Recent Advances in AI for Computational Sustainability
Douglas H. Fisher

Preprint of Sustainability column in *IEEE Intelligent Systems*, *Volume 31, Issue No. 4* (July-August), 2016.

In 2012 I surveyed papers on computational sustainability at the AAAI 2011 conference [1]. Computational sustainability encompasses research on ***computational contributions to decision making*** that affects environmental and societal sustainability.

In 2012 computational sustainability was a nascent field, with its own conference, having been recently introduced to AAAI, and a short time later to IJCAI, the two flagship conferences of artificial intelligence. My survey classified the papers of AAAI into computational categories and sustainability categories. The computational categories were

- optimization and search;
- planning, control, and scheduling;
- policy and action learning; and
- prediction learning, both supervised and unsupervised.

The sustainability categories to which I labeled computational research and development were the

- natural environment;
- natural resources (for human consumption);
- socio-economic dimensions of sustainability;
- transportation; and the
- built environment, other than transportation.

In addition to these two dimensions, the earlier survey discussed what amounted to a third dimension, descriptive of the decision making process itself, to include processes of

- data collection and processing that inform the decision making process;
- model-based reasoning and simulation, with models constructed from data and/or from "first-principles";
- the extent that decision making was constrained by, and reactive to, established physical and social designs; and
- the extent to which decision making included deliberative design, ideally that would beneficially constrain later reactive decision making.

Decision making typically involves non-trivial interactions between humans and computation, which might well be more explicit in other conferences on human-computer interactions than they are in AI-dominant conferences that emphasize the non-human, computational components.

In this article, I similarly frame the papers of AAAI-16 along the same dimensions, but sustainability challenges are the primary organizing dimension in this survey. Decision making and computational methods are ***not independent***, and I use decision making (and not computation *per se*) as the vertical dimension in the matrix of Figure 1. Several papers are concerned with building models and knowledge from data (center, along the vertical), and make no strong commitment to the interventions that this knowledge will inform. Of the remaining papers, some study interventions that are strongly constrained by existing designs (bottom, along the vertical), such as traffic interventions on existing road networks. Other papers focus on more open design (top, along the vertical), much less constrained by existing "infrastructure," such as the design of new road networks. Importantly, there is not a sharp boundary between constrained intervention and open design. Thus, rather than a simple matrix, the top and bottom of Figure 1 should be viewed as connecting, like a cylinder. To suggest the intended wrap around, Figure 1 references one paper, shown in lighter blue, that lies in the gray area between constrained intervention and open design.

