as an evolving, existential threat posed by AI to the human race. Among those expressing alarm include Stephen Hawking, Elon Musk, and Bill Gates.

Instead, our symposium focused on whether AI could help humans by reducing human error. Already, a few cars are being built with lane control; collision control; and control against harming a child while in reverse. For some years, fighter jets have been able to take control when a pilot passed out from excessive g-forces.

From an existential perspective, our symposium convened to discuss what could be done to protect human life with AI. Following are three examples of where we believe AI could have intervened. A Germanwings copilot committed suicide in 2015, killing all 150 aboard; an Amtrak train wrecked in the Northeast corridor of the United States in 2015 killing 8 and injuring more than 200; and a USS submarine commander in 2001 overrode his subordinates to perform a rapid ascent that destroyed a Japanese tour boat, killing nine.

What about drunk drivers? There are about 2.91 alcohol-related fatalities per 100,000 miles driven in the USA; 4.59 in Germany; and 12.25 in Russia. Can AI safely intercede, acting like an Uber driver by taking control? Can a hybrid team of cars and humans be constructed on the fly to navigate intersections? Generally, when should AI override a human making an error?

The symposium had four invited speakers. Julie Adams, Vanderbilt University, spoke on mitigating human error through intelligent decision support systems. James Llinas, Center for Multisource Information Fusion and SUNY at Buffalo, spoke about an argumentation-based system support toolkit for intelligence analyses. Stephen Russell, chief, Battlefield Information Processing Branch, U.S. Army Research Laboratory, spoke on error mitigation. Lastly, Martin Voshell, Charles River Analytics, spoke on multilevel human-autonomy teams for distributed mission management.

We had two panels. The first panel was on the Internet of things and error reduction. Llinas addressed the Internet of things — impacts to data and information fusion process design and realization. Russell addressed the Internet of things and error mitigation.

The second panel addressed general intelligence and Open AI. Its speakers were Vladimir Gontar, Ben-Gurion University of the Negev; and Olivier Bartheye, CREC St-Cyr, France.

The symposium was coorganized by Ranjeev Mittu (Naval Research Laboratory), Don Sofge (Naval Research Laboratory), Gavin Taylor (U.S. Naval Academy), and W. F. Lawless (Paine College). This report was written by W. F. Lawless. The papers from the symposium were published by AAAI Press as technical report SS-16-01, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Challenges and Opportunities in Multiagent Learning for the Real World

Multiagent systems are one of the most promising applications of artificial intelligence because many real-world problems can be thought of as consisting of multiple decision makers. However, programming multiagent systems is notoriously difficult. As a result, the need for learning agents in multiagent systems is even greater than in the single-agent case. Developing efficient methods for multiagent learning has therefore been a long-standing research focus in the artificial intelligence, game theory, control, and neuroscience communities.

As a growing number of agents are deployed in complex environments for scientific research and human well-being, there are increasing demands to design efficient learning algorithms that can be used in real-world settings. These domains are challenging due to reasons that include accounting for uncertainty, partial observability, sequential settings, and communication restrictions. These challenges exist in many domains, such as underwater exploration, planetary navigation, robot soccer, stock-trading systems, and e-commerce.

Many approaches to multiagent learning, however, focus on restricted aspects of the learning problem, or ignore certain difficulties that are important to address in real-world applications. This symposium brought together researchers to push research on multiagent learning to overcome some of these simplifying assumptions and address the fundamental issues that hinder the applicability of multiagent learning for solving complex, real-world problems.

The result of the symposium was a very exciting and interesting program with contributions and speakers on a great variety of aspects of multiagent learning. Applications ranged from large-scale multirobot teams in Amazon's semiautonomous warehouses, to coordinating with people in telepresence robotics, combating crime with security games, effective air traffic management, coordinated soft robotics for planetary navigation, and socially appropriate self-driving cars. On the theoretical side, people used a great variety of formal frameworks to describe their multiagent systems, such as grid games, MDPs, decentralized POMDPs, team decision problems, various control-theoretic models, and bioinspired or swarm models. Great progress has been made on applying these formal frameworks and machine-learning methods to large, challenging systems, developing systems that are both principled and perform well.

