Exploiting the Environment: Urban Navigation as a Case Study

Nicholas Kushmerick
Department of Computer Science & Engineering University of Washington Seattle
nick@cs.washington.edu

The Situated Action approach to AI emphasizes the role of the environment in the generation and control of behavior; see (Norman 1993) for an introduction. Work to date has focused mainly on activity within spatially and temporally localized environments such as kitchens and video games (Agre & Chapman 1987; Agre & Horswill 1992). How useful is this perspective when larger-scale activities are considered? I attempt to answer this question by considering some issues related to navigation in urban environments. I identify several constraints on the structure of street grids that make navigation much easier than arbitrary graph search. The ultimate goal is a theory of the relationship between features of an urban environment and the computational complexity of navigation. This work extends the sort of analysis advocated by (Agre & Horswill 1992; Horswill 1993).

Navigation is an attractive task for studying the ways agents might exploit the structure of their environment. There is no doubt that people make and use elaborate mental representations when they navigate. But city street grids are constrained in ways that make navigational problems relatively simple, and these constraints are poorly understood. The constraints take a variety of forms, ranging from the physical structure of space to cultural phenomena such as neighborhoods. • Street grids are *physically stable*. New buildings and streets are constructed, but the time scale at which street grids change is several orders of magnitude slower than the scale at which navigation occurs. • Navigation is much simpler than arbitrary graph search because streets are *topologically sensible*. It is impossible to drive along a street and suddenly end up on the other side of town; *culs-de-sac* are relatively rare, so hill-climbing tends to work; one-way streets never completely isolate regions. • Navigation occurs in a *topographically translucent* environment: some information is available by virtue of the 3-D nature of the environment, although not all. Often one can see more than just the immediate surroundings. Deciding which highway exit to take can involve simply looking at the buildings in the distance to decide when the correct exit is approaching. • Some cities are *coherently labeled*. Streets might be numbered, alphabetized or follow some other regular pattern. • Most cities are *informatively labeled*. Downtown Seattle is filled with signs guiding one to Interstate 5; dead-end streets have signs so indicating; freeway exits indicate the places the exit serves; fast-food restaurant billboards guide one to the nearest franchise. • Most street grids *facilitate optimization*. Arterials are easily recognizable and uniformly distributed. Near-optimal navigation is thus simplified because a simple policy of using the nearest arterial is easy and effective. • Finally, cities are composed of *neighborhoods*, with just a few major streets running through each. It is often sufficient to identify a goal location by neighborhood; navigation within the neighborhood can then be done with a combination of visual and exhaustive physical search.

I take navigation to be physical search over a highly constrained graph. Environmental features map directly to constraints on the graph, thereby illuminating each feature's computational significance. Thus, physical stability means that the graph is static, while topological sensibility requires that the graph be strongly connected and nearly planar. Street signs and topographic constraints are especially interesting; they are treated as node labels informing the agent about remote parts of the graph. Further research will clarify the relationship between the environmental and graph constraints. Ultimately, I seek a formal theory of the relationship between the complexity of navigation and graph constraints derived from the environment. The following cases illustrate that the complexity of the agent and the environment are in some sense equivalent: an agent with a complete street map can navigate in an arbitrarily complicated city; an agent with a compass but no map can navigate if all streets are numbered. A theory describing these tradeoffs will help elucidate the general principles governing interactions between agents and environments in a wider variety of circumstances. To validate my analysis I have begun to formally model a complete street map of Seattle.

References

Agre, P., and Chapman, D. 1987. Pengi: An implementation of a theory of activity. In *Proc. 6th Nat. Conf. on A.I.*

Agre, P., and Horswill, I. 1992. Cultural support for improvisation. In *Proc. 10th Nat. Conf. on A.I.*

Horswill, I. 1993. Analysis of adaptation and environment. *Artificial Intelligence*. To appear.

Norman, editor, D. 1993. Special issue on Situated Action. *Cognitive Science* 17(1).