Reports of the AAAI 2010 Spring Symposia

Thomas Barkowsky, Sven Bertel, Frank Broz, Vinay K. Chaudhri, Nathan Eagle, Michael Genesereth, Harry Halpin, Emily Hamner, Gabe Hoffmann, Christoph Hölscher, Eric Horvitz, Tom Lauwers, Deborah L. McGuinness, Marek Michalowski, Emily Mower, Thomas F. Shipley, Kristen Stubbs, Roland Vogl, Mary-Anne Williams

■ *The Association for the Advancement of Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2010 Spring Symposium Series Monday through Wednesday, March 22–24, 2010, at Stanford University. The titles of the seven symposia were Artificial Intelligence for Development; Cognitive Shape Processing; Educational Robotics and Beyond: Design and Evaluation; Embedded Reasoning: Intelligence in Embedded Systems; Intelligent Information Privacy Management; It's All in the Timing: Representing and Reasoning about Time in Interactive Behavior; and Linked Data Meets Artificial Intelligence.*

Artificial Intelligence for Development

The Symposium on Artificial Intelligence for Development was organized to explore opportunities for using machine learning, inference, planning, and perception to enhance the quality of lives of disadvantaged populations. Over the last several years, a community of researchers with interest in applying computing and communication technologies in developing regions has come together under the label Information and Communication Technology for Development (ICT-D). However, ICT-D efforts to date have rarely focused on opportunities to harness machine learning, reasoning, and perception to create intelligent systems, services, models, and analyses. Beyond exploring research projects and directions, we hoped that bringing together a critical mass of researchers who share an interest in applying AI to development challenges would serve to help launch a new vibrant subfield of ICT-D on artificial intelligence for development (AI-D).

The symposium program[1] included a lively mix of invited keynotes and presentations, with sessions on insights about people and behavior, health, education, and welfare, information access, and infrastructure and agriculture. Panel discussions explored the terrain of the field in a panel on Grand Challenges in AI for Development and causality and intervention in a panel on Learning, Causation, and Effective Action.

The talks, panels, and open discussions highlighted numerous opportunities to apply methods developed within the AI community to enhance the quality of lives and promote the socioeconomic development of people in the poorer regions of the world. Directions of research discussed at the meeting include the use of machine learning and reasoning to extend medical care to remote regions through automated diagnosis and effective triaging of limited medical expertise and transportation resources. Machine intelligence may one day assist with detecting, monitoring, and responding to natural, epidemiological, or political disruptions. Methods developed within the artificial intelligence community may even help to unearth causal influences within large-scale programs, so we can better understand how to design more effective health and education systems. And ideas and tools created at the intersection of artificial intelligence and electronic commerce promise to provide new directions for enhancing and extending novel economic concepts like microfinance and microwork.

Discussions included reflection about the particular promise for harnessing machine learning to help enhance the quality of life of people living in developing regions. Unprecedented quantities of data are being generated in the developing world on human health, commerce, communications, and migration. Automated learning methods developed within the AI community can help to tease out insights from this data on the nature and dynamics of social relationships, financial connections and transactions, patterns of human mobility, the dissemination of disease, and such urgent challenges as the needs of populations in the face of crises. Models and systems that leverage such data might one day guide public policy, shape the construction of responses to crises, and help to formulate effective long-term interventions.

At the end of the symposium, discussions focused on next steps. Attendees expressed a desire to arrange future meetings within and outside of the ICT-D conference venues. Discussion also focused on developing solutions to challenges with pooling and sharing data that is properly anonymized, and for sharing experiences, tools, and components with people interested in AI-D challenges. To help to promote the establishment of a vibrant AI-D research community, the cochairs set up a nonprofit organization (AI-D) and associated website for collecting and distributing datasets and information on people, projects, and data access. We invite readers to explore the content and resources at the AI-D site.[2]

We hope that the organization of the symposium stimulates additional efforts on opportunities to harness machine learning, reasoning, and perception to enhance the quality of life within disadvantaged populations. We thank the authors, attendees, and program committee for their creativity, effort, and energy in organizing this meeting, and we invite people who could not attend the symposium to explore the interesting and tantalizing challenges in AI-D.

