Maps Beyond Diagrams: Real-World Spatial Orientation

Clare Davies

Research & Innovation
Ordnance Survey
Romsey Road
Southampton SO16 4GU

United Kingdom

clare.davies@ordnancesurvey.co.uk

Abstract

Cognitive models of orientation using aerial-view 2D diagrams such as maps could help us to predict where their contents and design help, or hinder, cognitive processes, resulting in more usable map designs for situations where orientation is a common problem. However, the complex realities of geographic space, and the large scale transformations required with maps, make it less easy to model even basic processes such as orientation. This paper sets the context for ongoing work at Ordnance Survey, Great Britain's national mapping agency, to attempt to apply and extend cognitive modeling in order to enable (and ultimately perhaps to automate) decisions about geographic information provision for assisting real-world orientation tasks.

Introduction

A very common event , for example when traveling, occurs where we find ourselves viewing a real-world scene - either in an unfamiliar place, or by viewing a photograph or other image - and comparing this scene with a map or other information representation.

It is often the case in this situation that we already know approximately *where* we are on the map. We may already know the name of the location - perhaps it is supplied with the photograph, e.g. in a guidebook or on a website - or we may have just emerged from a subway or rail station, or alighted from a bus or tram. We may even have walked there ourselves, without needing to refer to a map until this point. In any case, the next question is "Which way am I - or is the camera - facing?" In this paper, this task of matching the reality to the map will be referred to as the orientation task.

Why This Problem?

Although in many cases the map user in this situation may be trying to navigate through the landscape, this problem does not only arise during wayfinding. It may also occur when:

o   trying to identify a specific building or object which is not explicitly labeled on the map, e.g. to visit or study it, or in an emergency scenario;

o   trying to match a historic image (e.g. of an old street scene) to a modern-day map;

o   trying to match a historic map to a modern-day scene or image;

o   making planning decisions based partly on viewing the current visual landscape;

o   trying to orientate with respect to cardinal directions in the absence of local cues, e.g. to establish sun and wind direction;

o   trying to understand the natural features of a landscape, with respect to the 'bigger picture' given in a regional map;

o   trying to decide on relative distances and directions to unseen distant locations (whether or not one intends to navigate to them).

Therefore, from urban and rural planning to risk management scenarios and recreational tourism, people need to be able to match 2D and 3D representations.

As a national mapping agency, and along with the rest of the geographic information industry, Ordnance Survey is increasingly moving towards a future where 2.5D and 3D representations of the world around us will be commonplace. So should we or anyone worry anymore about conventional 2D cartographic design? We believe that the orientation problem outlined above will still be relevant into the future because:

o   2D mapping is inherently more portable (e.g. it can be printed, scanned and copied more easily).

o   The aerial view of a 2D map gives the 'bigger picture' more easily (a whole area can be viewed at once).

o   Limitations are continuing on the feasible computing power of portable devices, and 2D mapping is a considerably less computationally demanding graphics format than 3D.

o   Technologies giving access to 3D will not be affordable or usable by every member of society for the foreseeable future.

o   Navigation within 3D immersed virtual environments presents its own problems of disorientation and even 'simulator sickness', at least for some users.

o   Part of the power of a map lies in the way it simplifies the enormous complexity of the real environment. Using symbolic conventions and ignoring a great deal of visual and physical detail, a map tries to distil only what really matters to its user. Even with 3D attributes and the opportunity to include realistic imagery (such as draped or rendered images of building fronts), the semiotics of map symbols will probably still have a place (e.g. MacEachren, 1995).

It should be added that any potential mapping provider in the 21st century is likely to be becoming less concerned with the ideal visualization that a map might have at a given scale, and more with the content that it must have in order to 'work' for a given purpose. Mapping agencies like Ordnance Survey provide most of our mapping as digital vector data for use in geographic information systems (GIS), where it is up to the GIS user to determine the final visual design.

Obviously, however, the content and design of any information visualization cannot be separated, particularly where some of that design is reliant on symbolism. In addition, national mapping agencies are often still tasked with producing paper mapping, and raster imagery of it, so there is obvious interest in improving these.

Factors in real-world orientation with maps

In the literature on diagrammatic reasoning, maps are frequently mentioned as a class of diagram. For example, some research has examined how maps, like other diagrams, can supplement text (e.g. Kulhavy et al., 1993), while other studies have considered issues such as recall for aspects of previously-viewed maps (e.g. Montello et al., 1994).