A number of themes resonated throughout the symposium, such as dealing with partial observability, distributed learning with communication restrictions, scaling to extremely large numbers of agents, system verification, and equipping artificial agents with the ability to learn about human behavior. These themes were identified as key bottlenecks for developing and deploying real-world multiagent systems. Much of the work in the symposium has begun overcoming these bottlenecks by methods such as using problem structure, extending or developing new formal models, and extending methods from single-agent learning and game theory, but many open questions remain. We look forward to contributing to and monitoring the future multiagent learning methods as approaches become more sophisticated and more systems are deployed.

The symposium was organized by Chris Amato, Frans Oliehoek, Miao Liu, Karl Tuyls, Peter Stone, and Jon How. This report was written by Christopher Amato, Frans A. Oliehoek, and Miao Liu. The papers from the symposium were published by AAAI Press as technical report SS-16-02, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Enabling Computing Research in Socially Intelligent Human-Robot Interaction: A Community-Driven Modular Research Platform

The field of human-robot interaction (HRI) has grown significantly in the last decade and a half, and actively brings together an interdisciplinary community of researchers across computing, AI, robotics, and social science. However, progress has been limited by the lack of affordable, general-purpose, modular robot platforms with software tailored for HRI to enable computing researchers to enter the field, develop and test algorithms on real robots, and conduct statistically significant user studies. This symposium served to kick off the National Science Foundation–supported community-informed design and development of Quori, a low-cost robot hardware and software platform to enable computing research in HRI.

The symposium was split into both a

hardware and software discussion of Quori. These topics were presented and discussed by a diverse group of academic researchers, industry professionals, and hobbyists from fields including computer science, mechanical engineering, design, and psychology. Diversity was also represented in the keynotes: Andra Keay (director, Silicon Valley Robotics) discussed the commercialization potential of socially intelligent robots; Doug Dooley (freelance animator) presented a series of animation principles to help bring social robots to life; and Nate Koenig (CTO, Open Source Robotics Foundation) overviewed two popular open-source robotics software tools — Robot Operating System (ROS) and the Gazebo three-dimensional simulation environment — and the potential for social robots in the ROS ecosystem.

Hardware discussions at the symposium focused on the applications of Quori in HRI computing research and the subsequent hardware requirements to facilitate this community's specific needs and desires. The objective of the modular Quori hardware platform is to meet many of these diverse interests in a robust and inexpensive way that makes Quori accessible as a research platform. The discussions were primed with an opening on the project objectives, hardware, and appearance of Quori by Mark Yim (University of Pennsylvania), followed by a review of prior community input through surveys, as well as an overview of the software framework being developed for Quori. This was followed by paper presentations on topics including low cost-reliable hardware, the impact of robot appearance on human-robot interactions, and methods to study the acceptance of a social robot. Sample discussions include Martin Gerdzhev (McGill University) on a low-cost modular system for wheelchair robots, Maartje de Graaf (University of Twente) on long-term acceptance of social robots in domestic environments, and Bertram Malle (Brown University) on inevitable psychological mechanisms triggered by robot appearance. The hardware discussion closed with a breakout session where participants provided answers to how Quori should look, what features they would require to do research, and what modules are desired. This session had a significant and direct impact on Quori's iterative design. The team plans to dramatically change its approach to Quori's appearance to a robust and extensible hardware module system, possibly featuring a stacked cylindrical design, that better facilitates the requirements of researchers in the field. Thus the final hardware should be both ready to use out of the box and configurable for more unique research.