Eric Horvitz and Nathan Eagle served as cochairs of this symposium. The papers were published as AAAI Technical Report SS-10-01.

Cognitive Shape Processing

The AAAI 2010 Spring Symposium on Cognitive Shape Processing was held at Stanford University, California, March 22–24, 2010. The goal of the symposium was to promote an understanding of how shape information is cognitively represented, retrieved, (re-)constructed, and integrated with other types of spatial information, and to gauge how this can inform AI approaches that process spatial information involving shape.

There has been a growing interest in recent decades in understanding and computationally modeling how spatial information is processed in natural intelligent systems. One AI-related motivation for this is the hope that one can improve the performance of artificial systems, such as in robotics, or of intelligent instruction or other interactive systems, by imitating representational and procedural principles found in natural intelligent systems. In spatial cognition research, numerous aspects of spatial knowledge are investigated, including knowledge about spatial reference systems and topology, route knowledge, knowledge about distances and directions, and so on. These investigations have led to the development of various representation and reasoning formalisms, many of which are specific to only one or very few aspects, or which deal with highly simplified spatial objects, such as points or basic geometric forms. Real-world situations, in contrast, typically deal with diverse types of spatial knowledge at the same time, and they involve complex objects with meaningful and specific shapes.

By *cognitive shape processing* we refer to all forms of knowledge processing that involve shape information and that are related to, inspired by, or derived from principles found in natural cognitive systems.

This symposium brought together a group of more than 20 highly motivated researchers from various fields in artificial intelligence, the cognitive sciences, mathematics, and design research. Contributions covered among other topics a range from formal descriptions of shape information (for example, as sets of qualitative spatial relations, as based on ordering information, or through combined spectral and spatial descriptions) to behavioral studies into the perceptual and cognitive foundations of shape (for example, as comparative explorations into modal aspects of shape, such as visual and haptic modalities, or as explorations into the role of shape information during the mental rotation of objects), to studies of shape aspects in very applied contexts (for example, for surgery or anatomy learning), to issues of cognitive modeling (for example, on the relationship between reasoning with spatial and shape information, or on "basic level" representations for shape knowledge), to investigations into the role of shape information for engineering contexts (for example, with respect to computer-aided design).

The symposium featured two invited keynote lectures. One lecture, given by Stephen Grossberg of Boston University, focused on cortical organization principles that are responsible for seeing and learning to recognize object categories and three-dimensional shapes, and on how the brain manages to direct eye movements and spatial attention during object learning and recognition. Philip Kellman of UCLA gave the second keynote lecture; he addressed issues of segmentation, grouping, and shape during human visual perception of contours and objects and argued that shape perception is deeply connected with object detection.

The symposium participants soon realized that while shape is a fundamental concept in many disciplines, the diversity of the presented approaches required a careful discussion of the different terminological and ontological bases of "shape" and of some related concepts. Topics that were also discussed in depth either in plenary sessions or in small groups included the specific importance of shape information, whether shape knowledge can be equally acquired and retrieved through different sensory modalities, the role of perceptual and cognitive simulations in shape processing, and how different descriptions of shape-related representations and processes on neural, psychological, and technical description levels can productively inform each other.

At the conclusion of the symposium, the participants discussed how the high degree of diversity in the audience had been unexpected but had turned out to be useful in gaining new perspectives and knowledge on dealing with shape information. It was found that the presented approaches from AI and cognitive psychology were often driven by similar problems. Last, the participants shared the impression that, in retrospect, "Cognitive Shape Representations" would have been a better title for this symposium as most contributions in fact emphasized representational aspects over procedural ones. "Cognitive Shape Processing" is thus proposed as the title for the symposium that should follow this one.

Sven Bertel served as chair of this symposium, assisted by Thmas Barkowsky, Christoph Hölscher, and Thmas F. Shipley. Madeleine Keehner of the University of Dundee delivered an inspiring report of this symposium at the Spring Symposia plenary session. The papers of the symposium were published as AAAI Press Technical Report SS-10-02. Additional support came from the Collaborative Transregional Research Center SFB/TR 8 Spatial Cognition. Funding by the German Research Foundation (DFG) is gratefully acknowledged.