Reasoning with diagrams and with maps often tends to be researched on the assumption that the representation is not being matched in real time with any real-world scene or object(s), but is having to 'stand alone' as an aid to a thinker or learner. Meanwhile, where geographic-scale spatial cognition is considered, it is often at the level of someone trying to learn or navigate around large areas that could not all be seen at one glance except from an aerial perspective (e.g. Rossano et al., 1999; Cornell & Heth, 2000; Hendricks et al., 2003).

This focus on movement, memory or both can easily overlook one of the most common situations in which map users have to employ spatial reasoning: the translation and rotation from the simplified, symbolized, aerial-perspective, map to the large scale scene in which they are immersed. This cannot always be solved using the extrinsic (e.g. cardinal directions) frame of reference that is often assumed for cognition of geographic space (Barkowsky, 2002), although such a reference frame may be used more in some cultures and groups than others (Lawton, 2001; Davies & Pederson, 2001). Whatever their preferred frame of reference when given a choice, in many situations people have to deal with immediate space where intrinsic, deictic and relative reference frames will often be available *without* the extrinsic, and yet they must compare it with a representation which may only include the extrinsic and relative. Therefore it is probably the relative reference frame which helps users the most in this situation: e.g. by observing the angle between two adjoining roads, or distinguishing between two farm buildings by their distance from the nearby wood.

Is this just a trivial translation and rotation problem? Arguably, it can be far more complex than that. It is reasonable to state that any diagram, by definition, always simplifies the information present in the real world. In the case of a map, the simplification may be quite drastic for simple reasons of scale: the cartographer must employ *generalization* (e.g. see also Regnauld, 2001; Barkowsky, 2002) as well as using symbolic forms to reduce some object identities to mere category tokens.

Particularly with a relatively small-scale map (e.g. 1:50,000), many objects that surround the user in the real-world scene are simply omitted from the map altogether. Roads with many bends and twists are reduced to a line with just a few representative kinks. For a user standing at a road intersection, this may mean that even the geometry of the junction itself is not accurately reflected in the map, although the 'gist' will be similar. Meanwhile separate buildings or trees may be merged into a generic shaded block, and street furniture and fences or walls may be totally absent from the map.

Therefore, the user must search for something that can be uniquely matched between the map and the scene. In the terminology of one previous paper on map use (Berendt et al., 1998), they must create an appropriate 'abstract map' that picks out the features of the mapped space that might be matchable to reality and hence suggest specific solutions to the orientation problem. If the geometry (the relative orientations and distances of features) is very regular, such as in many grid-pattern US cities, then there may not be much to go on, unless the map represents some local landmarks. Sometimes there may not be anything distinguishable other than street names, or a distant landmark such as a hill (if relief is represented on the map at all).

To make the task harder yet, unlike the scene-viewing tasks of many cognitive science experiments, the user is immersed in the scene and is small relative to most of its major objects, so the scene is spread not only around 360 degrees of rotation but also some way up the vertical axis (less on a prairie, but more in a city).

Overall, then, the impact of scale (of both the map and the environment) on the usability of a map for orientation is a factor whose effects are more profound than perhaps they would be among other diagrammatic representations.

What Do We Already Know?

Cognitive mapping research has already established that the orientation task is harder if the map is not aligned to the direction that the person (or camera) is facing (e.g. Aretz, 1989; Werner & Long, 2003). Some people are more willing and skilled at mental rotation than others; others seem to find the simultaneous change of perspective and scale too overwhelming to attempt; yet others solve the problem physically by rotating the map in their hands. Nevertheless, all could probably be helped by increasing the clarity of the correspondence between the map and the visible reality.

We also know that people appear to use the configuration of visible objects such as streets, buildings, trees or hills to help them solve the task. Clearly, the extent to which the different visible objects are distinct from each other, both in reality and on the map, either in their placement or shape or both, will affect the ease or difficulty of the task.

A closely related area to this one is the work on wayfinding, and the potential use of landmarks to help enhance route directions and other aids. Experimental and computational modelling work (e.g. Burnett et al., 2001; Raubal & Winter, 2002; Winter, 2003), has identified certain factors that make a landmark more visually salient compared to its surroundings, and hence more quickly recognisable and usable as a cue for wayfinding decisions. The outputs of such research could obviously enhance any kind of representation intended to help the traveller, even if it took the form of a 2D map within which perspective images of the salient landmarks were included at road intersections (or were symbolised in some way).

