The Ingredients of an Exosomatic Cognitive Map: Isovists, Agents and Axial Lines?

Alasdair Turner

Bartlett School of Graduate Studies, UCL, Gower Street, London
a.turner@ucl.ac.uk

Abstract. There is some evidence that an axial map, as used in space syntax, may be related to an underlying cognitive map in humans. However, the axial map is derived strictly from the mathematical configuration of space rather than any property of people. Hence there is a question of how a person might have embedded such a map.
In this paper we report the results of several experiments which aim to improve the correlation between agent and pedestrian movement. We use a database of external occlusion points derived from isovists constructed throughout the system to provide a lookup table for agents to guide their movement. Since the table is external to the agents, we refer to the visual architecture as exosomatic. The results do improve on previous studies, but are still far from a good simulation of pedestrian movement.
However, there is a philosophically important outcome from the experiments. When the agents are tuned to best performance, their movement patterns correspond to the axial structure of the system. This can be shown to be a mathematical result of their movement strategy; that is, the manifestation of movement, or the 'memory' of an agent experiment, relates to the combination of the internal structure of the agent and its engagement with the environment in the form of an axial map.
There are two unresolved steps from the relationship between individual and environment to human cognition: one, it cannot be shown that people do actually use occlusion points for movement, and two, even if they were to, it cannot be shown that they would use the resultant axial structure for higher level navigation decisions. Nevertheless, our results do provide evidence for a link between the individual and the axial map through embodiment of an agent-environment system, and our theory provides a mechanism for a link between the embodied map and preconditions for cognitive structure, which may in turn provide a basis for the future research into the means by which space syntax may be related to spatial cognition.

1 Introduction

Space syntax [1] is the study of the relationship between space and society. One of the central tools of space syntax is the *axial map*. An axial map is an abstraction of the space between buildings (or rooms and corridors of a building) to straight lines drawn through it according to a formal algorithm [2, 3] (figure 1). It has been discovered that graph measures of the resulting network of lines correlate with aggregate pedestrian movement, particularly at the urban scale [4, 5, 6]. However, the means by which the axial map translates movement at the level of the individual to a measure of a graph has not been fully explained. Therefore, the aim of this paper is to explore the nature of the axial map through its relationship to the individual. This does not carry us directly from individual to graph measure, but it does provide a link from an individual in continuous space to an axial, or at least, axial-like, structure. To achieve this aim, we employ agent-based models to simulate pedestrian movement, and look at various parameters that affect their movement, before returning to the aim of understanding the axial structure.

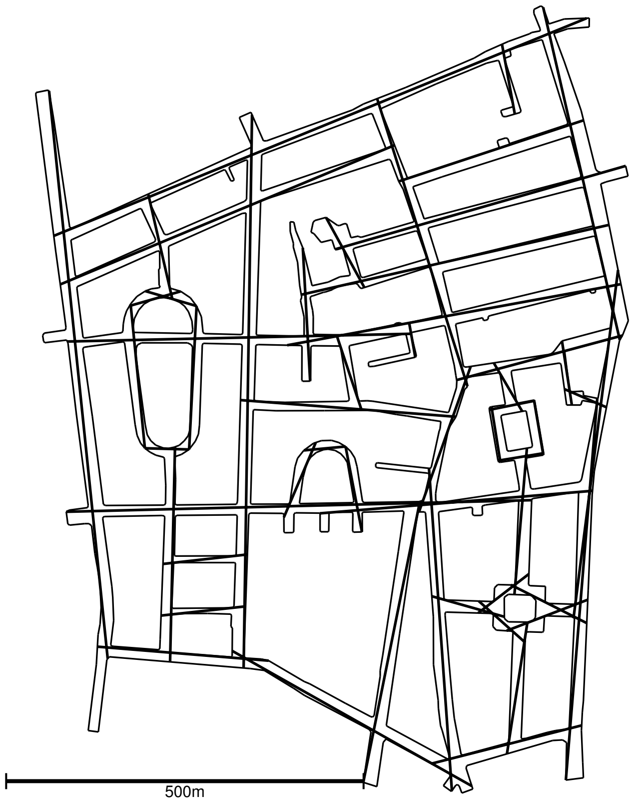

Fig. 1. An automatically generated axial map of Barnsbury in North London, within the open space of the area.

There is discussion about whether the movement correlation observed with axial maps relates to the internal structure of an individual, or whether it relates to the physcial constraints of the map.

Since the axial map definition derives from purely spatial constraints without reference to the occupants of a space it would seem natural to suggest that the overriding factor is the physical constraints of the space. The core premise is that placement of buildings or rooms leads directly to the movement observed. This is the theory of natural movement [5]. It is a theory that does not require an individual, rather the linkage is between society and space; the observed movement is the result of the aggregation of individuals, rather than any individual alone. It is this feature that allows the construction of a theory of the movement economy [7]: the spatial configuration self-perpetuates through the action of building shops in response to movement, which then leads to further movement, and further shops around it; that is, to Hillier, in his earlier work, the space is *the* 'machine', the mechanism in itself that leads to societal organisation.