Educational Robotics and Beyond: Design and Evaluation

The goals of the Educational Robotics and Beyond: Design and Evaluation Symposium were to present new, educationally focused tools and to discuss methods for evaluating educational experiences that use such tools.

For more than 20 years, robots, electronics, microcontrollers, and other physically instantiated devices have been used as educational tools in both formal (in-school) and informal (out-of-school) settings. Following a nearly annual tradition, this symposium provided a gathering place for researchers working on the design of novel systems for education, as well as for those interested in evaluating the outcomes of educational interventions using these tools. As such, the main themes of discussion at the symposium were educational evaluation, curriculum design, and tool design.

Evaluating the effectiveness of new educational tools on student performance is a difficult process. Furthermore, frequently those who design new tools or curricula are not formally trained in methods for educational evaluation. As such, we devoted a large fraction of the symposium to discussing educational evaluation. We kicked off this portion of the symposium with an invited speaker from the University of Pittsburgh's Learning Research and Development Center (LRDC) who has personally worked closely evaluating a number of robotics programs. We then had a combined panel and breakout session in which symposium participants were able to build off one another's experiences with evaluation to create a wide-ranging set of suggestions for how to better evaluate educational tools.

We focused on curriculum design through a number of our author presentations, specifically by holding another combined panel and breakout session. This session was devoted to discussing ideas for avoiding situations in which students are excited by newly energized and reformed introductory courses, only to be discouraged by follow-on courses that don't continue the themes that motivated students in the introductory course. While we looked in detail at how this problem can play out in the computer science curriculum, we recognized that it is a problem affecting much of science, technology, engineering, and math (STEM) education.

Finally, no educational robotics (and beyond!) symposium would be complete without offering participants a chance to use one another's cool new tools. To this end, we devoted Wednesday morning to a hands-on workshop where six participants brought their own educationally focused systems. These systems were set up around the room and participants were invited to use them.

The systems presented were widely diverse in terms of the types of educational experiences targeted, including introduction to computer science, kinematics, citizen science, secondary school engineering, and arts education. Participating in the Wednesday morning workshop, the authors were struck by the potential of these electronic tangibles to improve education in a large number of different fields. We suggest that the coming challenge in this field will be how to create a virtuous, iterative cycle in which evaluation, curriculum design, and tool design are aligned to create ever more effective and enjoyable educational experiences.

Tom Lauwers, Kristen Stubbs, and Emily Hamner served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical report SS-10-03.

Embedded Reasoning:Intelligence in Embedded Systems

Embedded reasoning incorporates the strengths of AI reasoning — planning, scheduling, controlling, learning, and diagnosing — into physical systems. This symposium combined these diverse fields to work toward common understandings and common definitions.

To accomplish that goal, this symposium brought together researchers in artificial intelligence, control systems, robotics, and human-machine interaction to explore the capabilities of each field and identify how they interact to form systems containing embedded reasoning. Recent advances in the field have been driven by improvements in processor capability and the desire to control increasingly complex systems safely, efficiently, and reliably. Embedded reasoning incorporates the strengths of AI reasoning — planning, scheduling, controlling, learning, and diagnosing — into physical systems. This advances system capabilities in solving complex tasks, in acting on high-level goals, and in adapting to changing and uncertain states. The integrated methods by which we approach these problems are rapidly evolving. As illustrated by the projects of the symposium attendees, these emerging capabilities require a tight integration of diverse techniques with a strong multidisciplinary understanding of their relationship. The traditional interfaces of the fields of AI reasoning, control, and human factors are becoming blurred, with control-system optimizations running on embedded processors, and artificial intelligence controlling autonomous vehicles.

We discussed challenges, approaches, and solutions for enabling systems of sensors, actuators, and processors to be adaptive, distributed, and robust. Many presenters discussed techniques for autonomously interacting with and understanding the surrounding environment with noisy sensor inputs and imperfect models of system behavior. A recurring theme was in the hierarchical nature of the representations of such systems. There was a trade-off in approaches. Although compartmentalization was used to simplify many problems, by breaking down these artificial barriers, the constraints on the solution space are relaxed enabling superior performance. We also discussed real-time execution and process concurrency, looking at the frameworks being developed in research. These frameworks increasingly make it possible to develop systems capable of meeting spatial and temporal requirements. There is a clear trade-off between ease of use and efficiency of multithreaded systems. Several speakers discussed architectures that compartmentalized real-time execution of lower-level tasks to enable more sophisticated higher-level tasks. Discussion of human-robot interaction demonstrated the importance of establishing a common ground between the user and the system.

