Qualitative Interpolation for Environmental Knowledge Representation

Antony Galton and James Hood[1]

1 Introduction

Environmental Knowledge Representation is concerned with representing and reasoning about the environments within which intelligent agents operate. To understand an environment is to have command of a structured model in which important features are highlighted and their mutual interrelations made available for inference. Such knowledge is largely *qualitative*: to understand an environment it is neither necessary nor sufficient to have detailed knowledge of the values of numerical variables at precisely specified spatial locations.

We are often concerned with *spread out characteristics* of the environment: built-up areas, grassland, hill country, flooded areas, forest, uneven ground, road surfaces, etc. In the context of AI, such knowledge may be represented using *spatial fluents*, i.e., functions from locations to values. A simple example [2] is the presence of a certain disease; this is represented by a proposition-like entity a such that, in a given model, a holds at some locations, and $\neg a$ at others.

Environmental knowledge is inevitably incomplete, and inferences from what *is* known to what is not are generally defeasible, e.g., *interpolation* and *extrapolation*, both forms of non-monotonic reasoning. The main regulating principles for such reasoning in the spatial domain are *continuity* and *persistence*, applying to continuous and discrete phenomena respectively. The general problem arises when we know the values of a fluent at certain points and we want to infer values at other points. For continuous fluents over continuous space, well-researched statistical techniques exist. But for *qualitative* environment understanding, where the fluents are typically discrete, other methods are needed; such methods are the subject of this paper.

2 Interpolation for fully discrete spatial fluents

A fully discrete spatial fluent f maps spatial locations into a set $V = \{v_1, \ldots, v_n\}$ of qualitative values, where mutual adjacency is unconstrained, i.e., for any $v_i, v_j \in V$, a region over which $f = v_i$ may be adjacent to one over which $f = v_j$. For non-monotonic reasoning over such a fluent, we use *spatial persistence* [2, 1, 3], as follows. Suppose we know $f(l) = v$, and let l' be some other location. Then unless we know of some barrier to propagation of value v between l and l', we assume $f(l') = v$ by default. For an example of such a barrier [2], let f be the presence or absence of some disease spread by non-swimming animals. Suppose we know the disease exists at l, but not whether it is present at l'. If there is a river between l and l', the default inference (that it is present at l') will be blocked.

As just described, spatial persistence leads to conflicts. Suppose we know that $f(l) = v$ and $f(l') = v' \neq v$, but the value of f at l'' is unknown. If there are no barriers, then applying persistence to both v and v' implies, absurdly, that f takes both values at l''. To prevent this, we must refine the spatial persistence procedure. One method, suggested by [2], is the following 'projection' method: assign to l whichever value of f holds at the nearest location to l at which its value is known.

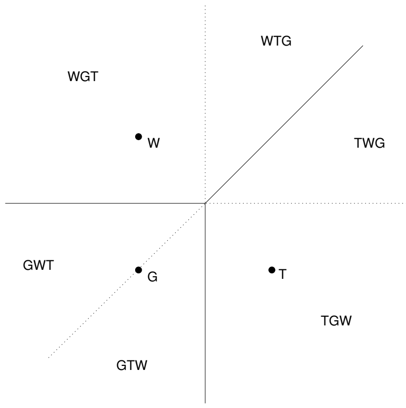

Figure 1. Voronoi diagram for three points arising from persistence-by-projection procedure.

This procedure assumes that different values of f are equally probable, which may be unrealistic. Shanahan [4] considers *spatial occupancy* in the context of reasoning about moving objects. He introduces a default principle that unless there is reason to believe otherwise, a point in space can be assumed to be unoccupied. The fluent in question has values 'occupied' and 'unoccupied', and Shanahan's principle is tantamount to privileging the latter over the former. Spatial persistence, as described above, is inappropriate here.

We modify simple projection as follows: to privilege value v over v' with factor $k > 1$, apply projection as before, except that when comparing distances from points at which $f = v$ with distances from points at which $f = v'$, first multiply the former by k. In $\mathbb{R}^2$, suppose $f(l) = v$ and $f(l') = v'$. With the unmodified principle, the dividing line between points assigned value v and points assigned value v' is the perpendicular bisector ll'. With the modified projection principle, it is the locus of points whose distance from l is k times their distance from l', i.e., a circle which encloses l' but excludes l. Inside the circle we assume $f = v'$, and everywhere else, $f = v$. The effect is to make v more 'likely' than v', by an amount dependent on k.

As a simple illustration, consider figures 1 and 2. In each figure the three points labelled G, T, and W, stand for Grass, Trees, and Water respectively. In Fig. 1, we apply the simple projection principle. The

[1] School of Engineering, Computer Science, and Mathematics, University of Exeter, UK. James Hood's attendance at ECAI is supported by a travel grant from ECCAI.

space is divided into sectors according to the relative distances of the three landmarks—the sector labelled TGW consists of those points for which the nearest landmark is T, the second nearest G, and the furthest W, and so on. The simple projection principle assigns to the whole of each sector the fluent that comes first in its label, so that WGT and WTG together constitute the W sector, consisting of those points for which the nearest landmark is W. The W, G, and T sectors together form the Voronoi diagram for the three landmarks.

