Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)

A Selected Summary of AI for Computational Sustainability

Douglas H. Fisher
Department of Electrical Engineering and Computer Science, Vanderbilt University
douglas.h.fisher@vanderbilt.edu

Abstract

This paper and summary talk broadly survey computational sustainability research. Rather than a detailed treatment of the research projects in the area, which is beyond the scope of the paper and talk, the paper includes a meta-survey, pointing to edited collections and overviews in the literature for the interested reader. Computational sustainability research has been broadly characterized by AI methods employed, sustainability areas addressed, and contributions made to (typically, human) decision-making. The paper addresses these characterizations as well, which will facilitate a deeper synthesis later, to include the potential for developing sophisticated and holistic AI decision-making and advisory agents.

Introduction

Sustainability is a vast concern, or should be, and presents challenges stemming from interactions between the natural and human-developed spheres across temporal and spatial scales. This has motivated computer science researchers to apply their trade to environmental and societal sustainability challenges. AI is among the most important technologies for enabling humans to deal with the complexity of sustainability challenges.

Gomes (2009) articulated research activity at the nexus of computing and sustainability, labeling it *computational sustainability*, with goals "*to develop new computational models, methods, and tools to help balance environmental, economic, and societal needs for a sustainable future.*" Since that articulation, a dedicated *International Conference on Computational Sustainability* has emerged, as have special tracks at AAAI, IJCAI, and other conferences. It seems time for a comprehensive survey of the field, and attention to possibilities for synthesis of projects that are currently being independently pursued. That comprehensive survey is beyond the scope of this conference paper, but this paper does point to edited collections and overviews elsewhere, as well as describing broad dimensions and categorizations for characterizing computational sustainability work to date.

Computational sustainability research in AI has been previously characterized by (a) the AI methods employed, (b) the sustainability areas addressed, and by (c) the contributions that are made to (typically, human) decision-making (e.g., Fisher, 2012, 2016; Eaton, Gomes, & Williams, 2014). Topic modeling has also been used to find topics in a collection of computational sustainability documents used in university coursework (Fisher, Bian, & Chen, 2016). In that preliminary study, sustainability areas dominated the topic definitions, as opposed to the computational approaches.

The expectation is that categorization of computational sustainability projects will help with a deeper synthesis of these projects down the road, with an eventual goal of this synthesis to be the development of holistic and sophisticated cognitive agents that advise on sustainability problems. The paper also points out potential gaps in the current computational sustainability landscape, to include the relative lack of research on AI for sustainable design and attention to issues of unanticipated consequences that may arise from AI interventions.

Background on Computational Sustainability

Background on computational sustainability can be found in a variety of sources (e.g., Gomes, 2009; Milano, O'Sullivan, & Sachenbacher, 2014; Lässig, 2016), but this section gives the briefest overview, using roughly the structure of the Introduction (written by the author of this AAAI summary paper) of the wikibook entitled "Artificial Intelligence for Computational Sustainability: A Lab Companion" (AIfCS, 2016). Importantly, the wikibook allows the research community to further elaborate these points. For example, if notable projects are missing in the history

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of computing and the environment before the "computational sustainability" moniker, as given in the wikibook, then the history can be updated using protocols associated with the wikibook.

There are three main points to be made here: (1) computing can help manage the very complex sustainability issues; (2) there is a special, even a dominant place for AI in the computational sustainability landscape; (3) computational sustainability conforms to use-driven basic research; and (4) "computational sustainability" has a history that predates the moniker.

Computational Sustainability and Complexity

Sustainability challenges are complex, typically involving many factors. An example of designing protected regions for Grizzly bears (e.g., Gomes, 2009) and other species is illustrated in Figure 1, using a framework that is adapted from Dietterich (2016). This framework, and similar evidence-based decision making frameworks found elsewhere (e.g., Evans & Fisher, 2002), indicates multiple, recurring steps of collecting and consolidating data, finding actionable patterns and other models from the data, and acting on these models, presumably for the benefit of society and environment. These abstract steps are indicated down the center of the figure.