Panel discussions led to engaging debates. To establish common language in this diverse group, the question naturally came up: what is an embedded system? Some said that an embedded system must be a part of something else — it's an algorithm incorporated into a larger system. Others added that an embedded system is composed of algorithmic processes that share interaction with physical processes. Hence, time and physical world effects such as constraints and uncertainty must be accounted for. Estimation and extrapolation of world state was identified time and again as a critical component of embedded reasoning, either explicitly, or implicitly by example of system limitations. This led to lively debate on the trade-off between two contrasting approaches to estimation — model based versus data driven. Although everybody agreed that some cases clearly call for one approach over the other, with ever increasing connectivity of devices, "crowd-sourcing" is shifting the balance toward more reliance on data than before.

The nature of knowledge representation across these disciplines was raised as a major need to enable efficient development of embedded systems. It is a potential waste of specialized skills to require one to become an expert in multiple domains to use the tools across those domains. For example, in the planning domain, practitioners have developed PDDL to define problems that their algorithms can solve. In the control-systems domain, practitioners have developed equations of motion with algorithms to create control laws. However, classical control's guarantees are only probabilistic, a challenge for a planner, and a control-system optimizer would be challenged to determine how to most efficiently achieve the intent of the planner. This was left as an open question and rich area for exploration — how can knowledge be exchanged across interfaces to enable more integrated approaches, beyond computationally expensive methods?

Gabe Hoffmann served as the chair of the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-10-04.

Intelligent Information Privacy Management

The goal of the Intelligent Information Privacy Management Symposium was to develop a new transdisciplinary understanding of privacy and personal information management issues by drawing from the key areas of law, artificial intelligence, and business.

This AAAI symposium attracted AI researchers, legal scholars, computer scientists, and people from business. Participants came from North America, Europe, and Asia. Forty papers were presented covering philosophical, legal, and technological aspects of forensic DNA profiles, smart environments, digital books, privacy in the cloud and mobile devices, sensor networks, autonomous systems, vehicles, social networks, the right to delete, and geospatial information. In addition to the technical papers the symposium provided a forum for a keynote address on privacy by design by Ontario's Information and Privacy Commissioner Ann Cavourkian.

A key objective of the symposium was to develop a new transdisciplinary understanding of privacy and personal information management issues by drawing from the key areas of law, artificial intelligence, and business. It focused on the need to develop effective information privacy management frameworks, tools, and techniques by addressing the underlying tension between transparency and disclosure. The motivation for organizing the symposium was the recognition that there is a significant and growing need to identify privacy requirements in application development and to use intelligent technology-enabled solutions to assist users to monitor and manage their personal information in a more transparent proactive fashion.

People derive significant benefits from sharing their personal details as they take advantage of relevant and useful services, particularly online. However, once the personal information is collected, businesses often seek to exploit and monetize it, and oftentimes it is disclosed. Individuals possess a digital footprint that computer applications can discover and use for a wide range of purposes including behavioral targeting for advertising products and services. Protecting and enforcing privacy is a major cost to business, but a lack of privacy protection creates risk for users and reduces trust. Trust plays an important role in the generation of innovation; without trust consumers tend to avoid engagement, they minimize or falsify responses, and as a result business opportunities can be missed and innovation retarded. The free flow of personal information that respects privacy can fuel and cultivate innovation.

Three lively panels were organized around the main symposium themes: what is privacy, privacy law, and the relationship between privacy and business innovation. Invited panelists included thought leaders from Stanford and the University of California, Berkeley, leading practitioners from law firms, and companies such as Facebook, LinkedIn, Adobe, and Google. The panelists were given the opportunity to identify research problems and field questions from the participants.

Many of the key ideas discussed during the symposium converged over the course of the symposium — for example, the notion that privacy is context dependent and that technologies need to be context aware in their management of privacy. Some of the challenging issues that arose were the need to offer users choice and control of their personal information, the difference between consent and informed consent, the need for systems to help users anticipate the impact of sharing information on their privacy, and the need to develop a better understanding of changing social norms with respect to privacy and information sharing.

