

Editorial Introduction

Computational Sustainability

Eric Eaton, Carla Gomes, Brian Williams

■ *Computational sustainability problems, which exist in dynamic environments with high amounts of uncertainty, provide a variety of unique challenges to artificial intelligence research and the opportunity for significant impact upon our collective future. This editorial introduction provides an overview of artificial intelligence for computational sustainability, and introduces the next two special issue articles that will appear in AI Magazine.*

Sustainable development — development that meets the needs of the present without compromising the ability of future generations to meet their needs (United Nations Environment Programme, 1987) — is a critical concern to current and future generations. The emerging interdisciplinary field of computational sustainability (Gomes 2009) draws techniques from computer science, information science, mathematics, statistics, operations research, and related disciplines to help balance environmental and socioeconomic needs for sustainable development. Artificial intelligence (AI) techniques play a key role in computational sustainability research, enabling the solution of sustainability problems that involve modeling or decision making in dynamic and uncertain environments. In turn, sustainability problems present unique challenges that further advance the state of the art of AI.

Since 2011, the main AAAI conference has included a special track on computational sustainability, encouraging AI research in this area and broader participation of sustainability researchers in the AAAI community. The International Joint Conference on Artificial Intelligence (IJCAI) included



The Three Pillars of Sustainability

Figure 1. The Three Pillars of Sustainability.

Sustainable solutions must balance between environmental, societal, and economic demands (United Nations General Assembly 2005). Together, these interdependent aspects are known as the three pillars of sustainability. These pillars have complex interactions at multiple spatiotemporal scales, and so must be analyzed and managed from both local and global perspectives across many contexts. These three pillars also have a nested relationship, with the life-sustaining environment supporting society, which in turn supports the economy (Griggs et al. 2013).

an equivalent track in 2013. Between these conferences, the number of publications on sustainability and AI has grown from 18 published papers in 2011 to 42 papers in 2013. These statistics do not include the numerous other papers on sustainability and AI published in the International Conference on Com-

putational sustainability or in numerous other workshops and symposia on sustainability.

The next two issues of *AI Magazine* will highlight recent AI research in computational sustainability, with an emphasis on projects that have had measurable impact to practical sustainability problems.

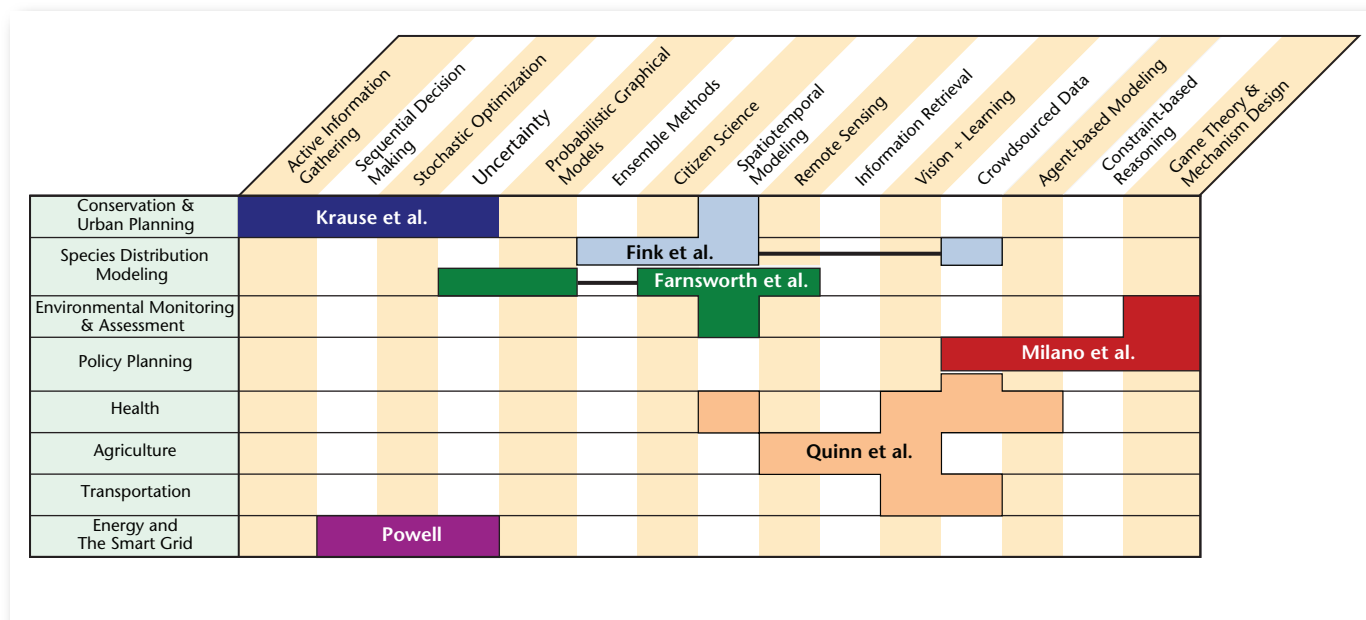


Figure 2. Article Topic Summary.