Figure 1 also illustrates that for a given problem domain, such as reserve design, there are many sub-problems (i.e., along to the right side of Figure 1), and many computing areas are implicated in the various decision making steps (i.e., on the left of the figure).

Computational Sustainability is Use-Driven

Computational sustainability research typically fits the ***use-driven basic research*** paradigm of Stokes (1997), in which real-world challenges motivate research that addresses the challenges, but from which approaches and results can also be abstracted so that the abstractions can be applied to problems other than the one that motivated the original work (Bryant, et al, 2011; Fisher, 2012b). Computational sustainability is also an ideal context for projects on AI for the social good (Wagstaff, 2012; AIfSG, 2016).

AI for Computational Sustainability

While computational sustainability admits the participation of all areas of computing, the various subfields of AI have played, and likely will continue to play, a dominant role in computational sustainability use-driven basic research.

At its core, computational sustainability is intended to facilitate improved human problem solving. AI methods can be "cognitive prostheses" (Ford, Glymour, & Hayes, 1997) that can power up myopic human decision making strategies, into hybrid human-computer decision making that is based on holistic understandings and the long view. This possibility applies to individual, cognitively enhanced human decision makers, and to collectives of humans and AIs.

Computational Sustainability History

As noted, the computational sustainability label was introduced in 2008 and 2009 (Gomes, 2009), but research at the nexus of computing and sustainability was underway well before then. This history is partially documented in AIfCS (2016). And as noted, the community can continue to expand the historical report because of the wikibook functionality.

Computational climate models are but one example of "computational sustainability" efforts that are decades old (Weart, 2016), as is work in social computing on human cooperation (Axelrod, 1984). The wikibook gives a richer history of research, courses, and policy developments, to include a list of computational sustainability collections, notably *computational sustainability conferences and conference tracks*. Eaton, Gomes, and Williams (2014) give brief statistics on computational sustainability tracks from AAAI and IJCAI from 2011-2013.

The sustainability areas covered across the history of computational sustainability include such diverse areas as ecological modeling and waste management. To return to just-considered themes of complexity and AI-enhanced decision-making, the future of computational sustainability will be in the holistic joint-consideration of "disparate" areas, such as ecological modeling and waste management. Of course, these and other areas are not disparate at all, but considering them (and others) in conjunction is complicated beyond society's current biases or abilities.

A very recent entry into the history of computational sustainability is *CompSustNet*, an NSF-funded network of many institutions of higher education and other non-government organizations, and government labs. The CompSustNet (2016) site contains descriptions of numerous projects, a blog and other social media pointers, and a schedule of and recorded talks from the *Computational Sustainability Virtual Seminar Series*, which is open to participants from across the planet.

Categorizing Computational Sustainability Research

Computational sustainability research has been previously characterized by (a) the AI methods employed, (b) the sustainability areas addressed, and by (c) the contributions that are made to (typically, human) decision-making. These categorizations may make one of these dimensions primary, and some collections may focus of on particular top-