Michael Genesereth, Roland Vogl, and Mary-Anne Williams served as cochairs of the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-10-05.

It's All in the Timing: Representing and Reasoning about Time in Interactive Behavior

When discussing representations of the physical world for autonomous agents, the passage of time is often overlooked or treated as an afterthought. However, especially in interactive domains, the timing of an action can be critically important to the action's intended meaning. This issue is important both to researchers who attempt to understand human interactive behavior and to those who design computational agents that interact autonomously with people. Human social behavior is highly dependent on a close feedback loop of simultaneous and coordinated activity between multiple interactors. Yet how to best represent these interdependencies and temporal relationships is just beginning to be explored and understood from a computational perspective. Speed, acceleration, tempo, and delay are concepts that AI and robotics researchers recognize as important in everything from motor control to verbal communication, but we do not yet possess a well-motivated framework for incorporating these temporal considerations into our designed systems.

This symposium brought together and discussed a range of approaches to analyzing and appropriately synthesizing temporal properties of interactive behavior. Participants came from a wide range of fields, including robotics, cognitive science, and theater. In addition to scheduled talks, attendees participated in focused breakout sessions discussing the role of time in the performing arts, planning, cognitive science, and language learning. Practitioners drew on a wide variety of tools for representing time in their systems, including graphical models, dynamical systems, and neural networks. There was also discussion of the conditions under which time must be represented explicitly, and when the timing governing an interaction may be implicitly represented through dynamics arising from the embodiment of the system. Participants working in cognitive science presented models based on the way the human mind processes time-related information.

Multiple participants suggested that human actors, with their refined understanding of the role of timing in nonverbal communication, may have much to teach the designers of socially interactive robots or software agents about producing expressive behavior. One of the presenters, Annamaria Pileggi (Washington University in St. Louis), led the participants in an elementary nonverbal acting exercise, with full-body explorations of starting or stopping behavior, changing direction, spatial relationships, leading or following, and tempo. This exercise provided hands-on insight into how actors learn to manipulate the timing of movement as an expressive tool, and breakout groups later studied recorded video and accelerometer data from the activity as an exercise in analyzing the timing of human behavior.

One common theme was the difficulty in generalizing certain concepts or techniques across the range of time scales (from microseconds to days) involved in social interaction. Another theme was the difference between analysis and synthesis of temporally appropriate behavior and the different considerations when these needed to be achieved online or offline. The symposium brought together people who use time in very different research endeavors using very different techniques, and it highlighted the goals shared across these groups. Participants agreed that increased attention and an explicit focus on timing is important in future symposia, with an even greater focus on psychological and cognitive models.

Frank Broz, Marek Michalowski, and Emily Mower served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-10-06.

Linked Data Meets Artificial Intelligence

The goal of linked data is to enable people to share structured data on the web as easily as they can share documents today. The basic assumption behind linked data is that the value and usefulness of data increases as it is interlinked with data from other sources using typed links. This emerging web of data includes data sets as extensive and diverse as DBpedia, Geonames, US Census, EuroStat, MusicBrainz, BBC Programmes, Flickr, DBLP, PubMed, UniProt, FOAF, SIOC, OpenCyc, UMBEL, Virtual Observatories, and Yago. The symposium was aimed at bringing together the researchers working on linked data and AI to create a new community interested in utilizing AI techniques such as ontologies, machine learning, data fusion, and visual analytics in exploring the linked open data.

The first day started with provocative big picture statements from researchers looking for future trends. An area of early impact of linked open data is its use in expanding the concept of open government by publishing government data using linked open standards. Demonstrations built using this newly available public data set show the increased level of transparency that is possible in the functioning of the government. One presentation from Li Ding and Deborah McGuinness provided numerous examples using publicly available government data in implemented demonstrations displaying the value of linking open data while minimizing reliance on background representations. Richard Fikes (Stanford) theorized how open government data could be combined with richer representations that are being investigated in AI-motivated projects such as Project Halo to help question answering about governments, thus giving a case study on how linked data can be combined with AI techniques.