However, the cognitive requirements of wayfinding, and the benefits of landmarks in enhancing it, are distinct from those of the 'static' orientation task. Wayfinding landmarks are generally assumed to act mainly as a 'cue for action', to warn the wayfinder that it is time to take the next turn, and this turn is assumed to be known in unambiguous terms e.g. "at the big grey church, turn left". The wayfinder is also assumed to be moving along a predetermined path to a specific destination, without having particular interest in the features or position of the current location.

It is not clear whether, once landmarks of the type envisaged are included in a representation, this would also be enough for orientation. It seems unlikely to be sufficient, because orientation is often needed at locations that are not wayfinding choice points. Also, what may be a good landmark for wayfinding may not be good enough for unambiguous orientation - e.g. a fountain in the center of a traffic island. A good landmark for orientation, therefore, would have a shape or character that was easy to match unambiguously to its aerial view or symbol on the map, but its visual appearance from ground level would not need to be particularly eye-catching (unlike a salient wayfinding landmark). Both, however, benefit from being unique.

Cognitive Modeling and Mapping Content

Recently Glenn Gunzelmann, a researcher at the US Air Force Research Laboratory in Arizona (Gunzelmann & Anderson, 2004), has built a cognitive model using the ACT-R architecture. This models human performance on one form of the orientation task. The task was simplified in being performed with a simple virtual environment - a flat circular world dotted with small conical landmarks - and a corresponding circular map on which each landmark appeared as a dot. Experiment participants had to match a highlighted landmark on the map with one of the 'real life' cones. The ACT-R model predicted, to a very high degree of accuracy, people's response times and error patterns where the following factors were manipulated:

1. Distance from observer to target
2. Number of landmarks in the same 'cluster' as the target
3. Alignment of map to 'real world' view

Problems of map alignment were raised earlier with respect to the real-world context. However, distance to 'target' and size of cluster are not so obviously relevant in normal, non-military, situations. This reflects a difference between Gunzelmann's task and the one described above: his task was to identify a specific entity in the landscape, and the only thing making this hard was the presence of several other identical items. In more familiar orientation situations, the task is not usually defined in terms of identifying a 'target', although it may well be if, say, the target is a specific road that one must turn onto.

Also, in Gunzelmann's virtual world the landmarks were small enough and the viewing angle high enough to be able to view all the items at once, even when directly occluding each other from a ground-level perspective. Therefore distance to target was relevant in his study in a way that may not still hold, if at least some of the objects in a geographic-scale visual scene were normal size and thus completely occluding each other (e.g. at an urban street corner, where only one line of buildings can be seen stretching away from each corner in each direction). Nevertheless, distance may well still play a part in orientation decisions.

Of course, this is not the first time that cognitive models have been considered for spatial rotation and/or orientation (e.g. Pellegrino & Hunt, 1991; Hermer & Spelke, 1994; Barkowsky et al., 2003). Much advanced AI work is also underway concerning related aspects of cognition in geographic space. The use of a cognitive modeling

approach based on a symbolic production system architecture may seem surprising to some, given the alternatives of connectionist modeling or agent-based simulations (e.g. Frank et al., 2001). However, there are several attractions of the production systems approach for our specific purposes. These include its initial simplicity, and its firm basis in actual human performance and in a psychological theory-grounded (albeit not universally accepted) cognitive architecture.

The grounding in such an architecture allows for the possibility of a series of models which all call on the same theoretical 'components' of the cognitive system, and which are particularly well suited to identify regularities in problem-solving strategies. These regularities may then permit prediction of changes in people's ability to orientate quickly if, say, landmarks were added to or removed from the map. Most important, therefore, is its potential (like other production-system models, e.g. Ritter et al., 2001) to be used specifically to evaluate and compare different design variations in real-world artifacts and tasks, in terms of their resulting impact on the human user. The performance of the model can be tested both quantitatively and qualitatively against results from experiments with human participants.