Software and autonomous robot social behaviors were the next topics to be covered by the symposium. The Quori team was primarily interested in tools required for HRI research, both existing and desired, and what level of complexity HRI researchers generally wish to manipulate platform behavior on. The discussion was primed with Ross Mead's presentation on Quori's software, behaviors, and the platform beta program, which will provide a Quori robot to eight different research institutions. As the day progressed, we discussed HRI software requirements to design powerful software packages without sacrificing the simplicity and reliability of the resultant application programming interface (API). Wizard-of-Oz (WoZ) studies, a commonly used design by HRI researchers, was merged with this topic discussion. Guy Hoffman (Sibley School of Mechanical and Aerospace Engineering, Cornell University) presented a generic WoZ system that enables researchers to easily perform WoZ studies, and provided new insights and possible directions for the Quori team moving forward. Finally, generation and design of autonomous behaviors for HRI studies was also presented, from Tiago Ribeiro (Intelligent Agents and Synthetic Characters, Instituto Superior Técnico Taguspark) and Rachel Holladay (The Robotics Institute, Carnegie Mellon University), and discussed.

The symposium concluded with future directions and two breakout sessions: research requirement analysis and robot prototyping. The results of these sessions are included above and final thoughts are provided here. Quori should provide a simple out-of-the-box experience for HRI researchers while still enabling highly modular hardware and software customization, allowing for diverse studies to be performed. The Quori platform has been well received by the community and participants look forward to working with Quori and continuing its iterative design process. The Quori team looks forward to overcoming the diverse set of design challenges this platform presents with the community. The organizers will continue to engage communities of interest through online surveys, symposia, and workshops to improve understanding of the needs and desires of potential Quori users.

Maja Mataric, Mark Yim, and Ross Mead served as coorganizers of the symposium. This report was written by Andrew Specian, Braden McDorman, and Ross Mead. The papers from the symposium were published by AAAI Press as technical report SS-16-03, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Ethical and Moral Considerations in Nonhuman Agents

The moral implications of our technological creations have long been a staple of science-fiction rumination. But now technology has reached a point where these moral concerns directly affect contemporary reality. Autonomously driven cars, drones that can be used to assassinate, social robots to assist the young, the sick, and the elderly, and sex robots, are examples of technology that bring ethical issues to the forefront. In this symposium, philosophers, roboticists, cognitive scientists, computer scientists, political scientists, and legal scholars came together for intense discussion of these issues, bringing exceedingly diverse viewpoints.

Broadly speaking, three concerns were discussed in the symposium. The first concern was the ethics of deploying autonomous systems in society. So, for instance, when a dialogue agent, a twitter bot, a military drone, or a socially assistive robot (whether for health care or for enjoyment) is deployed, what are the issues that need to be addressed by the humans who deploy them, and to what extent are the systems

autonomous and responsible. The second concern was incorporating ethical reasoning into the agents themselves, determining both the means (perhaps special-purpose languages) and the extent to which agents themselves should conduct moral or ethical reasoning, and how human values should be communicated to such intelligence. This topic includes programming paradigms for implementing moral reasoning. The final concern focused on studying how moral values like guilt, forgiveness, and cooperation emerge, and what utility they have in society. This research is conducted by running simulations such as agent-based models or multiagent systems, or by using evolutionary game theory directly.

In the discussion, we identified three paradigms of artificial intelligence ethics: (1) AI agents can be conceived directly as moral subjects (agent, patient, or both); (2) AI can be applied as a cognitive enhancement for humans directly to improve our ethical decision making; and (3) AI agents can be conceived of strictly as tools or prosthetic intelligence, serving human goals and maintaining human responsibility, but not incorporated into human bodies. This can include assistive technology that helps a human see a situation from different ethical perspectives.

Notice that these three paradigms are not necessarily mutually exclusive, nor is there anything essential or technological to determine which approach we take. Rather, the question is which approach is to be recommended — which is the most moral with respect to our obligations to ourselves, to our society, to our culture, and (possibly) to our artifacts. Some papers and participants argued passionately for or against passing moral responsibility to AI, including whether true or full intelligence can be achieved any other way. Some argued that their twitterbots were already morally responsible for their own actions, others that the idea of passing moral responsibility to artifacts is incoherent, for example, that any capacity related to suffering is meaningless since it can be intentionally built or excluded, or even built then excised, which cannot be done to humans or other animals.

