

Reports of the AAAI 2011 Spring Symposia

Mark Buller, Paul Cuddihy, Ernest Davis, Patrick Doherty, Finale Doshi-Velez, Esra Erdem, Douglas Fisher, Nancy Green, Knut Hinkelmann, James McLurkin, Mary Lou Maher, Rajiv Maheswaran, Sara Rubinelli, Nathan Schurr, Donia Scott, Dylan Shell, Pedro Szekely, Barbara Thönssen, and Arnold B. Urken

■ *The Association for the Advancement of Artificial Intelligence presented the 2011 Spring Symposium Series Monday through Wednesday, March 21–23, 2011, at Stanford University. This report summarizes the eight symposia.*

The Association for the Advancement of Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2011 Spring Symposium Series Monday through Wednesday, March 21–23, 2011, at Stanford University. The titles of the eight symposia were Artificial Intelligence and Health Communication, Artificial Intelligence and Sustainable Design, Artificial Intelligence for Business Agility, Computational Physiology, Help Me Help You: Bridging the Gaps in Human-Agent Collaboration, Logical Formalizations of Commonsense Reasoning, Multirobot Systems and Physical Data Structures, and Modeling Complex Adaptive Systems As If They Were Voting Processes. This report summarizes the eight symposia.

Artificial Intelligence and Health Communication

The goal of the Artificial Intelligence and Health Communication symposium was to advance the conceptual design of automated systems that provide health services to patients and consumers through interdisciplinary insight from artificial intelligence, health communication and related areas of communication studies, discourse studies, public health, and psychology.

There is a large and growing interest in the development of automated systems to provide health services to patients and consumers. In the last two decades, applications informed by research in health communication have been developed, for example, for promoting healthy behavior and for managing chronic diseases. While the value that these types of applications can offer to the community in terms of cost, access, and convenience is clear, there are still major challenges facing design of effective health communication systems. The goal of this symposium was to create an interdisciplinary task force to address these challenges by the use of artificial intelligence tech-

niques in combination with empirically-based theoretical frameworks from the field of health communication and related areas of communication studies, discourse studies, public health, and psychology.

The symposium was organized around five main concepts: (1) Patient empowerment and education focused on the definition of concepts such as empowerment and health literacy, on the ways to enhance them through human computer interaction, and on the barriers that might prevent the acquisition of adequate knowledge for self-management of a health condition. (2) Training of health-care professionals centered on the presentation of AI solutions for professional training. Here, the main challenge rests on designing systems that can represent the complexity of interactional scenarios with real patients. (3) Communicating health risks focused on the challenges of empowering people in understanding risks (such as absolute versus relative risks). Different conceptual frameworks for individual and public health risk communication were presented. (4) Conversational agents for health communication included presentation of AI systems to engage with patients for enhancing self-management (for example after hospital discharge). (5) Health informatics focused on the accessibility of structured health information and the related techniques of translation across different terminologies.

The session on designing health communication was a more conceptual session on the ways of reaching patients and consumers on the basis of what is relevant to them, through appropriate language and formats.

Two main themes across the different sections were information versus persuasion and immediacy. To enhance patient participation, it is important that interventions of health promotion provide the right information, in the right format, without attempting to persuade users to comply with it. Argumentation, as a process of supporting a standpoint by giving reasons for its acceptance, was presented as a promising communication technique to involve users in critical discussion toward decision making. Also, there was agreement on the need to enhance immediacy in health communication. Immediacy refers to communication that promotes aspects including physical and emotional closeness, comfort, engagement, and personal involvement. Indeed there are several important outcomes in achieving immediacy at the level of participation including communication of feedback as well as identification and reduced resistance.

The symposium hosted two keynote speakers. Gary Kreps (George Mason University) discussed the role of immediacy in health communication. Eric Horvitz (Microsoft Research) discussed the

cognitive biases at the origin of people's misunderstanding of risk-related health messages and implications for health-oriented web searching.

Overall, the participants found the format of the symposium engaging and constructive, and they expressed the desire to continue this initiative in further events. By popular request, we have established the LinkedIn professional group AI and Health Communication and an online repository for slides, papers, and so on. Subsequently, another symposium will be held in a few years to assess achievements and explore next steps.