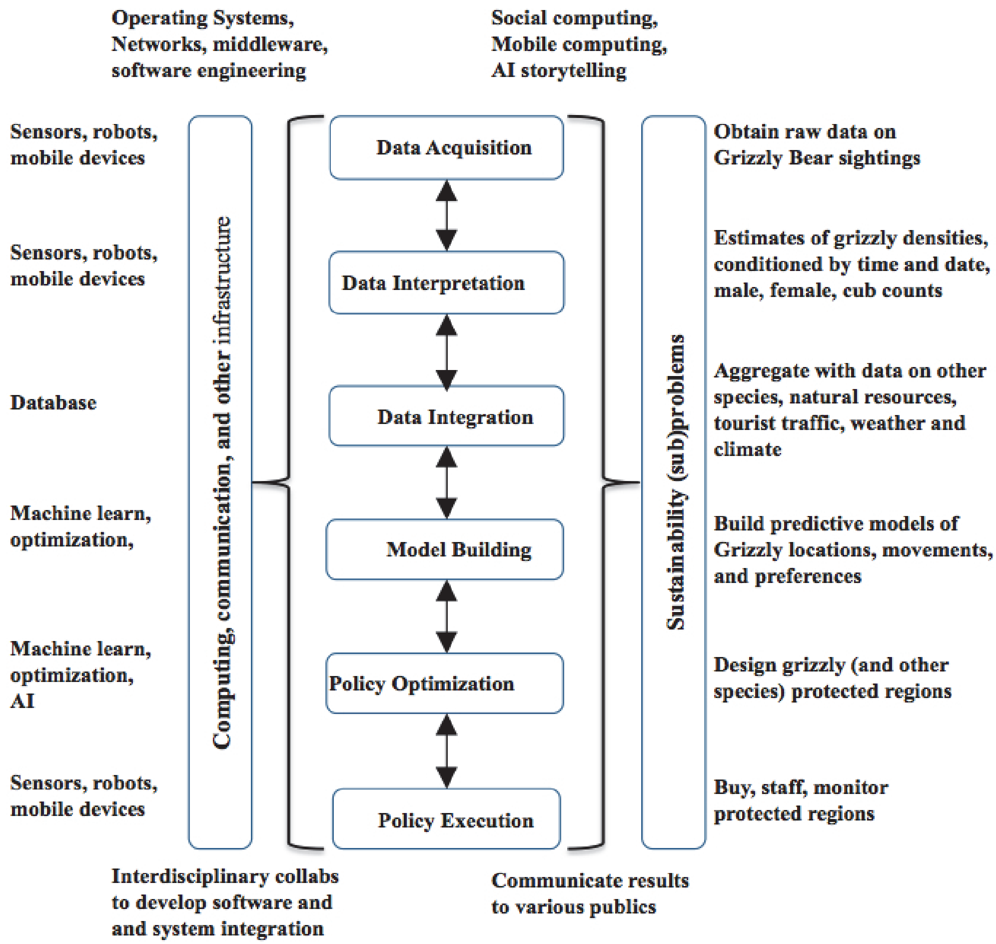

Figure 1: *A systems view of sustainability decision-making. Computing and communications technologies (on left) can be involved all along the decision-making system. The various aspects of a sample application are on the right. Adapted (with heavy annotation) from Tom Dietterich (2016) presentation at AI for Social Good, with original examples along all margins.*

ics along a dimension. For example, Lässig, Kersing, and Morik (2016) is an edited collection in computational sustainability with a significant (but not exclusive) emphasis on energy issues.

Computing-Centric Categorization & Abstraction

Computational sustainability AI research is across the spectrum of AI subtopics. A few works will make explicit categories of AI methods that are implicated. For example, an image found at the bottom of the CompSustNet (2016) webpage, uses the metaphor of a metrorail system, with each rail line (e.g., in blues, red, orange, brown, and black) representing a different type of computational method (e.g., Pattern Decomposition with Big Data; and Agents: Mechanism Design). The stops along these lines correspond to particular projects within the Computational Sustainability Network.

The metaphor and visualization are engaging, though the visualization only scales to a limited extent. In contrast, an interactive version of the rail system allows users to select stops that reveal information about the project, as well as to enable developers to define sustainability-based abstractions over the individual projects.

Fisher (2012) categorizes the papers of the Computational Sustainability special track of AAAI-2011 by four broad AI categories: (1) Optimization/Search; (2) Planning, Control, and Scheduling; (3) Machine Learning: Policy and Action; (4) Machine Learning: Prediction and Classification.