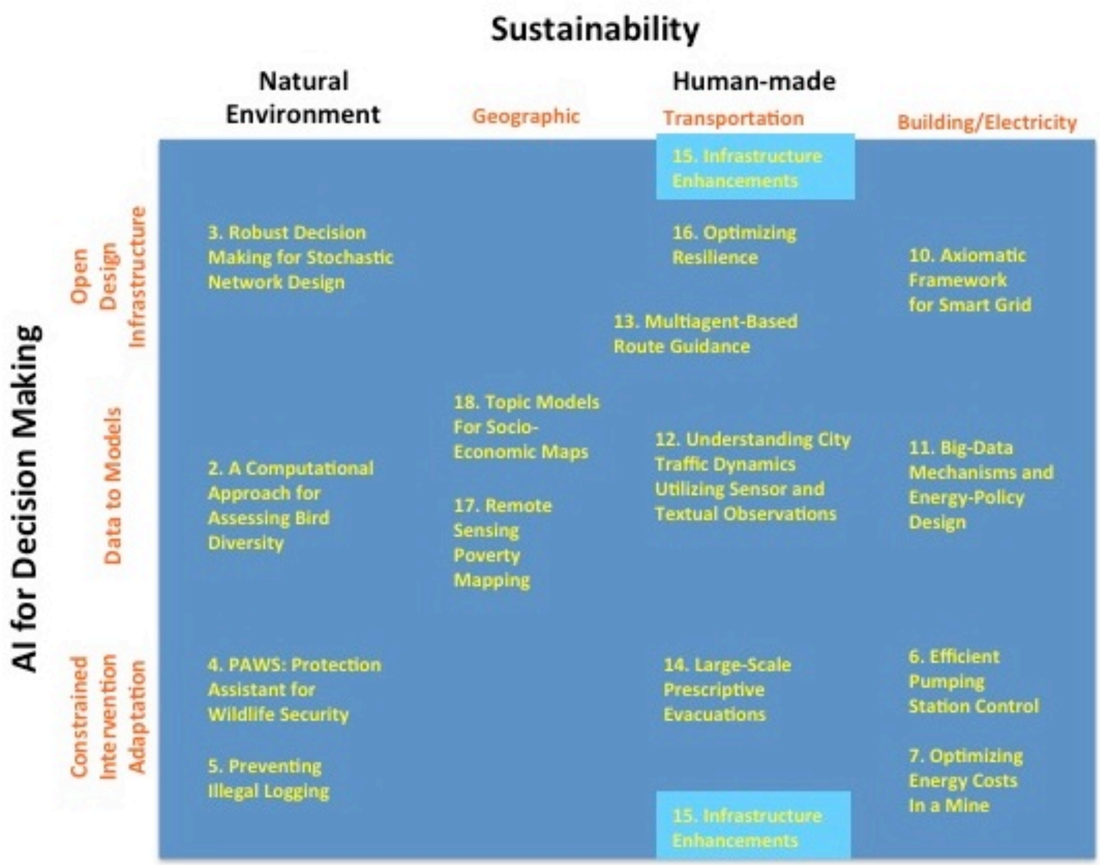

Figure 1: Papers of the AAAI-2016 special track on Computational Sustainability are organized horizontally by a main sustainability theme and vertically by their contribution to decision making processes. Each paper is indicated by an abbreviated title, and by the number of the paper as it is listed in the References section. It is intended that open design of new infrastructure (top of vertical) and constrained decision making under existing infrastructure (bottom of vertical) form a continuum, implying wrap-around, with paper 15 in light blue being illustrative of a paper that falls between open (top) and constrained (bottom).

AI across Sustainability Challenges

My survey of AAAI-2011 identified papers clearly concerned with the natural environment (i.e., acquisition of land reserves to protect species, ocean monitoring, estimating species distributions); papers that addressed natural resources (i.e., timber harvesting; water, both for hydroelectrically and residential use; and crop disease); papers identified as socio-economic (i.e., on infection propagation; and distributed agents in smart electrical grids); papers in transportation, all by personal vehicle; and papers on other aspects of the built environment (i.e., product design, building energy control).

This year's papers inhabit most of the same general regions, even though they differ from earlier papers in many specifics.

AI and the Natural Environment

Four papers at AAAI-16 were most concerned with AI that benefits decision making about the natural environment. Briggs, et al [2] developed unsupervised machine learning algorithms that substantially ease the task of labeling bird vocalizations in the wild by identifying audio samples to be labeled by human experts. After a relatively small number (e.g., 100, 1000) of informative audio clips are expert labeled, the labeling can be extended to a larger archive that is two-to-three orders of magnitude greater than the human-labeled sample. These acoustic records have clear applicability to characterizing the distribution of bird species, arguably doing a better job than visual sightings, which may be more difficult.

Kumar, et al [3] addresses conservation planning for preserving wildlife species habitat. Given parcels of land that can be purchased to ensure the safety of a species, only some of which are initially occupied by representatives of the species, [3] identifies a network of parcels, subject to budget constraints, that maximize the expected diffusion through the network by the species. The paper addresses the problem of estimating parameters of the network, notably the probability that representatives of the species will move from parcel to parcel, which prior work has taken as given.

Fang, et al [4] report on the deployment of a game-theoretic decision aid to protect African wildlife from poaching. The Protection Assistant for Wildlife Security (PAWS) schedules patrols by wildlife agents in a manner difficult to anticipate by poachers, while also taking into account the varying priorities of species and locations. This is one example of what the authors call "*green security games*", a special case of "*security games*" used in other applications (e.g., protecting airports). A second green security application presented at AAAI-16 also employed game-theoretic techniques to the problem of impeding illegal logging [5]. While the green security games [4, 5] are similarly intended, each paper presents its own theoretical and practical advances.

The distinction between applications that protect the natural environment from those that manage natural resources for human consumption can be subtle. Both of the green security applications might have been primarily labeled as natural resource applications, for example, if the intent of the project was to protect legal timber harvests or tourism dollars, rather than flora and fauna *per se*.

AI and Electricity Usage

There were four papers at AAAI-16 on *energy* in various human-made settings. Papers by Kanters et al [6] and Kinsella et al [7] address the problem of reducing energy costs to industries and utilities, particularly for pumping water, by using planning and optimization methods that exploit fluctuating electricity pricing.