Nancy Green, Sara Rubinelli, and Donia Scott served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-01.

Artificial Intelligence and Sustainable Design

The Artificial Intelligence and Sustainable Design symposium focused on domains, problems, and challenges of sustainable design and the role that AI can play in sustainable design.

Long-term environmental and societal sustainability requires that artifacts, materials, systems, and processes be designed to minimize energy and waste and to maximize reuse and utility; we should hope that the days of designing things that are ultimately thrown away are rapidly coming to an end. The "design for X" paradigm considers downstream objectives, such as reusability, early in the design process. Designers are being challenged to consider factors that had been previously given little attention, like life-cycle costs along many dimensions, including energy requirements during manufacture, use, and end-of-use phases, and material loss and environmental damage at the end of a product's life. A vision for sustainable design is cradle-to-cradle design, in which products are designed and built in ways that enable full reuse at low costs (for example, energy), with nothing thrown out and nothing degraded. Our motivation to organize the symposium stemmed from our presumption that the increased complexity of design necessitated by a desire for very long-term planet sustainability requires application of and advances in artificial intelligence.

The symposium brought together researchers from three primary areas: AI and design, computational sustainability, and design for sustainability. AI and design is an established field already, with conferences and journals; AI has provided computational approaches to design processes and the representation of design knowledge. Design for sustainability is not so much a field as it is a set of principles, which are implemented on an ad hoc basis. Computational sustainability is a nascent field, but already influential in moving AI and oth-

er computational fields into addressing sustainability questions.

About 25–30 people physically attended the symposium. There was also a virtual participation option, which was made available to coauthors, colleagues, and students of the authors of the papers as a way of broadening participation without requiring additional travel, and as a result about 10 avatars attended in second life at various times. A total of 18 papers were presented over the course of the two-and-a-half days. Presentation topics included smart buildings, environmentally smart occupants, decision support systems to support energy analyses, optimization to inform design decision making, and theoretical models of cradle-to-cradle design.

The symposium included two invited talks. Alice Agogino (University of California, Berkeley, Department of Mechanical Engineering) described the Smart Lighting project—a wireless sensor network for customizable commercial lighting control. This application requires decision making in the face of uncertainty, with needs for system self-configuration and learning. Kirstin Gail (Autodesk) spoke on Bloom, a recyclable laptop developed by graduate students of a Stanford University course on design innovation; Bloom represents a new class of electronic products that can be easily disassembled for recycling by the consumer at end of product life. In both of these talks, the speakers noted that sustainability was achievable only by considering human behavior as a significant factor. In fact, this recognition of the importance of social influences and implications of human behavior was an overarching theme of the symposium generally.

There were several major sustainability themes in the papers and discussions, including energy conservation, material recycling, life cycle, and environmental impact analysis. The AI and computational themes included approaches to modeling sustainable design knowledge and processes, sustainability through computational support for biologically inspired design, characterizing and managing uncertainty, machine learning and case-based reasoning for sustainable design, and computational tools for estimating sustainability costs in design alternatives.

Participants also engaged in three separate breakout groups to brainstorm on forward-thinking research directions at the intersection of AI and sustainable design. Many themes were discussed, including (1) the role of AI in modeling the effect of human behavior on sustainable design of artifacts and systems; (2) the role of AI in biologically inspired sustainable design, and (3) the role of customization in fitting designs to individuals and groups in a manner that improves sustainability indices.

The symposium was successful in achieving a

common interest among the participants, which we hope will be maintained through a continued effort to communicate and share research on this topic. Our expectation is that AI will provide formal models, languages, and methods of sustainable design, thus helping to push sustainable design from the status of interest area into the status of a discipline.

Douglas Fisher and Mary Lou Maher were cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-02.

Artificial Intelligence for Business Agility

The goal of the AAAI symposium on Artificial Intelligence for Business Agility was to identify appropriate AI methods and to assess their potential for increasing the agility of organizations.