Concerns were raised as to how AI is affecting the moral landscape. For example, twitterbots can create fake followers and can generate new ideas (some of which may be offensive and hurtful). Drones afford more selective assassinations, which might reduce collateral civilian casualties compared to a conventional bombing attack, but create oppressive fear and uncertainty. Social robots can display deceptive behavior to help or please their user. Here again we discussed technology that is already on the market and in our society, and which we expect to become more pervasive.

As previously mentioned, we debated whether autonomous systems should be considered moral agents. Around the time of the symposium there was an accident in Palo Alto when Google's autonomous car bumped into a public bus. Google reiterated its position in a related statement, that the company (rather than any owner or the cars themselves) will be liable for all damage caused by its autonomous cars. The Swedish car company Volvo took a similar stance some months ago, taking all responsibility for accidents that may happen while its cars are in the autonomous mode. For the time being, this pragmatic approach may be the best method to approach the issue of liability. Though note too that there was some discussion that corporations should also be thought of as artificial intelligences, whose goals may not be aligned with those of individual humans' or of humanity at large.

We felt a pressing need to involve more people with background in jurisprudence and law in the discussion. We intend to rectify this in a follow-up conference. More generally, we were concerned how to disseminate what we had learned in three days of discussion out to our colleagues and into the public as soon as possible. AI is growing at an accelerating rate, and is already creating ethical and moral dilemmas in the real world, affecting humans, corporations, and governments. It is essential that we academics engage with policy makers, legislators, and lawyers as well as with industry.

Bipin Indurkhya and Georgi Stojanov served as cochairs of this symposium. This report was written by Bipin Indurkhya, Georgi Stojanov, and Joanna Bryson. The papers from the symposium were published by AAAI Press as technical report SS-16-04, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Intelligent Systems for Supporting Distributed Human Teamwork

Distributed teamwork has become more common as technology enables groups of people distributed over vast distances, with fewer opportunities for synchronous interaction to work together on complex tasks extended in time. Existing technologies, however, rarely provide intelligent support to help teams work together more effectively. Many teamwork challenges, such as coordination, trust, and communication, persist despite developments in technology.

The symposium brought together AI and human-computer interaction researchers to identify challenges to developing intelligent systems for supporting human teamwork along with potential multidisciplinary approaches to overcoming them. Cross-disciplinary expertise is essential for pushing forward the boundaries of systems for supporting distributed human teamwork. For example, systems might benefit from intelligent algorithms that reduce coordination overhead, but assumptions AI methods make for computer-agent environments often poorly match people's capabilities. Integrating key ideas from social science and HCI research into the design of AI methods will enable the development of systems that address people's core needs, adequately consider cultural and organizational factors, and make reasonable assumptions.

Position papers submitted to the symposium tackled a wide range of teamwork problems in a variety of application domains. Several papers addressed teamwork problems in the context of health care, such as forming teams to combat HIV spread among homeless youth, supporting interdisciplinary medical team meetings, and supporting information sharing in complex care teams. Another common application domain was education, where papers addressed problems such

as the development of conversational agents to facilitate distributed learning and methods for crowdsourcing learners' activities in a way that would provide benefit to other learners.

During the symposium, participants engaged in several working group sessions, aimed at identifying key teamwork problems and opportunities to integrate insights and methods from AI and HCI to address these problems. To provide common ground for discussion, an invited talk by Gary Olson reviewed research on teamwork from the human-computer interaction and computer-supported cooperative work communities, and one by Liz Sonenberg and Milind Tambe reviewed research on teamwork from the multiagent systems community.

In the working group sessions, participants discussed a variety of teamwork problems, including team formation, coaching of teams, team coordination, and reconciliation of different perspectives of team members. Several overarching challenges to the development of effective intelligent systems emerged in these discussions. One such challenge is the need to evaluate such systems not simply in terms of accuracy or efficiency, but also in terms of their effects on people's activities. A related challenge is the importance of ensuring that intelligent systems empower people, rather than diminish people's sense of control or abilities. To make this the case will require that we develop transparent systems that can clearly communicate and explain their reasoning to the people they aim to support.