Without overtly stating it, these papers make evident that there is a difference between *energy savings* and *energy cost savings*, though they are related. Variable pricing across minutes, hours, days, and seasons, are a class of mechanisms for balancing the demand for electricity, and other resources, and in particular to subtract demand from historic peak periods. Because electricity providers ramp up for these peak periods and other anticipated "worst case" scenarios, reductions in predictable peaks save energy because providers need not generate as much. With an increasing portfolio of energy sources, both fossil and renewables, the volatility of energy costs will magnify both challenges and opportunities for producers and consumers to operate cost effectively, as these papers show.

It remains to be seen how consumer abilities to save on energy costs by adjusting to variable pricing will impact overall energy demand and use. Rebound effects, acknowledged at least since Jevons 1866 [8], result when energy efficiency increases, leading to lower energy "per unit" energy costs, which in turn leads to greater overall usage of energy across sectors [9]. How might we anticipate rebound effects, if any, of new AI-enabled consumer capabilities beyond the very short term? Two other papers are suggestive in this regard.

Bandyopadhyay, et al [10] introduce an axiomatic framework to guide dynamic pricing mechanisms on smart electrical grids. For example, two (of five) axioms are stated abstractly as

1. "The higher is the consumption of electricity by a customer, the higher is the cost she has to pay." [10, p 2]
2. "The cost of electricity should be designed in such a way that it would encourage consumers to reduce consumption during peak demand hours." [10, p 2]

This statement of axioms – of explicitly laying all assumptions on the table – is a substantial contribution of this paper. Even if some might take issue with some of these axiomatic prescriptions, the attitude that one should state assumptions explicitly is critical for reasoning through the implications of a technology intervention. The paper shows that not all of the five axioms, in their formal instantiations, can be satisfied simultaneously, though certain proper subsets of the five can be simultaneously satisfied. For example, under some schemes, some users might use more electricity than others, but because the former do so primarily during non-peak hours, the former pay less (violating axiom 1, while adhering to axiom 2).

A paper by Pat, et al [11] explores the effectiveness of incentives that an electrical utility can offer to customers that allow that utility to control home thermostats for purposes of reducing an/or energy usage. A utility's policy is a subset of incentives that it implements to influence customer behavior. As with [10], this work requires a statement of assumptions, in the form of expectations and preferences, by utilities and by customers.

These papers [10, 11] do not illustrate consequences of technological interventions beyond the very short-term, but the potential of axiomatic and policy approaches for mitigating human myopia, and thereby mitigating rebound effects and other unanticipated consequences, motivates this work nonetheless. Indeed, it is probably for long-term projections and long-term planning that society most needs AI and computational sustainability [9].

AI and Transportation

Five papers at AAAI-16 addressed transportation. Two of these were concerned with traffic in urban settings, while three were concerned with evacuations as a result of disasters.

Anantharam et al [12] combine machine learning from quantitative sensor data with qualitative data from tweets and official traffic reports to reason about traffic. The hybrid approach learns piecewise linear models, one for each of the 168 hours of a week, from sensor data to represent "normal" traffic patterns in a region. Sufficient deviations from a normal pattern signals a possible anomaly, triggering a search for qualitative data that will explain the anomaly. While the paper is limited to showing that traffic patterns can be characterized, there are clearly applications to emergency and routine route planning.

Cao et al [13] describes an approach to routing vehicles so as to maximize the likelihood of on-time arrival, relative to each vehicle's desired arrival time at a destination. Importantly, this in not a system hyper-focused on one driver's desires, but dynamic routing is intended to yield a good expected solution across all vehicles. For example, by directing some drivers with more relaxed deadlines along less traversed paths, other drivers with tighter deadlines can benefit from less congestion. If this work is viewed as purely reactive to the current state of traffic systems, there would be many who would rightfully question whether most drivers would opt for globally optimal solutions over their own locally optimal solutions; and others who would suggest that rebound effects would result from more time-efficient individual travel, with global (spatial) routing akin to the electrical (temporal) load balancing referenced above. Though its not explicit in the paper, this work is better viewed as an approach that is appropriate to a future of autonomous vehicles and/or for collecting data (from the project's traffic simulations) to inform the design of future road networks. For these reasons, the paper has been placed between the open design and data-to-knowledge "rows" of Figure 1, though a different perspective might reasonably place it under constrained, reactive intervention.