The current state of the art in browsing linked data is rather primitive and limited to using techniques such as faceted browsing. The presentation from Eric Bier (PARC) on the visualization inspired many younger researchers at the symposium by opening up a whole new range of techniques that could be applied to linked data.

Another major theme was the exploration of how the ever-increasing amounts of public structured linked data transformed both our understanding of data-driven computing and artificial intelligence. Peter Norvig (Google) gave an invited talk on the "Unreasonable Effectiveness of Data," demonstrating how problems that were formerly intractable using standard artificial intelligence techniques could be solved by using statistics over massive amounts of data. Othar Hansson and Kavi Goel (Google) showed how Google can use structured RDF data in its "Rich Snippets" search engine feature. In the lively panel discussion afterward, two distinct insights emerged. The first insight was that the chasm between "unstructured" statistical machine learning and highly structured knowledge representation can be bridged by the use of "lightweight" structure in linked data such as input vectors to machine learning. The second insight was that it seemed that the use of AI with linked data was not just to understand human intelligence on a level of abstraction but to "amplify intelligence" by creating hybrid human and computational collective intelligence on the backbone of the web. Pat Hayes (Institute for Human and Machine Cognition) brought this point home by reminding us that "We're intelligent, aren't we?" in his summarizing commentary.

Far from presenting only the theoretical, the symposium participants also discussed in detail the practical details of releasing large data sets, as demonstrated by the invited talk of Robert Kaye (MusicBrainz) on how he helped release the world's music knowledge. Another panel had businesses, ranging from Freebase to Microsoft, discussing the future of research on AI and linked data, demonstrating that many of the techniques developed in AI using linked data could have commercial significance. The symposium brought together participants from a wide range of areas and helped to create a new emerging community, making a historic link between artificial intelligence and linked data.

Vinay Chaudhri (SRI International), Harry Halpin (University of Edinburgh/W3C), Deborah L. McGuinness (Rensselaer Polytechnic Institute), and Dan Brickley (Vrije Universiteit Amsterdam) served as cochairs of the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-10-07.[3]

Notes

1. See ai-d.org/program2010.html.
2. AI-D.org.
3. The workshop site is available at www.foaf-project.org/events/linkedai.

Thomas Barkowsky is an assistant professor at the Cognitive Systems Group at the University of Bremen and scientific manager of the Collaborative Transregional Research Center SFB/TR 8 Spatial Cognition at the Universities of Bremen and Freiburg.

Sven Bertel is a postdoctoral research associate at the Human Factors Division and the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign.

Frank Broz is a research fellow in the Adaptive Systems Research Group in the Computer Science Department at the University of Hertfordshire.

Vinay K. Chaudhri is the program director of ontology management in the Artificial Intelligence Center at SRI International.

Nathan Eagle is the CEO of txteagle Inc. and a visiting assistant professor at the MIT Media Laboratory.

Michael Genesereth is an associate professor in the Computer Science Department at Stanford University.

Harry Halpin is a postdoctoral research associate at the University of Edinburgh and W3C Fellow.

Emily Hamner is a senior research associate at the Robotics Institute, Carnegie Mellon University.

Gabe Hoffmann is a researcher at the Palo Alto Research Center.

Christoph Hölscher is a lecturer at the Center for Cognitive Science at the University of Freiburg.

Eric Horvitz is a distinguished scientist at Microsoft Research and immediate past president of AAAI.

Tom Lauwers is a postdoctoral researcher at the CREATE lab in the Carnegie Mellon Robotics Institute.

Deborah L. McGuinness is the Tetherless World Senior Constellation chair and professor of computer science and cognitive science at Rensselaer Polytechnic Institute.

Marek Michalowski is a graduate of the Robotics Institute at Carnegie Mellon University and cofounder of BeatBots LLC.

Emily Mower is a Ph.D. candidate at the University of Southern California.

Thomas F. Shipley is an associate professor in the Department of Psychology at Temple University and the Spatial Intelligence and Learning Center.

Kristen Stubbs is the STEM outreach program manager at iRobot Corporation.

Roland Vogl is executive director of the Stanford Program in Law, Science, and Technology, and a lecturer in law at the Stanford University Law School.

Mary-Anne Williams is a professor and director of the Innovation and Enterprise Research Laboratory at the University of Technology, Sydney, Australia.