Eaton, Gomes, and Williams (2014) give the most detailed categorization scheme of AI areas in computational sustainability:

(1) Active Information Gathering
(2) Sequential Decision Making

(3) Stochastic Optimization
(4) Uncertainty
(5) Probabilistic Graphical Models
(6) Ensemble Methods
(7) Citizen Science
(8) Spatiotemporal Modeling
(9) Remote Sensing
(10) Information Retrieval
(11) Vision + Learning
(12) Crowdsourced Data
(13) Agent-Based Modeling
(14) Constraint-based Reasoning
(15) Game Theory and Mechanism Design

This metrorail visualization and the matrix representations used by others (Fisher, 2012; Eaton, Gomes, & Williams, 2014) illustrate how *each computational concept and method spans multiple sustainability areas*. Regardless of the particular means of visualizing and otherwise expressing the computation-centric categorization, this strategy facilitates *abstraction* over similar computational approaches that are used in different sustainability contexts. Abstraction is an important mechanism for both characterization, and for later synthesis and integration in decision-making and cognitive architecture frameworks.

Sustainability-Centric Categories & Composition

Characterization of computational sustainability research also is in terms of the sustainability problem or area that it addresses. Recall that Figure 1 shows a decision making pipeline that is annotated by the various computational technologies (more comprehensive than AI per se) that are useful at each step in the pipeline (on the left), and how the various aspects on one sustainability challenge (e.g., the design of ecological protected spaces) would be composed using the framework. *Composition*, like abstraction, is an important mechanism for later synthesis and integration in decision-making and cognitive architecture frameworks.

This pipeline illustrates how each sustainability area spans multiple computational concepts and methods. This is also the case with matrix representations used by others. Fisher (2012) uses a broad categorization of (1) Natural Environment; (2) Natural Resources; (3) Socio-economic areas; (4) Transportation; and (5) the other Built environment. Eaton, Gomes, and Williams (2014) again have the more categories:

1. Conservation & Urban Planning
2. Species Distribution Modeling
3. Environmental Monitoring and Assessment
4. Policy Planning
5. Health
6. Agriculture
7. Transportation
8. Energy and The Smart Grid

Some of these could be mapped into the coarser categories of Fisher.

Other Characterizations and Visualization

There are other categorization schemes that are possible. Fisher (2016) uses the aspects of human decision making that are aided by different computational sustainability projects reported at the AAAI-2016 conference. There appears to be a significant correlation between aspects of human decision making that are implicated by a project, and the AI methods used.

Topic modeling has also been used to find topics in a collection of computational sustainability documents used in university coursework (Fisher, Bian, & Chen, 2016). In that preliminary study, sustainability areas dominated the topic definitions, as opposed to the computational approaches.

Perhaps the real advantage of further developing the topic modeling strategy will be that the representation of topics, and documents in terms of these topics, opens the door to other kinds of data transformations and visualizations as well, involving continuous dimensions, which may prove to be informative and compelling.

Under-Represented Areas

Characterizations of computational sustainability research to date also suggests areas that are under-represented in the current portfolio, as well as those that are better covered and listed in the previous section. This section touches on but a few under-represented areas – a computing area, a sustainability area, and a methodological concern.

Data-Driven Storytelling

There is no research in AI storytelling that this author knows of that is concerned with telling stories about sustainability problems and solutions per se, though there are some promising possibilities, such as the Science of *Data-Driven Storytelling* (SoDDSW, 2016). This could be an area that is important for communicating the science of computational sustainability to different publics, and generalizing the role of citizen scientist to citizen journalist and citizen educator.

Sustainable Design

While there is some attention in sustainable design (e.g., Sundaravaradan, et al, 2011; Oehlberg, Shelby, & Agogino, 2010), it is not a significant footprint in computational sustainability. AI, however, could greatly benefit aspirations towards cradle-to-cradle design (McDonough & Braungart, 2002). The complexity of design for full-reuse of products, homes, materials and the like, in a way that is energy and materially efficient in both manufacturing and recycling, as well as the usage phase, is an important challenge for AI (Fisher & Maher, 2011), which requires a holistic view of the entire lifecycle of human-made artifacts.