The remaining three transportation papers are concerned with evacuation planning. Romanski and Van Hentenryck [14] apply a divide and conquer methodology to the task of determining evacuation plans (e.g., for town and city settings) in advance of disasters. Here, candidate plans are "*convergent*" in the sense that the plans do not fork evacuees onto diverging roads, since this latter practice can slow evacuation due to "driver hesitation". These prescriptive evacuation plans, versus ad hoc self-evacuations, can then be promulgated to regional residents in advance. The paper highlights an important principle – that in addition to being computationally rich, *sustainability planning requires attention to human factors* (e.g., convergence), if it is to be adopted and successful.

Kumar et al [15] assume the same basic planning problem as [14], but address the issue of how to upgrade road networks (e.g., through the addition of lanes and the heightening of roads above the flood plane) so as to improve evacuation plans. As such, [15] takes an important step into infrastructure design under budget constraints, albeit design that is limited to revisions of existing road networks. This paper has been placed in the middle ground between open design and constrained intervention in Figure 1, though it would probably be placed in the latter if a binary choice was forced.

Wu et al [16] takes a further step into deliberative design of road networks under the prospect of evacuations necessitated by disasters, again under budget constraints. This paper [16] employs different formalisms than the other evacuation-motivated papers [14, 15], and [16] is concerned with the design of a road network "from scratch", by introducing road linkages between pre-determined pairs of locations (e.g., hospitals) and optimizing for *resilience* of the network to certain failures (e.g., a flooded road).

Taken collectively, the transportation papers highlight important distinctions that those who work in computational sustainability must attend to, most notably differences along a continuum, ranging from decision making that is constrained by existing infrastructure to deliberative design from scratch. Particularly in the case where AI is being used to remediate a bad existing design, we should worry about rebound effects and other unanticipated consequences. Making dysfunctional downtown traffic somewhat more livable in the short term with AI, for example, may not be nearly so beneficial as banning parking downtown and using AI to create "optimal" parking centers at the outskirts of a city, which are frequently serviced by public transportation.

There are other important differences stemming from the nature of the tasks expressed across the transportation papers. For example, approaches to evacuation planning [14, 15] constrain routes for vehicles to be convergent. We can imagine that other approaches might simply favor such convergent routes as a soft constraint, rather than as a hard constraint. Contrast this with [13], which presumably has no explicit bias towards convergence, and indeed, given the desiderata of on-time arrival, may route vehicles in highly divergent ways. Explicitly looking at the differences between these transportation papers begs questions like those that follow.

- Can approaches to on-time arrival planning be adapted to the task of *ad hoc* self-evacuation planning, which might complement or conflict with prescriptive *en masse* evacuation planning?
- What designs of road networks emerge when planners are simultaneously concerned with prescriptive evacuation routing and expected on-time arrival routing during routine commuting?

AI and Socio-Economic Maps

While most papers above were on problems that included important, integral social components, two papers in the AAAI-16 computational sustainability track were primarily social in nature, and both of these were focused specifically on creating maps of socio-economic status, albeit from very different data sources. Xie et al [17] used data from satellite imagery, extracting features such as identifiable roads and nighttime light intensity, to identify poverty areas using deep learning with artificial neural networks. In contrast, Hong et al [18] used data from mobile phones, with geo-location data on the source of calls and the location of the target of calls, over time, to identify more general socio-economic maps using topic modeling.

The approaches taken are very different, and presumably complementary, begging questions as to how the results would compare when applied to the same regions. Both papers are motivated by a desire to inform decisions by policy makers. It would be a wonderfully important longitudinal study to see how these and similar approaches actually do inform policy, not just in the short term by reducing the need for expensive and infrequent interview-based censuses, but how decision makers of the future use this knowledge. Is it possible to predict this future usage now, with associated uncertainties, or are future practices rife with *truly unanticipated consequences*? These are the big questions that we need AI to help us address.

Concluding Remarks

In the 2012 summary, I wrote that when the papers were considered collectively, they suggested the potential for an environmentally minded cognitive agent of some sophistication, rather than simply disparate decision support tools. This remains true of this year's crop of papers. Even if such agents are not on the near- or medium- horizon, the agent perspective gives us a systematic way of reflecting on the portfolio of computational sustainability tools that have been built, and that should be built.