Ofra Amir, Krzysztof Gajos, Barbara Grosz, Gary Olson, and Judy Olson served as cochairs of this symposium. This report was written by Ofra Amir and Barbara Grosz. The papers from the symposium were published by AAAI Press as technical report SS-16-05, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Observational Studies through Social Media and Other Human-Generated Content

While using the Internet and mobile devices, people intentionally and unintentionally create data through their interaction with messaging services, websites, and other applications. Such human-generated content in general, and social media in particular, are a rich repository of data for observational studies across many areas: public health, with research on prevalence of disease and on the effects of media on the development of disease; medicine, showing the ability to detect mental disease in individuals using social media; education, to optimize teaching and exams; and sociology, to prove theories previously tested on very small populations. The result is that studies can be performed in a variety of topics with heretofore unprecedented large populations.

While many past studies demonstrate correlations between variables of interest, some study designs are able to show causal relationships through natural experiments and mathematical inference methods. This AAAI symposium investigates such causal inference from human-generated content.

The symposium brought together researchers from a wide array of fields such as economics, computer science, sociology, and public health. Applications of observational studies across these domains was one of the major themes of the symposium. Susan Athey (Stanford University) discussed how experiments with web search logs and social media helped provide insight into the power of information intermediaries such as search engines and social networks. Jure Leskovec (Stanford University) presented studies of antisocial behavior in online forums; and Amit Sharma (Microsoft Research) showed techniques to better estimate the true influence of online recommender systems. Other presentations covered observational studies in public health scenarios, political science, crisis informatics, linguistics, and economics.

A primary focus of presenters was the many challenges inherent in causal inference from large-scale, observational data sets. Dean Eckles (MIT) presented approaches to modeling and reducing confounding bias when dealing with high-dimensional observational data. Virgile Landeiro and Aron Culotta (IIT) discussed how commonly used text-classification techniques can inadvertently bias observational studies. Isaac Johnson (Univsity of Minnesota) and Han Zhang (Princeton) presented challenges, biases, and opportunities in using geo-referenced human-generated content. Adam Glynn (Emory University) discussed a thought-provoking application of Judea Pearl's front-door criterion, combined with prior knowledge from experience and social theory, for causal inference without control units. Emre Kicman (Microsoft Research) presented on an effort to build a more general-purpose analytical system that could be reused for a large class of analyses from social media. Adding to purely technical challenges of experimental design and analytical techniques, Brian Keegan (Harvard) proposed a framework for reasoning about ethics and the potential harms in studies of interactions among multiple parties.

The cochairs would like to thank the AAAI Spring Symposium chairs and organizers, the OSSM program committee, and especially the presenters and attendees at the symposium for the engaging, thought-provoking, and lively discussions. Munmun De Choudhury (Georgia Tech), Emre Kiciman, and Elad Yom-Tov (Microsoft Research) served as cochairs of this symposium.

This report was written by Emre Kiciman. The papers from the symposium were published by AAAI Press as technical report SS-16-06, which was included in The 2016 AAAI Spring Symposium Series: Technical Reports.

Well-Being Computing: AI Meets Health and Happiness Science

Well-being computing is an information technology that aims to promote psychological well-being (that is, happiness) and maximize human potential. Our environment escalates stress, provides unlimited caffeine, distributes nutrition-free fast food, and encourages unhealthy sleep behavior. For this issue, well-being computing 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. Today, great advances are being made both in the science of health and well-being and artificial intelligence (AI). Synergy between these two fields can bear fruits in well-being computing. It is now very important to share these scientific findings and AI methodologies for better human-centered system design. Well-being computing is where AI meets health/happiness sciences.

In this symposium, we explored the methods or methodologies for four topics: (1) methods for quantifying our health happiness and well-being; (2) methods for analyzing the health and wellness data to discover the new meanings (discovery informatics technologies and cognitive and biomedical modeling); (3) methods for designing better health and well-being space; and (4) applications, platforms, and filed studies.