Business agility is a key attribute of successful private and public organization. Agility is the ability of enterprises to quickly adapt to changes in their environment and seize opportunities as they avail. It is a topic of increasing interest in a globalized business environment, in particular boosted by the recent global financial crisis. It is a challenge for any business to sense opportunities or threats, prioritize potential responses, and act efficiently and effectively. Agility requires both ad hoc reactions on what is happening in a specific situation and also adaptation of the organization in the long run. For adaptations, two principles can be distinguished: change and maturation. Whereas the first is an explicit, planned transformation of an enterprise, the latter can be viewed as being a continuous improvement.

Artificial intelligence can contribute in several ways to increasing the agility of organizations. In the first presentation, Richard Colbaugh presented an approach for detecting and characterizing emerging topics of discussion and consumer trends by spotting emerging memes—distinctive phrases acting as tracers for topics—which appear in a large number of blog posts. To predict trends in an early phase, measures are identified that should be predictive of meme success. The metrics form the basis for learning a classifier that predicts, for a given meme, whether or not it will propagate widely.

Simon Brander presented five approaches of process mining that could further increase the agility of processes. The basis of this research is the knowledge-intensive service support developed at FHNW, which tackles agility by combining process modeling and business rules.

The research of Jörg Leukel and colleagues is to build for agile supply chains a logistics knowledge infrastructure that enables enterprises to respond quickly to individual customer demand. They

identified basic and complex logistics patterns that supply-chain actors can reuse for decision making. The logistics knowledge infrastructure is grounded on related AI concepts that are a logistics ontology, knowledge extraction adding facts from semi- and unstructured resources to the knowledge base, and agents representing supply-chain actors.

In their position paper, Hans-Georg Fill and colleagues investigate the opportunities of using functionalities provided by social network sites for the collaborative formalization of semantics in the domain of health. He illustrated the approach with a use case for the eleventh revision of the International Classification of Diseases (ICD-11).

Kristin Glass presented two computational methods for estimating the sentiment social media content, for instance, of blog postings or online product reviews. The first algorithm is a semisupervised sentiment classifier that combines knowledge of the sentiment labels for a few documents and words with information present in unlabeled data. The second algorithm assumes existence of a set of labeled documents in a domain related to the domain of interest and leverages these data to estimate sentiment in the target domain. The proposed methods were shown to outperform several standard methods for the task of inferring the sentiment of online reviews of movies, electronics products, and kitchen appliances.

Elaheh Momeni proposed a methodology and analysis of components required for combining and integrating information data sets from different web data sources into machine-processable ones. The approach is based on a suitable e-commerce ontology. As an example the process of wrapping and matching the data from two public data sets was discussed.

The final presentation of Johannes Magenheim and Knut Hinkelmann was about the role of sociofacts for knowledge maturing. Sociofacts are introduced as the subject-bound knowledge manifestation of social actions. The relations of sociofacts to concepts like artifacts, cognifacts, and action are represented as an ontology. Additionally a formal notation for sociofacts is established. The case of a merger between two companies is used to describe the advantages of sociofact analysis for business agility.

The presentations were supplemented by several lively discussion sessions where specific topics were addressed. In a first session, appropriate AI methods were identified and related to objectives and challenges of business agility. Other sessions specifically dealt with AI methods for process mining and knowledge support for agile processes.

Knut Hinkelmann and Barbara Thönssen served as cochairs of the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-03.

Computational Physiology

The Computational Physiology AAAI symposium examined how insights from physiology and medicine can aid practitioners of artificial intelligence in the discovery of early clinical and physiological predictors of health state.

Automated human health-state monitoring aims to identify when an individual moves from a healthy to a compromised state. For example, changes in diet or physical activity can lead to life-threatening hypo- or hyperglycemia in diabetics. Similarly, elderly individuals managing multiple chronic conditions may experience rapid changes in physical and cognitive health state that must be caught quickly for treatments to be most effective. Even in healthy individuals, heavy exertion in extreme climates can quickly lead to life-threatening situations.

The emergence of inexpensive and unobtrusive health sensors promises to shift the health-care industry's focus from episodic care in acute settings to early detection and longitudinal care for chronic conditions in natural living environments. While these sensing systems are able to provide a wealth of physiological information, these measurements are often quite different from the limited, but high-quality, clinical data used by physicians. Data monitored continuously from sensors requires us to draw conclusions from large quantities of lower-quality data from subacute environments where these measures are often not specific to health states of interest and can reflect the output of multiple latent variables.