Unanticipated Consequences

The ability of AI planning and search technologies to consider many possible futures is a cognitive capability that would greatly benefit human problem solving and decision-making. In particular, the motivation and ability to explore the space of consequences of technology (e.g., Köhler & Erdmann, 2004) and policy interventions is little studied, but unanticipated consequences are not necessarily unanticipatable consequences. It is surprising, for example, that energy efficiency of products is still promoted without regard to the potential rebound effects that per-unit energy savings are associated with an increase in collective energy footprints (Jevons, 1866; Fisher, 2012b; Fisher, 2016).

A relevant challenge problem for computational sustainability, akin to challenges posed by Wagstaff (2012) for machine learning that matters, which would also implicate data-driven story telling perhaps, is an *AI-generated environmental impact report* for a non-trivial intervention on the environment.

Concluding Remarks

Computational sustainability has taken hold as a vibrant area of use-driven basic research for AI. The paper has presented broad characterizations of the space that is covered by AI for computational sustainability and provided pointers to collections of computational sustainability research. The paper also articulates some gaps in the current portfolio.

Apropos the decision-making framework of Figure 1, and the challenge for an AI that creates environmental impact reports, is a broader vision for creating comprehensive environmental decision support systems (Cortes, Anchez-Marre, & Ceccaroni, 2000) that are cast in cognitive architecture frameworks (Langley, Laird, & Rogers, 2008). Comprehensive environmental advisors of the future would rely on many, if not all, of AI's sub-disciplines. Their importance would be to help mitigate human myopia and the unanticipated consequences that result from technological interventions in the environment, to include AI technologies.

Acknowledgements

This work was supported by NSF Grant #1521672 to Doug Fisher at Vanderbilt University for research, education, and reporting on advances in computing and sustainability. This is part of the NSF Expeditions collaborative award that supports CompSustNet (2016).

References

AIfCS (2016). *Artificial Intelligence for Computational Sustainability: A Lab Companion* (2016, June 17). Wikibooks, The Free Textbook Project. Retrieved November 30, 2016 from https://en.wikibooks.org/w/index.php?title=Artificial_Intelligence_for_Computational_Sustainability:_A_Lab_Companion

AIfSG (2016). Artificial Intelligence for Social Good, Information retrieved December 1, 2016 at http://cra.org/ccc/events/ai-social-good/

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books. Adaptation retrieved from http://www-ee.stanford.edu/~hellman/Breakthrough/book/pdfs/axelrod.pdf

Bryant, R. et al., (2011). *Science, Engineering, and Education of Sustainability: The Role of Information Sciences and Engineering*, Computing Community Consortium of the Computing Research Association, 2011. Retrieved November 30, 2016 at http://archive2.cra.org/ccc/files/docs/seesit/RISES_Workshop_Final_Report-5-10-2011

CompSustNet (2016). Retrieved on November 30, 2016 at http://www.compsust.net.

Cortes, U., Anchez-Marre, M. S., Ceccaroni, L. (2000). Artificial Intelligence and Environmental Decision Support Systems, Applied Intelligence 13, 77–91, Kluwer Academic Publishers

Dietterich, T. (2016). Presentation at AI for Social Good, Retrieved on November 30 at http://cra.org/ccc/artificial-intelligence-social-good-speakers/

Eaton, E., Gomes, C., & Williams, B. (2014). Editorial Introduction to the Summer and Fall Issues on Computational Sustainability, *AI Magazine*, V. 35, No. 3. Association for the Advancement of Artificial Intelligence.

Evans, R., & Fisher, D. (2002). "Using decision tree induction to minimize process delays in the printing industry," in W. Klosgen and J. Zytkow (Eds.), *Handbook of Data Mining and Knowledge Discovery*. Oxford University Press, 874--881.

