# Environmental Knowledge Representation

Antony Galton and James Hood*

February 2004

Environmental Knowledge Representation is concerned with representing and reasoning about the surrounding spaces, or environments, within which intelligent agents go about their businesses. A key facet of this is an understanding of qualitative features of the landscape. To understand a landscape is to have command of a structured model in which important features are highlighted and their mutual interrelations made available for inference; in constructing such a model we must acknowledge that our cognitive engagement with large-scale environments differs in important ways from the understanding of 'table-top' spaces which has been prevalent in AI. We advocate an approach in which knowledge about objects (e.g., landmarks and paths) in the landscape is supplemented by knowledge of spread-out characteristics of the terrain, represented by means of spatial fluents or qualitative fields. Using knowledge of the landscape in this way, we focus on two areas in particular: qualitative spatial interpolation (a form of non-monotonic spatial reasoning); and qualitative representation of routes over the landscape.

The need for spatial interpolation arises from the fact that an agent's knowledge of a landscape is inevitably incomplete. Typically, the knowledge an agent has of various spatial fields will be more detailed at landmark points (or places), less detailed along paths, and considerably less detailed still at the areas between paths and places. But on the basis of what *is* known, an agent can make plausible inferences about what is not yet known. Such inferences are almost invariably defeasible and hence considered non-monotonic.

In recent work, we were concerned with the interpolation of fields that are qualitative, as opposed to quantitative. Interpolation of quantitative fields over continuous space is the classic problem of spatial statistics, and there already exist a number of well-established methods for attacking this problem – kriging, for example. In contrast to quantitative spatial fields, qualitative spatial fields are discrete, and how we handle them depends on how the different values they can take are related to one another. The simplest assumption is that there are no constraints on mutual adjacency for any pair of values, i.e., that for any two values $v$ and $v'$ which the field can take, it is possible for a region over which the field takes the value $v$ to be adjacent to a region over which it takes the value $v'$. For example, for a land-cover field with value-set $\{trees, grass, sand, water\}$ we may reasonably assume that any two of these types can be mutually adjacent. We have examined the problem of spatial interpolation both for qualitative spatial fields of this type, and for qualitative spatial fields where mutual adjacency is constrained. The latter incoroporated Galton's previous work on adjacency spaces [1].

Following on from the interpolation of qualitative spatial fields, current work is looking at how the knowledge extracted from this process can be integrated into environmental models

*School of Engineering, Computer Science, and Mathematics, University of Exeter, UK

used for travelling, or exploring, the landscape. Such models are usually based upon graph theory, which can provide an abstract representation of paths, places and the relations between them. Through labelling the edges of a graph (representing the paths), we can represent the distances between places (represented as nodes). Graph theory, though, cannot describe the areas between paths and places, nor can it readily describe the variations which can occur along a single path segment. This limits its utility as a tool for representing an environment; how, for example, could one use graph theory to describe a path that cuts through a forest, or a path that goes up a hill?

To answer this question, consider a network of paths across a landscape. From the traveller's point of view, it is important to know the variation of field-like properties such as gradient or surface quality along the paths. If we confine our attention to a particular path, then the sequence of gradients (considered qualitatively as, say, steep down, gentle down, level, gentle up, steep up) encountered along the path forms a *profile* which provides a useful overview of the route to be followed. Kulik and Egenhofer [2] have developed a formal language for describing such profiles, but in the context of interpreting the visual horizon. By assigning profiles to all the component stretches of a network of paths one could build up a qualitative overview of the variation in elevation across the landscape in so far as it affects the user of the paths. Exactly the same procedure could be applied to other field characters such as the surface quality, and an agent's knowledge of a landscape may consist largely of elements of such profiles, obtained in the course of actually traversing the paths. Part of our current work is investigating the development of a formal system that incorporates these profiles.

# References

[1] A. P. Galton. Continuous motion in discrete space. In A. G. Cohn, F. Giunchiglia, and B. Selman, editors, *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000)*, pages 26–37, San Francisco, California, 2000. Morgan Kaufmann.

[2] L. Kulik and M. J. Egenhofer. Linearized terrain: Languages for silhouette representation. In W. Kuhn, M. Worboys, and S. Timpf, editors, *Spatial Information Theory: Foundations of Geographic Information Science*, Lecture Notes in Computer Science, pages 118–135, Berlin, 2003. Springer-Verlag. Proceedings of COSIT 2003.