Summary of the primary computational sustainability topics (rows) and AI-related topics (columns) discussed in each article, showing the diversity of interconnections between AI and computational sustainability research. Each article in these special issues is listed by the authors' names. This figure depicts only the primary topics for each paper; articles may include secondary discussions of topics that are not highlighted (for example, all articles include some aspect of uncertainty, but may not focus on it).

	2011	2012	2013	
Conference and Track	AAAI-11 <i>Computational Sustainability and AI</i>	AAAI-12 <i>Computational Sustainability and AI</i>	AAAI-13 <i>Computational Sustainability and AI</i>	IJCAI-13 <i>AI and Computational Sustainability</i>
Number of Accepted Papers	18	21	16	26

Table 1. Publishing Statistics of Special Tracks on Computational Sustainability at Recent Major AI Conferences.

These articles span a large variety of AI techniques and sustainability problems, from learning large-scale spatiotemporal models for bird migration, to optimal control for energy storage, to constraint-based interactive social policy planning. Together, these articles provide a snapshot of the field that we hope will encourage further participation in computational sustainability.

Overview of The Issues

Problems in sustainability are inherently interdisciplinary, requiring that any solution balance between environmental needs, societal demands, and economic constraints (figure 1). These problems exist across multiple scales in highly dynamic domains with high levels of uncertainty. AI methods for modeling and decision making can enable the creation of

optimal, or near-optimal, policies for sustainable development, but the multiscale, dynamic, and uncertain aspects can present significant computational challenges. Additionally, sustainability solutions must balance between the needs of individuals and groups that interact in highly interconnected, complex manners.

The issues feature articles that discuss problems across a variety of key computational sustainability research areas, including conservation planning, species distribution and ecological modeling, environmental monitoring and assessment, policy planning, health, agriculture, transportation, and energy and the smart grid.

Additionally, these articles span a wide variety of AI techniques and topics, from machine learning and optimization to agent-based modeling and mechanism design. Figure 2 depicts the relationship

between the computational sustainability topics and AI-related topics across all articles, demonstrating the diversity of interconnections between sustainability and AI.

Many of these articles also discuss issues of broader interest that span different AI techniques and applications. Some of these broader issues include compensating for skewed data distributions in crowdsourced data and citizen science (Fink et al.), embracing the unique research opportunities and challenges associated with problems in developing nations (Quinn, Frias-Martinez, and Subramanian), exploiting structural properties such as submodularity to produce efficient solutions with optimality guarantees (Krause, Golovin, and Converse), and unifying problem formulations for stochastic optimization across different subfields, including reinforcement learning, optimal control, and stochastic programming (Powell).

In the following subsections, we briefly describe each article that will appear in this and the subsequent issue of *AI Magazine*.

The Environment

Sequential decision problems that exhibit adaptive submodularity in their structure can be solved efficiently using greedy policies with provable near-optimality. The article *Sequential Decision Making in Computational Sustainability Through Adaptive Submodularity* by Andreas Krause, Daniel Golovin, and Sarah Converse describes the use of adaptive submodularity for decision making in uncertain, partially observable environments. It focuses on two sustainability problems: optimally gathering information for decision making in adaptive conservation management, and dynamically selecting land patches for species conservation in Washington State's South Puget Sound region.

In *Crowdsourcing Meets Ecology: Hemispherewide Spatiotemporal Species Distribution Models*, Daniel Fink, Theodoros Damoulas, Nicholas E. Bruns, Frank A. La Sorte, Wesley M. Hochachka, Carla P. Gomes, and Steve Kelling describe an ensemble method for learning multiscale spatiotemporal models that adapts to variations in spatial sampling densities. They applied this approach to learn species distribution models from eBird data, yielding the first hemispherewide models of bird migration patterns. This article also addresses issues of citizen science and learning from crowdsourced data.

In addition to meteorological monitoring, weather radar provides a remote sensing platform for detecting the movement of birds, bats, and insects. Andrew Farnsworth, Daniel Sheldon, Jeffrey Gevorgyan, Jed Irvine, Benjamin Van Doren, Kevin Webb, Thomas G. Dietterich, and Steve Kelling describe, in *Reconstructing Velocities of Migrating Birds from Weather Radar — A Case Study in Computational Sustainability*, a Bayesian approach for

learning large-scale models of bird migration from radar data, advocating the benefits of joint inference over stage-based pipelined AI systems. The resulting models of large-scale nocturnal bird migration reveal a variety of insights at various spatial and temporal scales that are of interest to ornithologists, ecologists, and policy makers.

Energy

The article *Energy and Uncertainty: Models and Algorithms for Complex Energy Systems* by Warren B. Powell investigates different sources of uncertainty across various problems in sustainable energy systems, while describing a perspective on stochastic optimization that unifies ideas from different fields, including reinforcement learning, optimal control, approximate dynamic programming, and other areas of operations research. These concepts are illustrated through their application to modeling and policy-based control of a solar energy storage system.