Clearly, for those of us involved in topographic mapping of more complex environments, work like Gunzelmann and Anderson's is only the beginning. Our own research group at Ordnance Survey is currently building a revised model based on Gunzelmann's, but incorporating some of the representational variables that matter in a more realistic world as discussed earlier.. However, this work still utilises virtual environments in testing (for now). They are easier to control for human experiments, but more immersive and hence realistic - and also potentially more interactive - than using photographs or video.

In our research, both the environments and the associated maps are derived directly from our national geographic feature database, OS MasterMap[1], which is among the most detailed and comprehensive national geographic datasets in the world. This includes not only every residential and business building in Great Britain, and the road, rail and water features, but also many other landscape features such as mailboxes, sidewalks, farm buildings, schools, cliffs, post offices, churches, country pubs, walls and fences, rural footpaths, woodland, marshes, farms and beaches.

Based on OS MasterMap data, it is becoming possible to construct highly realistic models of the British urban and rural landscape, via enhancement with 3D attributes such as building heights, shapes and frontages. Indeed, some computer games have already used this to provide extremely realistic backdrops to car chases around large areas of London. Ironically, in an echo of the arguments made earlier that some people continue to want 2D representations even with 3D visualisations and navigation systems, at least one review of one such game actually criticized it for not including a map (Smith, 2003):

> I really, really, really wish there were a map in this game. It wouldn't have to clutter up the in-game interface... Considering the amount of time that was spent on all the rest of this game, someone could have sat off in the corner putting together a real cracker of an in-game map, with indicators for the traffic directions and patterns, major landmarks pointed out... A map like that would be a brilliant aid to the player's appreciation of the virtual city.

3D models of real-world landscapes are also beginning to be used in many other applications, such as environmental research. Meanwhile the base topographic data continues to underpin most British geographic databases in urban planning, national and local government, utilities, retail, and many other industries and applications.

Of course it is also possible, using GIS, to produce customised 2D maps for specific purposes, directly from this data. Yet, for both these and traditional published maps, there is still no systematic model for extracting and appropriately depicting suitable features of map design so as to optimise the orientation task for map users working with ground-level outdoor reality (or with images of it). At smaller scales, where much information has to be removed or simplified as discussed earlier, knowledge of what users minimally need for orientation becomes even more useful. Unlike most artificial intelligence research in this field, the present work is aiming to model the (real human) user for the sake of being able to improve performance and hence the acceptability of mapping (particularly for people who struggle with the orientation task). Despite the many complexities, and the need to account for individual differences among map users at some point, it is hoped that a successful model would lead to automatic geographic feature extraction capability, so that a rich topographic dataset can be effectively reduced to a suitable set of content for use in generalizing and creating a map at a specific scale, that would in turn be usable for orientation at most locations where this was likely to be needed (e.g. intersections). The key distinguishing attribute of this work is not its modeling approach so much as its point of departure: trying to study and model the human user, rather than making broad heuristic assumptions about human needs.

---

[1] See http://www.ordnancesurvey.co.uk/oswebsite/

References

Aretz, A. 1989. Spatial Cognition and Navigation. In *Perspectives: Proceedings of the Human Factors Society 33rd annual meeting*, 8-12. Vol. 1. Santa Monica CA, USA: Human Factors Society.

Barkowsky, T. 2002. Chapter 2: State of the Art. In *Mental Representation and Processing of Geographic Knowledge: A Computational Approach*, (T. Barkowsky ed.), 21-64. Lecture Notes in Computer Science, Vol. 2541. Berlin: Springer-Verlag.

Barkowsky, T., Bertel, S., Engel, D., and Freksa, C. 2003. Design of an Architecture for Reasoning with Mental Images. Presented at: International Workshop on Spatial and Visual Components in Mental Reasoning about Large-Scale Spaces, September 2003, Bad Zwischenahn.

Berendt, B., Barkowsky, T., Freksa, C., and Kelter, S. 1998. Spatial Representation with Aspect Maps. In *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, (C. Freksa, et al. eds.), 313-336. Lecture Notes in Computer Science, Vol. 1404. Berlin: Springer-Verlag.

Burnett, G., Smith, D., and May, A. 2001. Supporting the Navigation Task: Characteristics of 'Good' Landmarks. In *Annual Conference of the Ergonomics Society: Contemporary Ergonomics 2001. Cirencester, UK, 10-12 April 2001*, 441-446. Hanson, MA: Taylor & Francis.