Five invited talks gave participants a new perspective on well-being computing. Michael Snyder (Stanford University) spoke on genomics and personalized medicine by introducing his iPOP (integrative personal omics profile) project. He also introduced his AI/deep-learning health informatics analyses from wearable device data. Michael Nova (Pathway Genomics Inc.,) discussed the issues on mobile cognitive health care using AI (such as IBM Watson), which can automatically read unstructured data and then dynamically learn to make personalized recommendations on personal health (for example, genetic tests, lab data, or wearable information). Monica Worline (Stanford University) introduced the research topics on the science of compassion in Stanford CCARE (the Center for Compassion and Altruism Research and Education), and then focused on her social scientific research on how we can create compassionate cultures in the workplace. Rafael Calvo (the University of Sydney) explained his concept of positive computing, which refers to the area of work on the design and development of technology to support psychological well-being and human potential, and then he discussed some related projects on AI innovation for psychological well-being. (for example, new systems that leverage natural language processing and AI to detect depression and anxiety from social media data.) Finally, Dennis P. Wall (Stanford University) spoke about big data analysis for autism research. He introduced his Google Glass projects aimed at early-stage autism diagnosis with wearable devices. He also explained how machine-learning technologies (for example, deep learning), computational tools of systems biology, and a crowdsourcing platform (human in the loop approach) can be applied in the big data analyses on genomic and phenotypic Autisms databases.

Twenty-on technical papers and three posters or demonstrations were presented over the course of the two-and-a-half days. Presentation topics included well-being computing framework, interaction, personal genome, visualization, health monitoring, mindful technologies, sleep, dementia support, positive computing and self-control, health data analysis, emotion analysis, and brain science and brain computer interface.

Takashi Kido introduced frameworks for connecting machine learning and personal genome informatics with happiness sciences, Mai Kosahara proposed a system to visualize tactile perceptual space of young and old people, Hannes Bend discussed the scientific and artistic framework for mindful technologies, Keiki Takadama proposed a next-generation sleep-monitoring system based on human body vibration analysis, Kenichi Shibata and Yoichi Takebayashi proposed an interprofessional collaborative system to raise awareness and understanding of dementia using an action observation method, Er Sin Khoo evaluated the mental time of older adults during conversations supported by coimagination method with expedition, Rafal Rzepka discussed the concept of global brain, which makes us think twice with crowdsourcing platforms, and Ray Lee reported experimental results on neural correlates of conscious flow during meditation.

The symposium provided participants with unique opportunities where researchers from completely different backgrounds came up with new ideas through innovative and constructive discussions. We believe the symposium will present important interdisciplinary challenges that may guide future advances in the AI community.

Takashi Kido and Keiki Takadama served as cochairs of this symposium and wrote this report. The papers from the symposium were published by AAAI Press as technical report SS-16-02, which was included in *The 2016 AAAI Spring Symposium Series: Technical Reports*.

Christopher Amato is an assistant professor at the University of New Hampshire.

Ofra Amir is a Ph.D. candidate at Harvard University.

Joanna Bryson is a reader at the University of Bath .

Barbara Grosz is Higgins Professor of Natural Sciences in the School of Engineering and Applied Sciences at Harvard University.

Bipin Indurkhya is a professor at Jagiellonian University, Poland.

Emre Kiciman is a senior researcher at Microsoft Research.

Takashi Kido is a research manager at Rikengenesis in Japan.

W. F. Lawless is a professor at Paine College.

Miao Liu is a postdoctoral associate at the Massachusetts Institute of Technology.

Braden McDorman is a Ph.D .student at the University of Southern California.

Ross Mead is the founder and chief executive officer of Semio.

Frans A. Oliehoek is a Veni researcher at the University of Amsterdam.

Andrew Specian is a Ph.D. student at the University of Pennsylvania.

Georgi Stojanov is a professor at the American University in Paris.

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