Fisher, D. (2012a) "Recent Advances in AI for Computational Sustainability" (overview of the papers of the AAAI-11 Special Track on Computational Sustainability). In *IEEE Intelligent Systems column on AI and Sustainability*, V. 27, N. 4.

Fisher, D. (2012b) "Sustainability" in *Leadership in Science and Technology: A Reference Handbook*, W. S. Bainbridge, Ed: SAGE Publications, pp. 201-209.

Fisher, D. (2016) "Recent Advances in AI for Computational Sustainability" (overview of the papers of the AAAI-16 Special Track on Computational Sustainability). In *IEEE Intelligent Systems* column on AI and Sustainability, V. 31, N. 4.

Fisher, D., Bian, Z., and Chen, S. (2016) "Incorporating Sustainability into Computing Education" In *IEEE Intelligent Systems column on AI and Sustainability*, V. 31, N. 5.

Fisher, D., & Maher, M. L. (2011). CCC Blog on AAAI 2011 Spring Symposium on Artificial Intelligence and Sustainable Design. Retrieved on November 30, 2016 from http://www.cccblog.org/2011/04/04/a-role-for-artificial-intelligence-in-sustainable-design/

Ford, K., Glymour, C., & Hayes, P. (1997). "Cognitive Prostheses", *AI Magazine*, V. 18, N. 3, p. 104.

Gomes, C. (2009). "Computational Sustainability: Computational Methods for a Sustainable Environment, Economy, and Society" *The Bridge, Frontiers of Engineering Vol 39* No. 4 Winter National Academy of Engineering.

W. S. Jevons, (1866). *The Coal Question: An Enquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-mines*. London, UK: Macmillan.

Köhler A. & Erdmann, L. (2004) "Expected Environmental Impacts of Pervasive Computing" *Human and Ecological Risk Assessment*, 10: 831–852.

Langley, P., Laird, J., and Rogers, S. (2008). "Cognitive architectures: Research issues and challenges", *Cognitive Systems Research*, Vol 10, No. 2, pp. 141-160.

Lässig, J. (2016). "Sustainable Development and Computing – An Introduction" In J. Lässig, K. Kersing, & K. Morik (Eds). *Computational Sustainability*, Springer, pp. 1-12.

J. Lässig, K. Kersing, & K. Morik (2016). *Computational Sustainability*, Springer

McDonough, W., & Braungart, M. (2002). *Cradle to Cradle: Remaking the Way We Make Things*. New York: North Point Press.

Milano, M., O'Sullivan, B., & Sachenbacher, M. (2014). "Guest Editors' Introduction: Special Section on Computational Sustainability – Where Computer Science meets Sustainable Development" In *IEEE Transactions on Computers*, Vol. 63, No. 1.

Oehlberg, L., Shelby, R., & Agogino, A. (2010). "Sustainable Product Design: Designing for Diversity in Engineering Education", *International Journal of Engineering Education*, V. 26, N. 2, pp. 489–498.

SoDDSW (2016). The Science of Data-Driven Storytelling Workshop, West Data Hub, Retrieved on November 30, 2016 from http://www.cccblog.org/2016/06/14/the-science-of-data-driven-storytelling-workshop/

Stokes, D. E. (1997). *Pasteur's Quadrant: Basic Science and Technological Innovation.* Washington, DC: Brookings Institution Press.

Sundaravaradan, N., Patnaik, D., Ramakrishnan, N., Marwah, M., Shah, A. (2011) "Discovering Life Cycle Assessment Trees from Impact Factor Databases," *Twenty-fifth AAAI Conference on Artificial Intelligence*.

Wagstaff, K. (2012). "Machine Learning that Matters" Appearing in *Proceedings of the 29th International Conference on Machine Learning,* Edinburgh, Scotland, UK.

Weart, S., *The Discovery of Global Warming, American Institute of Physics*. Retrieved on Nov. 30 at from http://www.aip.org/history/climate/GCM.htm