Representing the search spaces of everyday human sustainability problems

M. Six Silberman and Bill Tomlinson
Department of Informatics, University of California, Irvine
msilberm | wmt @ uci.edu

Summary. Some sustainability-relevant "everyday" human problems are usefully supported by data and algorithms. However, everyday problems tend not to admit optimal solution in practice, due to complexity and/or ill-posedness. In practice, people use heuristics to prune the search spaces of these problems. Since people's lives are similar, their search spaces are similar. Some tools exist to provide people with access to information to support efforts to solve these problems, but many problems are not (well) supported. When search costs exceed collaboration costs, systems that aggregate information from multiple people can outperform both monolithic and strictly individualistic approaches to sustainability by drawing on local expertise. We present a series of case studies of deployed systems and everyday problems for which no systems yet exist, including a project we are deploying for real users. We conclude with a representation of the space of these problems and relevant systems.

Detail. Alice identifies as an environmentally responsible consumer, and has heard about "food miles" and the relation between locally produced food and sustainability. Without necessarily being aware of the underlying mathematics, she decides to attempt to solve the problem $\min D$, where D is the distance her food has traveled to reach her, subject to the constraints $c \leq B_f$ and $\mathbf{f} \in P_f$, where c is the cost of the food, B_f is her food budget, $\mathbf{f}$ is the vector describing the characteristics of the food (taste, nutrition, and so on), and P_f is the region in the space of foods which includes the foods she is willing to eat. How can Alice solve this problem? This may be the problem she intends to solve, but in practice she faces additional constraints that complicate the task of representing the problem she *actually* solves. These include search costs, time constraints, the cost of transporting herself to the food (these can be interpreted as transaction costs), and the demands of her own schedule.[1] Additionally, reduction of food miles may occur at the expense of increasing the number of miles she must transport herself; is the tradeoff one-to-one? Reductions in D may be less important to Alice after a certain point. To specify the relevant tradeoffs would require a complete representation of Alice's preferences, which is impossible.

The point here, as Runge [1] (slide 5) notes in discussing natural resource management problems, is that the problems individuals face in everyday life are not well-posed: rather they are multiple-objective, spatially explicit (but not necessarily crisply spatially bounded), recurrent, and involve vastly imperfect information. Taking a decision may reveal new information about the search space or prompt a reframing of the problem. Each decision may be an opportunity to solve a different problem.

So, how does Alice deal with this problem in practice? She uses a heuristic to prune the search space. Because most places that sell food don't advertise the distance their ingredients have travelled, search costs are high and the distribution of options is unknown. We can expect a heuristic like "only buy food from the local farmers' market" to be a good one. But how good (especially if Alice has a relatively low value for B_f)? How do we discover and evaluate these heuristics for other everyday problems?

Consider Alice, Bob, Carol, and David, neighbors who have all agreed to minimize D, even though they have different values for B_f and P_f (and for other unrepresented constraints). Their search spaces overlap substantially, so if the cost of sharing information is less than the cost of searching, collaboration can reduce total search costs. In particular, a shared representation (a database or "map") of the search space can often reduce search costs. We suggest that there is an underexplored design space of information systems that reduce search costs for everyday human problems in which many people have overlapping search spaces. Further, a map of the search space for one problem may contribute to a map for a different problem. The role of the researcher who designs and maintains such a map is not to solve one problem or another definitively but to provide a resource for the practical, everyday activity of finding good-enough solutions to—and occasionally reframing—the problems at hand.

Many such maps already exist. Geographic maps (and utilities with path finding like Google Maps and

[1] e.g., if Alice must be at A at t_1 and at B at $t_2 > t_1$, and can travel between them at a speed no greater than r, and some food she wants is available at C such that the distance she must travel to stop at C on her way from A to B exceeds $r(t_2 - t_1)$, she will not get food at C in the interval $[t_1, t_2]$ even if doing so would reduce D).