However, there is another view, and that is that the axial line is a feature of the internal representation of space used by people[1]. It has been realised that the axial map is an extremely economic representation of space. Each line can cover many street segments, and requires, mathematically on a surface, just four variables (x, y location to x, y location or x, y location, direction and length). This economy, and the relation of axial lines to movement, has led to the suggestion that a cognitive map of the space might utilise similar features [8, 9]. This mirrors research in robotics by Kuipers that suggests that a skeletal framework just needs the notion of 'to the left of' or 'to the right of' in order to allow navigation from it [10]. Thus, Kim and Penn [11], investigate the relationship of axial maps to sketch maps as drawn by people. They find parallels between the graph importance of certain axial lines and the regularity with which the line is included in the sketch map. There is also a move to tie plausible cognitive functions to measures of the axial lines. For example, Dalton [12] has introduced angular analyses of axial maps in order to move towards experimental evidence of how people use minimal angular strategies to navigate from location to location [13]. The inclusion of these features has lead to improved correlation between measures and observed movement [14]. However, in all these studies, the existence of axial lines as a cognitive stucture on which to build the analyses is assumed. Thus, in this paper, we are interested in how the structure itself might arise, not through consideration of cognitive function, but through the engagement of agents with the environment.

In the next section, we will review the recent research in the direction of an embodied model of agent behaviour in the environment, including the relationship to active perception in robotics, situated cognition and further extensions to sensorimotor contingencies and phenomenology. We will then move on to practical experiments involving agent models, and show that there are deficiencies

[1] More evidence for this mode of interpretation exists than is presented here: in particular, the argument in the next section about eigenvectors can be used to show that if agents used strict natural vision, then movement could only be related to the connectivity of the axial line. As observed, movement is actually related to a second degree measure of each axial line, and thus must be due to some secondary effect: either people following people, or using memory in order to navigate.

in current models where large open spaces are concerned. Previous models have used first order, that is, direct, indicators of further configuration to guide their movement. In this paper, we apply a second order indicator, that of an occluding edge, so as to attempt to create a more 'human-like' pattern of movement from space to space. There is considerable experimentation to be done in this area, and we show the results of the application of several combinations of parameters to how the agents might respond to these environmental clues, both for the urban example of Barnsbury in North London (see figure 1), and the Tate Britain Gallery in London. It is discovered that the occlusion edge mechanism creates agents which correlate well with aggregate pedestrian movement in the urban environment. In addition, these agents create a pattern of lines through their movement, linking occlusion point to occlusion point. We show that these linkages reflect the mathematical mechanism used to create an axial map, and therefore that there is an innate association between the axial map and the embodied process of the agent movement in the environment. We suggest that this link provides evidence that an axial map is in fact the embodiment of movement in the environment, and thus inherently related to both the structure of the agent and the structure of the environment. However, we also demonstrate that this is not the whole story. We show that people in building environments *cannot* navigate by occlusion-edge alone if they are to produce the patterns of movement observed. There is a further connection with the environment not caught by these new agents, which is perhaps better caught by the original first order agents, or perhaps a feature of the environmental task not included in our models, such as, in the art gallery example used, a move to view paintings. Therefore, we are forced to conclude that axial models of building environments are also insufficient both representationally and practically. A rereading of the original space syntax literature relating to the Tate Britain Gallery shows that this restriction is implicitly acknowledged: the model used is *not* purely axial, but also includes convex areas (rooms), to equate movement to the model. The conclusion is that if a cognitive map represents the entire movement process, then it must be more complex than an axial map alone.

2 Background

An isovist is the visible polygon from a location in a plan of an urban or building environment, as shown in figure 2. Benedikt introduced the concept to architecture in order to try to quantify the perceptual experience of a place [15]. In particular, he recognised that the way people moved around a space might be influenced by the shape of the isovist, not simply by the objects within it. He supposed that people would be guided by isovist properties, following Gibson's suggestion that people may be guided by direct (or active) perception, that is, simply respond directly to the affordances offered by the environment rather than through any higher cognitive function [16].

Recently, with Penn, the author introduced simulation agents guided by the simple isovist property of visible area in the direction of movement [17]. We

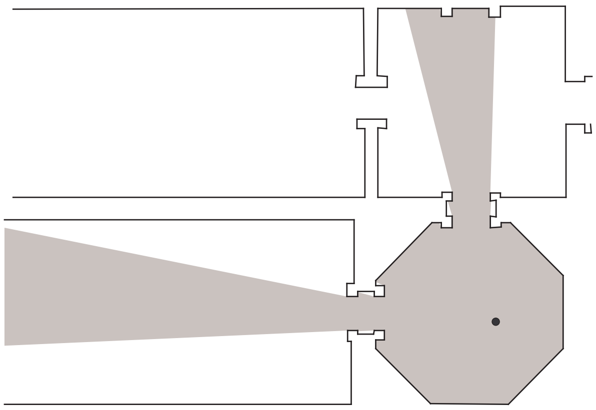

Fig. 2. An isovist generated from a location, or conversely, the isovist which converges on a location

constructed a dense grid visibility graph of a floorplan and allowed agents to choose their next destination from a selection of grid points in the agents' field of view. We discovered that if the agents are given a field of view of 170° and allowed to progress three steps before reselecting a destination, then aggregate movement of agents in a gallery space correlates well ($R^2 = 0.77$) with observed movement of people [18].