The symposium brought together researchers from the fields of artificial intelligence, machine learning, engineering, physiology, and medicine. The symposium's paper presentations and invited speakers addressed areas that covered: (1) feature extraction—especially addressing the problems associated with finding meaningful signals in the noisy data from free-living individuals; (2) onset detection of clinical/physiological events; (3) activity recognition, both as contextual information and as a means of discovering a change in health state; (4) models and their use in the detection of clinically significant events; and finally, (5) the symposium addressed what could be done to better foster interdisciplinary collaboration.

Feature Extraction. Finding meaningful patterns in noisy data is crucial for detecting events in free-living individuals. Techniques to better filter ECG and EEG data were examined, along with methods to extract clinically relevant features from these variables.

Onset Detection. The presenters and speakers detailed clinical health-state problem areas such as sepsis, epileptic seizures, fatigue, hypoglycemic shock, and frustration levels in individuals with autism spectrum disorder. The speakers discussed

both the key clinical aspects of these conditions and how artificial intelligence approaches could prove beneficial.

Activity Recognition. Activity recognition was viewed as important both in providing contextual information and in its own right in determining patient health.

Models. Physiological models developed from empirical data, physiology, and physics were discussed as important components in providing an underlying framework for observed physiological variables. Specifically, presenters discussed both an automated approach to discovering model parameters and the significance in developing an appropriate model to the problem of hyperthermia.

Collaboration. Finally, the symposium discussed how researchers in physiology and medical practitioners are often unaware of the artificial intelligence solutions that can be applied to their open problems. Similarly, the artificial intelligence community has limited awareness of the critical medical and physiological problems that can benefit by collaboration. To address this problem the symposium concluded that assembling several comprehensive medical and physiological data sets for specific problem areas would be of particular interest to the AI community.

Mark Buller, Finale Doshi-Velez, and Paul Cudihy served as organizers and cochairs of the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-04.

Help Me Help You: Bridging the Gaps in Human-Agent Collaboration

The goal of this symposium was to discuss the challenges involved in creating effective interacting systems of human and nonhuman agents. From traffic lights to household cleaning robots, automation is a ubiquitous and growing aspect of our lives. However, the transition from automation to agency, where a software or robotic agent conducts a role and has the responsibility that might have once been trusted only to a human, is still lagging. The aim of this symposium is to begin to form a community of researchers working on different aspects of the problem. The agents and AI communities have tackled this problem from the agent's point of view, while the intelligent user interfaces community has begun to address this problem from the human's point of view. This symposium brought together researchers from these and other communities to foster collaborations, share accomplishments, identify common issues and challenges, and propose directions to move forward.

One major theme of the symposium was the interaction between humans and robots. The

issues discussed included learning behavior with human feedback, enabling a single operator to have supervisory control over multiple entities, developing an understanding of spatial interactions between humans and robots, using cognitive models to help interaction, and social robotics. Another major theme was the area of personal assistant agents. Issues included communication aspects, such as speech recognition and teaching verb phrases, and decision-making aspects, such as human instruction and reasoning under uncertainty. A third major theme was the general notion of collaboration. This involved issues such as learning intentions, action, and plans along with understanding context. The symposium featured several distinguished invited speakers: Tina Beard (NASA), Henry Lieberman (MIT), Karen Myers (SRI), Mark Riedl (Georgia Tech), Paul Scerri (CMU), and Milind Tambe (USC).

There were three panels. The first on the fundamental challenges of human-agent collaboration resulted in an extensive discussion on the appropriate methodology and value of evaluation for problems in this sphere. The second on the "killer apps" of human-agent collaboration involved several panel members who had successfully transferred their technologies. The overarching theme was the significant value in addressing real-world problems. The final panel addressed next steps for the community. There was strong sentiment that the issues and ideas of the symposium should be continued and plans are being made to continue with another similar event and to create venues for disseminating information and developing the field of human-agent collaboration research.

Rajiv Maheswaran, Nathan Schurr, and Pedro Szekely served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-05.