Policy Making and Development

In their article *Sustainable Policy Making: A Strategic Challenge for Artificial Intelligence*, Michela Milano, Barry O'Sullivan, and Marco Gavanelli describe aspects of the e-Policy project, a decision support system for policy makers that integrates global and individual perspectives on economic, social, and environmental impacts of different decisions. The article poses policy making as a multifaceted domain containing a variety of challenging AI-related problems, from integrating and balancing between multiple competing objectives, to impact assessment using agent-based modeling, to the use of game theory and mechanism design for sustainable policy creation.

Sustainability problems in the developing world have a variety of unique requirements, stemming from such challenges as scarce resources and lack of infrastructure. In their article *Computational Sustainability and Artificial Intelligence in the Developing World*, John Quinn, Vanessa Frias-Martinez, and Lakshminarayan Subramanian examine the use of mobile devices, social computing, and multiple AI methods in various sustainability applications, with a focus on developing-world issues. These applications range from automated microscopy-based diagnosis of malaria, to cropland monitoring and disease spread tracking using remote sensing, to preventing traffic congestion in developing cities.

Making an Impact on Our Future

These articles are only a sample of current research in computational sustainability and AI; there are a large number of other current projects applying AI to sustainability problems, and a wealth of sustainability problems that have yet to be addressed. One hallmark of sustainability research is the focus on having a measurable impact on our collective future by

addressing current, important problems. To ensure that this research has practical significance, many research groups have partnered with regional government offices (such as the U.S. Fish and Wildlife Service) to deploy the developed technology in small-scale studies. In other cases, AI technology for sustainability is beginning to reach consumers directly through commercial applications (for example, the Green Driver project and smartphone app [Apple et al. 2011]). Due to the interdisciplinary nature of sustainability problems, computational sustainability research is also injecting computational thinking into other fields and fostering the cross-fertilization of ideas. Within the AI community, we hope that continued research in this area will help broaden the AAAI community, while providing a rich source of important new problems to further advance the field of artificial intelligence.

References

Apple, J.; Chang, P.; Clauson, A.; Dixon, H.; Fakhoury, H.; Ginsberg, M.; Keenan, E.; Leighton, A.; Scavezze, K.; and Smith, B. 2011. Green Driver: AI in a Microcosm. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.

Gomes, C. 2009. Computational Sustainability: Computational Methods for a Sustainable Environment, Economy, and Society. *The Bridge* 39(4): 5–13.

Griggs, D.; Stafford-Smith, M.; Gaffney, O.; Rockström, J.; Öhman, M.; Shyamsundar, P.; Steffen, W.; Glaser, G.; Kanie, N.; and Noble, I. 2013. Policy: Sustainable Development Goals for People and Planet. *Nature* 495: 305–307. dx.doi.org/10.1038/495305a

United Nations Environment Programme. 1987. Our Common Future: Report of the World Commission on Environment and Development. Annex to General Assembly document A/42/427, Development and International Cooperation: Environment. Report of the World Commission on Environment and Development. Nairobi, Kenya. New York: United Nations.

United Nations General Assembly. 2005. Resolution adopted by the General Assembly on 16 September 2005: 2005 World Summit Outcome, Resolution A/60/1, adopted by the General Assembly on 24 October 2005. New York: United Nations.

Eric Eaton is a lecturer in the Department of Computer and Information Science and a member of the General Robotics, Automation, Sensing, and Perception (GRASP) laboratory at the University of Pennsylvania. Prior to joining Penn, he was a visiting assistant professor in the computer science department at Bryn Mawr College and a senior research scientist in the Artificial Intelligence Lab at Lockheed Martin Advanced Technology Laboratories. His research focuses on lifelong machine learning, knowledge transfer, and interactive AI.

Carla Gomes is a professor of computer science and the director of the Institute for Computational Sustainability at Cornell University. Her research themes include constraint reasoning, mathematical programming, and machine learning for large-scale combinatorial problems. Recently, Gomes



Support AAAI Programs!

Thank you for your ongoing support of AAAI programs through the continuation of your AAAI membership. We count on you to help us deliver the latest information about artificial intelligence to the scientific community, and to nurture new research and innovation through our many conferences, workshops, and symposia. To enable us to continue this effort, we invite you to consider an additional gift to AAAI. For information on how you can contribute to the open access initiative, please see www.aaai.org and click on "Gifts."

has helped found the new field of computational sustainability, which is her current main research focus. Gomes is a fellow of AAAI and a fellow of the American Association for the Advancement of Science (AAAS).

Brian Williams leads the Model-Based Embedded and Robotic Systems group within the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology. His research concentrates on model-based autonomy, model-based programming, and cooperative robotics. He is a fellow of AAAI, has served as guest editor of the *Artificial Intelligence Journal* and has been on the editorial boards of the *Journal of Artificial Intelligence Research* and The MIT Press.