OpenStreetMap) help users minimize the time spent moving between physical locations. Database applications like Yelp help users solve many complex problems involving food preferences and budget constraints. Applications that support secondary markets like Craigslist and eBay help users in particular locations with budget constraints fulfill a wide variety of demands that they would not otherwise be able to. OneBusAway aims to "improve the usability" of Seattle public transportation by representing the search space of bus routes [2]. SourceMap aims to help consumers and businesses understand and represent the provenance of the goods they buy and sell [3]. Yet many everyday problems lack useful search space representations.[2] Alice's effort to reduce the carbon dioxide emissions associated with her food consumption presents one example. Some search space characteristics change slowly, and representations can be updated, say, yearly while remaining useful. Others vary rapidly, rendering old representations useless.[3] Such rich detail often falls into the category "local knowledge," "the commonly held understanding of the customs and information shared by a group of people with a shared interest in a particular space" [5], and is therefore not broadly available. If we wish to represent search spaces relevant to problems motivated by interests in sustainability, we must ask, first, which problems to support solution of and which search spaces to represent; second, how to collect the data; and finally, how to represent it appropriately. The first question cannot be answered in general; it is likely to be answered by particular interpretations of "sustainability" in particular contexts. (What is to be sustained?) Methods for collecting and representing the data may be as diverse as the problems.[4]

Mapping existing resources can yield shifts in the distribution of resource use. Further, some representations can cause people to perceive things as resources that they would not otherwise have perceived as such. Consider a ladder that sits in Federico's garage. In a neighborhood without existing sharing practices, the ladder's main function is to take up space. It occasionally helps Federico put up and take down holiday lighting. In a neighborhood with robust sharing practices, Federico's neighbors can use his ladder while he isn't using it. Based on this approach, we are building a database-backed web application to support sharing of goods and services within a neighborhood. We are examining effects on resource use of a variety of design features (reputation system, interface design elements, privacy controls, etc.). This can be seen as an attempt to help people understand their assets as shareable resources, and illustrates an important development—the aggregation of small amounts of data contributed by many people, sometimes called "peer production" (see, e.g., [6])—in resource representation.[5] We conclude with a map of the space of projects of this type, listing existing exemplars and the lessons learned from them as well as open questions and unsupported problems.

References

[1] Runge, M. C. Challenges in framing the problem: just what are we trying to optimize anyway? *CompSust '09*. Slides at `cis.cornell.edu/svideo/ICS/CompSust09Slides/Runge.pdf`.

[2] Ferris, B., et al. OneBusAway. *Proc. CHI '10*: 1807-1816.

[3] Bonanni, L., et al. Small business applications of SourceMap. *Proc. CHI '10*: 937-946.

[4] Benkler, Y. Sharing nicely. *Yale Law J.* **114**: 273-358.

[5] Geertz, C. *Local Knowlege: Further Essays in Interpretive Anthropology*. Basic, 1983.

[6] Benkler, Y. Coase's penguin, or, Linux and *The Nature of the Firm*. *Yale Law J.* **112**: 369-446.

[7] Turner, A. OpenStreetMap Haiti. *opensource.com*. `opensource.com/life/10/1/openstreetmap-haiti`.

[8] Maron, M. Haiti OpenStreetMap response. *Brain Off*. `brainoff.com/weblog/2010/01/14/1518`.

[2] And of course, as anyone who has used these services knows, the representations they offer of the search spaces in question constrain the problems they can easily support solution of; there may be a tradeoff between the number of problems such a representation can support solution of and their utility in solving any particular problem.

[3] Consider carpooling organized around a regular workday commute: times and places of departure and arrival change infrequently in such a situation. A one-time cross-country moving trip, however, might produce similar surplus capacity, but may require a more frequently updated representation for this capacity to be used. For more on carpooling, see [4].

[4] For example, how would you go about asking restaurants and grocery stores where their food came from? Packaged foods in grocery stores may have information on their labeling that divulges some of the desired information, but collecting this information from restaurants would require the cooperation of proprietors, which might or might not be forthcoming.

[5] Notably, OpenStreetMap owes their definitive map of Port-au-Prince, Haiti, to this approach (see e.g., [7], [8]).