Logical Formalizations of Commonsense Reasoning

Endowing computers with common sense is one of the major long-term goals of artificial intelligence research. One approach to this problem is to formalize commonsense reasoning using representations based on formal logic or other formal theories such as algebraic representations. Formal representations have the advantages that they can easily express partial knowledge; they are extensible; they have a well-defined semantics; they support inference in multiple directions; and the theory of inference is well understood. The challenges to formalizing commonsense knowledge include the accumulation of large amounts of knowledge about our everyday world, the representation of this knowledge in suitable formal languages, the integration of different representations in a coher-

ent way, and the development of reasoning methods that use these representations.

The tenth meeting of the biennial International Symposium on Logical Formalizations of Commonsense Knowledge included 33 papers and posters by authors from 15 countries as well as two invited talks and a panel. Thirteen of the authors were students; eight of these were able to attend the meeting, thanks to a generous grant from *Artificial Intelligence Journal*.

The chief focus of the symposium was on the use of formalized common sense in applications of all kinds. A particularly active area is cognitive robotics, the use of high-level, commonsense reasoning to guide physical robots, either alone or in teams. Gerhard Lakemeyer in an invited talk discussed his work on using formal representations of preservation goals in robots that can plan to maintain their own well-being while carrying out their required tasks (Asimov's third law of robotics). Other talks described the use of commonsense reasoning in housekeeping robots and soccer-playing robots. The closely related issue of interactions among agents—human or robotic—was also much discussed, with talks on commitments and on interpersonal relations.

The issue of applications generally was the subject of a panel, chaired by Leora Morgenstern. Jerry Hobbs spoke about representations used in machine reading; Michael Theielscher spoke about game playing; and Gerhard Lakemeyer spoke about cognitive robotics. Other kinds of applications discussed at the symposium included reasoning about biological processes, and narrative understanding.

On the more theoretical side of the field there were papers on spatial and temporal reasoning, motion, folk psychology, individuation, vagueness, uncertainty, evaluation, and efficient reasoning.

There was also discussion of the field at a higher level. Richmond Thomason gave a personal overview of work on practical reasoning. Hector Levesque proposed an alternative to the Turing test: a machine would be required to disambiguate anaphoric reference in a corpus of sentences constructed so that the correct disambiguation cannot be found without understanding the text and applying commonsense knowledge.

One of the high points of the meeting was the invited talk by Peter Norvig, director of research at Google and coauthor, with Stuart Russell, of the textbook *Artificial Intelligence: A Modern Approach*. The talk was open to the participants of all the sessions at the AAAI Spring Symposium. Google's current research focuses on leveraging its vast data corpora to solve AI problems; this kind of technique has achieved remarkable successes, and Norvig foresees many more. At the same time, he

acknowledged that there are limitations to what can be accomplished this way, and that deeper forms of commonsense knowledge and inference will have to be integrated before true AI can be achieved. We were privileged to be joined that evening by John McCarthy, one of the founders of artificial intelligence in general and the chief founder of the field of automated commonsense reasoning.

Ernest Davis, Patrick Doherty, and Esra Erdem served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-06.

Multirobot Systems and Physical Data Structures

Multirobot systems can store computational data in unique ways. For example, their positions and actions in the physical world can be used to store intermediate state of algorithmic progress. The goal of this symposium was to begin to reason about performing computation using these physical data structures.

Multirobot systems can represent and manipulate physical representations of information by modifying their environment or changing their positions. Information necessary for agent coordination or collective task performance can be externalized (written) into the environment and then used (read) by other robots when and where necessary. Systems are able to exploit the locality of this external information to complete tasks with higher efficiency or enable radically different solutions.

Observations and discussion of physical information representations are found throughout the robotics literature, but it is seldom the primary goal of the work. As a result, there is no unified theory of the cost and complexity of communicating or storing information in this way. The goal of this symposium was to bring together researchers to study aspects of using external information in algorithms for multirobot systems. Ultimately, these techniques should pave the way for multirobot systems that exploit their embodiment and their environment to accomplish tasks more effectively than the current state of the art.