Cornell, E. H., and Heth, C. D. 2000. Route Learning and Wayfinding. In *Cognitive Mapping: Past, Present and Future*, (R. Kitchin, and S. Freundschuh eds.), 66-83. London and New York: Routledge.

Davies, C., and Pederson, E. 2001. Grid Patterns and Cultural Expectations in Urban Wayfinding. In *Spatial Information Theory: foundations of geographic information science. International Conference, COSIT 2001, Proceedings*, 400-414. Lecture Notes in Computer Science, Vol. Berlin: Springer.

Frank, A. U., Bittner, S., and Raubal, M. 2001. Spatial and Cognitive Simulation with Multi-Agent Systems. In *Spatial Information Theory: foundations of geographic information science. International Conference, COSIT 2001, Proceedings*, 124-139. Lecture Notes in Computer Science, Vol. Berlin: Springer-Verlag.

Gunzelmann, G., and Anderson, J. R. 2004. Spatial Orientation Using Map Displays: A Model of the Influence of Target Location. Presented at: CogSci2004: 26th Annual Meeting of the Cognitive Science Society, August 2004, Chicago, IL, USA.

Hendricks, M. D., Egenhofer, M. J., and Hornsby, K. 2003. Structuring a Wayfinder's Dynamic Space-Time Environment. In *Spatial Information Theory: foundations of geographic information science. International Conference, COSIT 2003, Proceedings*,

75-92. Lecture Notes in Computer Science, Vol. Berlin: Springer.

Hermer, L., and Spelke, E. S. 1994. A Geometric Process for Spatial Reorientation in Young Children. *Nature* 370: 57-59.

Kulhavy, R., Stock, W., Verdi, M., Rittschof, K., and Savenye, W. 1993. Why Maps Improve Memory for Text - the Influence of Structural Information on Working-Memory Operations. *European Journal of Cognitive Psychology* 5: 375-392.

Lawton, C. A. 2001. Gender and Regional Differences in Spatial Referents Used in Direction Giving. *Sex Roles* 44: 321-337.

Montello, D., Sullivan, C., and Pick, H. 1994. Recall Memory for Topographic Maps and Natural Terrain: Effects of Experience and Task Performance. *Cartographica* 31: 18-36.

Pellegrino, J. W., and Hunt, E. B. 1991. Cognitive Models for Understanding and Assessing Spatial Abilities. In *Intelligence: Reconceptualization and Measurement*, (H. A. H. Rowe ed.), 203-225. Mahwah NJ, USA, and London: Lawrence Erlbaum.

Raubal, M., and Winter, S. 2002. Enriching Wayfinding Instructions with Local Landmarks. In *Geographic Information Science: Second International Conference, GIScience 2002, Boulder, CO, USA, September 25-28, 2002*, 243-259. Lecture Notes in Computer Science, Vol. 2478. Berlin: Springer-Verlag.

Regnauld, N. 2001. Contextual Building Typification in Automated Map Generalization. *Algorithmica* 30: 312-333.

Ritter, F. E., Baxter, G. D., Jones, G., and Young, R. M. 2001. User Interface Evaluation: How Cognitive Models Can Help. In *Human-Computer Interaction in the New Millennium*, (J. Carroll ed.), 125-147. Reading MA, USA: Addison-Wesley.

Rossano, M. J., West, S. O., Robertson, T. J., Wayne, M. C., and Chase, R. B. 1999. The Acquisition of Route and Survey Knowledge from Computer Models. *Journal of Environmental Psychology* 19: 101-115.

Smith, D. 2003. The Getaway: A Beautiful Movie, but How Is It as a Game? Dated Jan 6, 2003. Retrieved Jan 29, 2004: http://ps2.ign.com/articles/381/381552p1.html

Werner, S., and Long, P. 2003. Cognition Meets Le Corbusier - Cognitive Principles of Architectural Design. In *Spatial Cognition III*, (C. Freksa ed.), 112-126. Lecture Notes in Computer Science, Vol. 2685. Berlin: Springer.

Winter, S. 2003. Route Adaptive Selection of Salient Features. In *International Conference, COSIT 2003, Spatial Information Theory: foundations of geographic information science. Proceedings. Kartause Ittingen, Switzerland, 24-28 September 2003*, 349-361. Lecture Notes in Computer Science, Vol. Berlin: Springer.