Papers also vary in the directness of the link between the stated goals of a project and environmental impact. This traceability to environmental consequences is not something that research communities should forget. Desiderata such as arrival time and energy *cost* savings have clear connections to environmental costs and benefits, but we should not forget to elaborate them, and to evaluate interventions in terms of environmental consequences, which could be once or twice removed, as well as projects' zeroth-order effects relative to immediate project goals. It is only in this way that we will remain conscious of rebound effects and other unanticipated (but not "unanticipatable") consequences, and the importance of engaging with social and behavioral scientists if we want our tools and methodologies to be adopted, adapted, and to benefit the natural environment and society.

References

1. D. Fisher, "Recent Advances in AI for Computational Sustainability," IEEE Intelligent Systems, vol. 27, no. 4, 2012, pp. 75-79. Preprint at: http://www.vuse.vanderbilt.edu/~dfisher/IEEEIS/RecentAdvancesInAIforCompSust.htm

2. F. Briggs, X. Fern, R. Raich, & M. Betts, "Multi-Instance Multi-Label Class Discovery: A Computational Approach for Assessing Bird Diversity," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp 3807-3813. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12214/12164 (June 1, 2016).

3. A. Kumar, A. J. Singh, P. Varakantham, & D. Sheldon, "Robust Decision Making for Stochastic Network Design," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3857-3863. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12224/12171 (June 1, 2016).

4. F. Fang, T. H. Nguyen, R. Pickles, W. Y. Lam, G. R. Clements, B. An, A. Singh, M. Tambe, A. Lemieux, "Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3966-3973. Retrieved from http://www.aaai.org/ocs/index.php/IAAI/IAAI16/paper/view/11814/12315 (June 1, 2016).

5. S. McCarthy, M. Tambe, C. Kiekintveld, M. L. Gore, & A. Killion, "Preventing Illegal Logging: Simultaneous Optimization of Resource Teams and Tactics for Security," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3880-3886. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12160/12174 (June 1, 2016).

6. T. V. Kanters, F. A. Oliehoek, M. Kaisers, S. R. van den Bosch, J. Grispen, & J. Hermans, "Energy- and Cost-Efficient Pumping Station Control," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3842-3848. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12035/12169 (June 1, 2016).

7. A. Kinsella, A. F. Smeaton, B. Hurley, B. O'Sullivan, & Helmut Simonis, "Optimizing Energy Costs in a Zinc and Lead Mine," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 4022-4027. Retrieved from http://www.aaai.org/ocs/index.php/IAAI/IAAI16/paper/view/12020/12323 (June 1, 2016)

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

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

10. S. Bandyopadhyay, R. Narayanam, P. Kumar, S. Ramchurn, V. Arya, & I. Petr, "An Axiomatic Framework for Ex-Ante Dynamic Pricing Mechanisms in Smart Grid," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3800-3806. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12032/12163 (June 1, 2016).

11. A. Pat, K. Larson, & S. Keshav, "Big-Data Mechanisms and Energy-Policy Design," AAAI Conference on Artificial Intelligence, AAAI, 2016, 3887-3893. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12251/12175 (June 1, 2016).

12. P. Anantharam, K. Thirunarayan, S. Marupudi, A. Sheth, & T. Banerjee, "Understanding City Traffic Dynamics Utilizing Sensor and Textual Observations," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3793-3799. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12047/12162 (June 1, 2016).

13. Z. Cao, H. Guo1, J. Zhang, & U. Fastenrath, "Multiagent-Based Route Guidance for Increasing the Chance of Arrival on Time," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3814-3820. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11797/12165 (June 1, 2016).

14. J. Romanski & P. Van Hentenryck, "Benders Decomposition for Large-Scale Prescriptive Evacuations," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3894-3900. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11949/12176 (June 1, 2016).

15. K. Kumar, J. Romanski, & P. Van Hentenryck, "Optimizing Infrastructure Enhancements for Evacuation Planning," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3864-3870. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11950/12172 (June 1, 2016).

16. X. Wu, D. Sheldon, & S. Zilberstein, "Optimizing Resilience in Large Scale Networks," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3922-3928. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12311/12180 (June 1, 2016).

17. M. Xie, N. Jean, M. Burke, D. Lobell, & S. Ermon, "Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3929-3935. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12196/12181 (June 1, 2016)

18. L. Hong, E. Frias-Martinez, V. Frias-Martinez, "Topic Models to Infer Socio-Economic Maps," AAAI Conference on Artificial Intelligence, AAAI, 2016, pp. 3835-3841. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11757/12168 (June 1, 2016).