Against a backdrop of practical applications, such as search, exploration, surveillance, foraging, construction, and pursuit-evader problems, we set out to explore four main questions with our discussion:

Analogous properties between physical and computational data structures. Papers included examples of searching robots employing treelike abstractions and formations using a doubly linked-list structure. Although operations like insertion, deletion,

pushing, and popping are realized by physical movements, at an abstract level they still satisfy pre- and postconditions, which enable proofs of correctness and amortized communication costs.

Quantitative versus qualitative information, and continuous versus discrete state. Different forms of stigmergy have been proposed in the literature reflecting whether a physically encoded signal affects either the degree or the type of response. We saw all varieties in the papers presented as part of the symposium and even significant differences when similar problems were addressed. For example, breadth-first search, depth-first search, and rapidly exploring random tree (RRT) algorithms were presented that explicitly store the search progress in the physical configuration of the robots. The encoding of that state varied considerably for these approaches.

Quantifying physically encoded information. Some examples (for example, a robot indicating a direction at a junction) permit straightforward quantification of information as negative entropy. In general, the number of bits encoded per robot also depends on the environmental constraints. One suggestion was that treatment from a communication channel perspective might be more constructive for this type of analysis.

Analysis of tasks from a physical data-flow perspective. O'Hara's paper identifies principles from computer architecture that can be used to structure robot system design: analysis of spatial and temporal locality, for example, influences the dimensions of mobility and pervasiveness. The analogy to Amdahl's rule of thumb for balanced computing systems resulted in much discussion of different preferences for trade-offs between different robot capabilities.

The symposium started with the organizers presenting a broad range of physical data structure examples from biology and robotics. Biological examples included ant decision making with scent trails, hymenoptera agreement and consensus through food sharing (trophallaxis), and the stigmergy used to store state in insect nest architecture. We coordinated with the related symposium on voting processes in order to attend Deborah Gordan's lecture on how low-level encounter patterns in harvester ants mediate colony-level behaviors. From the robotics literature we selected examples such as search and exploration using a physical breadth-first search, construction of modular robotic assemblies with graph grammars that depend on the current physical state of the assembly, and a pursuer-evader problem using a physical line-sweep algorithm.

There were several themes visible in the submitted papers. Perhaps most significant was the proportion of papers that employed the bodies of the robot themselves as the data to be stored or mes-

sage to be communicated. A common theme was using this physical information to enable systems to be successful with minimal sensing. Frequently, information encoded physically obviates the need for the robots to store this information in memory, which in turn allows a minimal controller and sensing model to perform the task.

We concluded the symposium by organizing our thoughts into a taxonomy of physical data structures. We believe the following three dimensions to be novel in this context: (1) Do robot bodies themselves encode the data? In several papers, robots physically positioned themselves so as to encode information. Other instances may lay some other object, for example, breadcrumbs. (2) Are the physically encoded data persistent, volatile, or something between? Things placed in the environment (for example, forming a cluster of bricks or pucks) will persist until some other action removes the items or erases the structure. In contrast, robots moving as part of a formation can encode information (for example, about the formation shape) that disappears as the robots move in the plane. Pheromones are an example at neither extreme because they persist for the short term, but also rely on physical processes to cause decay over time. (3) Are the data structures active or static? When physically deployed items are capable of their own embedded computation (for example, wireless sensor network nodes) then they are active rather than static. The analogy is that they individually and collectively become objects rather than mere records for storing data.

Dylan Shell and James McLurkin served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-07.

Modeling Complex Adaptive Systems as if They Were Voting Processes

The goal of this symposium was to investigate the theoretical properties of modeling voting time and collective agent inference in social insect, robotic, and human systems.

Voting is not often modeled as if it were a complex dynamic process. When information is represented using votes, then these votes can be fused across a network and can transcend incommensurate semantics when trying to integrate and understand multilevel interactions between the many parts of a complex systems model. Voting processes can reveal complicated, volatile, and unexpected insights that can lead to a deeper understanding of highly complex systems that defy traditional approaches. Although voting models have been used to design data fusion systems for engineering solutions, the theoretical properties of voting

methods have not been fully understood or exploited to improve the state of the art. Current models of voting used in explaining social behavior for social insects and robots typically explicate observed patterns of behavior without specifying the connection between vote allocation and a physical process or explaining the timing and dissemination of the collective outcome. The goal of this symposium was to investigate different concepts of voting processes in different types of complex dynamic systems that display high error resilience and to seek explanations by using a voting metaphor.