In Fig. 2 we employ instead the modified projection principle. With the same three landmarks, we have applied a weighting of 3:2:1 to G, T, and W respectively. The smaller circle containing landmark W consists of the points exactly three times further from G than from W. Points inside this circle are assigned value W in preference to G. Similarly, the larger circle containing W consists of points exactly twice as far from T as from W, and hence the points inside the circle would be assigned W in preference to T—except that, outside the smaller circle, this preference is overridden by the preference for G over W. The four areas into which the space has been divided by the circles are labelled GWT, WGT, GTW, and TGW in accordance with the preference orders for the points they contain. Finally the value T is assigned to the points in the lighter shaded circle, W to those in the darker shaded circle, and G everywhere else.

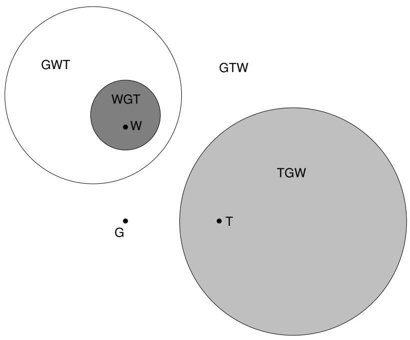

Figure 2. Modified projection procedure.

Such projection procedures provide a means for provisionally determining the character of a place from what is already known. With the modified procedure, results will vary with the fluent weightings. These could be empirically tuned to maximise agreement with real environments. Such tuning will also depend on other terrain properties: e.g., standing water is less extensive on sloping than on level ground and will accordingly be more heavily weighted against in the former context.

3 Interpolation in quasi-continuous spatial fluents

Some discrete-valued fluents are *quasi-continuous*, i.e., an adjacency relation is defined on values so that adjacent regions must have equal or adjacent values. A simple example arises if a continuous real-number value such as elevation is discretised by applying a 'banding' procedure, e.g., below 200m is 'Low', above 1000m is 'High', and other values are 'Medium'. If the variation in elevation is truly continuous, then a 'Low' area cannot be immediately adjacent to a 'High' area, whereas both types can be adjacent to a 'Medium' area.

For a fluent of this kind, the projection procedures used above can lead to 'illegal' adjacencies, e.g., a 'High' region abutting a 'Low' region. Persistence must be supplemented by a continuity principle to ensure that illegal adjacencies do not occur. The leftmost diagram in Fig. 3 illustrates an interpolation problem for qualitative elevation. Five points are given at which the qualitative elevation (H, M, or L) is known. The problem is to infer a plausible distribution of qualitative elevations for other locations in the map.

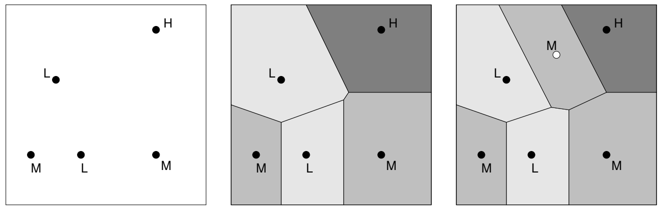

Figure 3. Interpolation of qualitative elevation

Simple projection yields the central illustration in Fig. 3. The L, M, and H areas are shaded light, medium and dark respectively. One of the 'Low' areas is adjacent to the 'High' area, violating continuity. In this case, where there is a single intermediate value that can be interpolated, we introduce a new 'Medium' point on the boundary between the adjacent 'High' and 'Low' areas, as shown, and then apply projection to the new set of points, (right hand illustration). This is consistent with continuity, and may be regarded as 'plausible'—but note that the procedure has 'decided' that there is a 'Low' valley separating the 'Medium' ground to the south-west from the 'Medium' and 'High' ground in the east, rather than, say, a 'Medium' ridge extending SW–NE separating the two 'Low' areas on either side.

The right-hand diagram in Fig. 3 may be regarded as expressing a set of inferences about the lie of the land in areas where it is not known, on the basis of a small number of places where it is known. These inferences are defeasible since acquisition of further knowledge may contradict them, requiring a fresh interpolation to be carried out. The true state of affairs may even violate continuity: real landscapes can exhibit discontinuities in the variables used to describe them. Continuity is a default assumption, which is reasonable because in real situations discontinuities tend to be isolated from one another, separated by areas in which continuity holds sway.

We could extend to this kind of fluent the modified projection procedure using weightings to privilege some fluent values over others. In a mainly low-lying landscape with sporadic peaks and ridges, we might privilege lower elevations over higher, e.g., by weighting L, M, H in the ratio 5:4:3. In a high plateau landscape incised by tongues of lower-lying land, these ratios might be reversed. Empirical studies could 'fine-tune' the ratios for best fit to different kinds of landscape.

REFERENCES

[1] Hans W. Guesgen and Joachim Hertzberg, 'Spatial persistence', *Applied Intelligence*, **6**, 11–28, (1996).

[2] Jérôme Lang and Nicholas Asher, 'A nonmonotonic approach to spatial persistence based on distances and boundaries', in *Time, Space and Movement: Meaning and Knowledge in the Sensible World*, eds., Pascal Amsili, Mario Borillo, and Laure Vieu, pp. 101–112. Groupe LRC Toulouse, (1995). Workshop of the 5th International Workshop TSM'95.

[3] Jérôme Lang and Philippe Muller, 'Plausible reasoning from spatial observations', in *Proceedings of the Seventeenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-01)*, eds., D. Koller and J. Breese, pp. 285–292. Morgan Kaufman, (2001).

[4] Murray Shanahan, 'Default reasoning about spatial occupancy', *Artificial Intelligence*, **74**(1), 147–163, (1995).