The symposium brought together researchers from a variety of subfields including voting theory, social insect voting behavior, risk management, customer preference classification, robotic decision making, agent cognition, and human-machine interactions. One major theme of the papers presented at the symposium was the design of data fusion systems based on the theoretical properties of voting rules and systems. The papers on this theme considered the properties of different voting rules in the context of the metaproperty of error-resilient collective outcomes (ERCOs). ERCOs predict the collective outcome at any point in the vote-collection process that would be produced if all the votes could be or were collected.

One paper critiqued the theoretical work on voting models in social insects and robotics research and suggested that studies based on theoretical predictions from analytic and Monte Carlo simulation predictions would improve the rigor of hypothesis testing and enable investigation of new theories about insect and robotic behavior. Another paper replicated existing studies of ERCO scenarios using the Choquet integral. This study, operating with a different probabilistic model of voting, produced consistent ERCO predictions.

Analysis was also done to determine the advantages/disadvantages of using the Choquet method over the modeled voting method. Empirical studies of insect voting behavior provided a strong experimental theme throughout the conference. One paper presented an analysis of honeybee voting for nest site location in which the voting process revealed how honeybees adapt to manage the trade-off between time and accuracy in choosing a nest site. In contrast, a paper on ant colony behavior suggested that communication in ant cultures had little to do with voting by analogy with human processes. Ants seem to act as if they follow local, decentralized rules without apparent coordination. Another major theme was using voting systems to analyze data in online business operations.

One paper discussed the use of alternative voting algorithms to understand the preferences of customers. This study recognized the merits of

comparing different interpretations of customer preferences using more granular voting analysis, but concluded that the computational complexity required for such comparisons was an impediment to experimentation. Another business study considered how voting data from online polls and revealed choices could be used to create a community of shared interests and learning associated with personal medical knowledge.

Applications of voting for multirobotic and man-machine cultures provided a significant design theme. Papers on this theme considered the practice of using simple majority rules to ascertain consensus in robotic and machine decision making without considering the underlying choices about design that can affect the likelihood of reaching a majority consensus or, if applicable, the probability of producing a reliable collective inference. These design issues can affect the robustness and resilience of the data fusion mechanisms. This issue was also discussed in the context of the use of multiagent voting to reach inferences about scientific knowledge from scanned documents that had multiple scanning errors. In this application area, the potential for applications of voting theory to enhance the efficiency and accuracy of analysis has not been extensively explored.

The symposium participants discussed the common practice of inventing plausible voting schemes without scrutinizing the design to determine if a proposed scheme is consistent with the goals of a decision process. It was useful to bring together many different perspectives on the design of voting systems to understand the undeveloped and fragmented theoretical work on fusing voting data in centralized networks. The participants highlighted the need for developing a culture of trust when voting data are collected and shared to increase common knowledge. Legal challenges against unauthorized sharing of personal data, now increasingly common in medical research, must be considered in developing online business cultures that make use of voter data. The participants shared the goal of increasing their awareness of applications of voting theory and brainstormed about options for continuing to learn from each other about the strengths and weaknesses of modeling complex dynamic processes as if they were voting systems. Clearly, understanding and modeling complex dynamic systems is a new, exciting, multidisciplinary and emerging area of research.

Arnold B. Urken served as the chair of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-11-08.

Mark Buller is a Ph.D. candidate in computer science at Brown University, RI, and a principal investigator at the U.S. Army Research Institute of Environmental Medicine, Natick, MA.

Paul Cuddihy is a researcher at General Electric Research, Niskayuna, NY.

Ernest Davis is a professor in the Department of Computer Science, Courant Institute of Mathematical Sciences, at New York University.

Patrick Doherty is a professor of computer science at the Department of Computer and Information Sciences (IDA), Linköping University, Sweden.

Finale Doshi-Velez is a Ph.D. candidate in computer science at Massachusetts Institute of Technology, Cambridge, MA.

Esra Erdem is an assistant professor in the Computer Science and Engineering Department at Sabanci University.

Douglas Fisher is an associate professor at Department of Electrical Engineering and Computer Science, Vanderbilt University.

Nancy Green is an associate professor in the Department of Computer Science at the University of North Carolina Greensboro, USA.

Knut Hinkelmann is a professor for Business Information Systems and dean of studies at the School of Business at the University of Applied Sciences Northwestern Switzerland FHNW.

James McLurkin is an assistant professor in the Department of Computer Science at Rice University

Mary Lou Maher is a senior research scientist at the Human-Computer Interaction Lab, University of Maryland.

Rajiv Maheswaran is a research assistant professor at the Computer Science Department at the University of Southern California.

Sara Rubinelli is group leader at the Department of Health Sciences and Health Policy of the University of Lucerne/Swiss Paraplegic Research, Switzerland.

Nathan Schurr is the artificial intelligence team lead at Aptima, Inc.

Donia Scott is professor at the Department of Informatics of the University of Sussex, UK.

Dylan Shell is an assistant professor in the Department of Computer Science and Engineering at Texas A&M University

Pedro Szekely is a research assistant professor at the Computer Science Department at the University of Southern California.

Barbara Thönnissen is a lecturer at the School of Business at the University of Applied Sciences Northwestern Switzerland FHNW and head of the competence center Information and Knowledge Management.

Arnold B. Urken is a research professor in the College of Engineering at the University of Arizona.

CAMBRIDGE

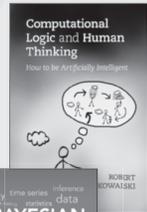
Hot New Titles in AI from Cambridge University Press

Computational Logic and Human Thinking

How to be Artificially Intelligent

ROBERT KOWALSKI

\$110.00: Hb: 978-0-521-19482-2: 336 pp.
\$46.00: Pb: 978-0-521-12336-5



Bayesian Reasoning and Machine Learning

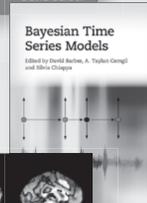
DAVID BARBER

\$90.00: Hb: 978-0-521-51814-7: 650 pp.



Bayesian Time Series Models

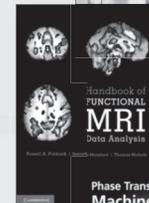
Edited by DAVID BARBER,
A. TAYLAN CEMGIL, SILVIA CHIAPPA
\$80.00: Hb: 978-0-521-19676-5: 432 pp.



Handbook of Functional MRI Data Analysis

RUSSELL A. POLDRACK,
JEANETTE A. MUMFORD,
THOMAS E. NICHOLS

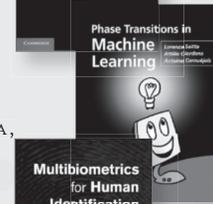
\$80.00: Hb: 978-0-521-51766-9: 450 pp.



Phase Transitions in Machine Learning

LORENZA SAITTA, ATTILIO GIORDANA,
ANTOINE CORNUÉJOLS

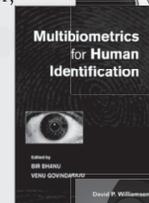
\$90.00: Hb: 978-0-521-76391-2: 410 pp.



Multibiometrics for Human Identification

Edited by BIR BHANU,
VENU GOVINDARAJU

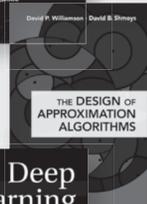
\$90.00: Hb: 978-0-521-11596-4: 408 pp.



The Design of Approximation Algorithms

DAVID P. WILLIAMSON,
DAVID B. SHMOYS

\$55.00: Hb: 978-0-521-19527-0: 516 pp.

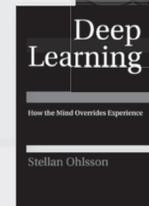


Deep Learning

How the Mind Overrides Experience

STELLAN OHLSSON

\$120.00: Hb: 978-0-521-83568-8: 540 pp.



Prices subject to change.



CAMBRIDGE
UNIVERSITY PRESS

www.cambridge.org/us/computerscience

800.872.7423