The power of this simple movement rule has led both Penn [9] and the author [19] to speculate as to the possible link between the agent movement and axial lines. Both of us point out that the information the agent sees is exosomatic: that is the agent samples the what Gibson calls the ambient optic array at any one location. The agent itself simply reacts to the environmental inputs at any time, producing a 'natural' outcome to stimuli. As such, the agents mimic those in Brook's subsumption architecture [20]. They have no explicit representation model, but the knowledge of the system (the natural outcome) is embedded in how the agents relate to a specific system. I have suggested that the system is essentially autopoietic [21] in nature: that there is an ongoing relationship between an organism and its environment [19]. Wheeler [22] calls such an ongoing relationship a 'hermeneutic dialogue'. He (Wheeler) suggests that this leads to a phenomenological understanding of the process; that is, that the experience of the agent can be accessed at the metalevel of the dialogue between agent and (to the author) the environment. Penn's analysis [9] is more clearly related to the notion of distributed or situation cognition [23, 24]. That is, that the knowledge exists out there, in the environment, and the agent is the receptor of that knowledge. In both the cases of Penn and myself, however, there is a desire to suggest that the axial map has some sort of relationship to the agents – either the result of an externalised memory or as a physically embodied system – which is not directly internal. O'Keefe, who proposed that the hippocampus may act as the cognitive map in humans [25], also points out that any brain structure must inherently relate to the physical structure, and that there is an essential intertwined relationship between representation and physical structure, be it in the brain or the environment [26]. It is an area that O'Keefe calls 'neurophilosophy' and it is echoed in the visual sensorimotor contigencies proposed by O'Regan and Noë [27]. In effect, O'Regan and Noë suggest that visual consciousness is only possible in relation to an environment and meaningless without it. The point is that there is a constant overlap between the representation in the brain and physical environment; where the information is stored is to a certain extent flexible, and features structural properties of both the environment and the individual. The emphasis herein is how structure might arise from movement and engagement with the environment.

At one level, Penn and my [17] agents are even more simple than we have already suggested. Since each agent only reacts to the information directly available to it at any one location, the probability of its next move at any time can be written as the initial distribution at each location (a vector for the inhabitants at each location) ω multipied by a Markov transition matrix, for example, M. That is $\omega_{t+1} = M\omega_t$. The probability of a move after two steps is given by multiplying the matrix together, $M \times M$ or M^2, and after n steps, M^n, so $\omega_{t+n} = M^n \omega_t$. An interesting property of the steady state of movement in the system is that $\omega_{t+1} = M\omega_t$. Thus, the steady state vector is simply an eigenvector of the matrix M. We might call the first eigenvector ω_0. It should be noted that $\omega_0 = M\omega_0 = M^\infty \omega_0$. That is, the steady state movement is simply an eigenvector of the standard transition matrix[2]. Now, the transition matrix for the agents is extremely complicated, but we can draw certain conclusions. Most importantly, that the matrix does not rely on the last direction chosen by an agent. That is, any directionality involved in the movement of the agents is inherent in the physical properties of the location at hand. This suggests that, so long is there is no direct or indirect feedback into the system (for example, agents that see other agents), the movement result is a local property that can be extracted from the environment. This is not as easy as it may sound given the matrix, and it is for the steady state limit, rather than agents walking a certain number of steps, but it does give a direct relationship between a property of the environment and the *internal* properties of the agent. Hence, although we cannot access the internal properties of a real agent (and we have failed to cater for learning) we can suggest that the static abstract output from an agent run is, to a certain degree, a manifestation of its internal structure. As we have discussed, both Penn and I independently suggested that the static abstract output we presented had a relationship to an axial structure. As I have suggested here, if this is the case, the axial structure is an inherent property of the internal structure of the agent and the external local properties of the environment when applied together. However, the earlier agents presented by Penn and myself have shortcomings which lead us to suggest that the properties they presented were

[2] This is a standard mathematical result, first raised with the author by Wagner [28]

not actually axial, although perhaps closely related. Herein, I show that with minor changes, those agents could embody axial structures.

3 Methodology

3.1 Practical Problems and Solutions

Application of Penn and Turner's agents [17] to larger scale spaces has proved less successful than our initial building scale experiments [18], perhaps due to a less controlled environment where entrance and exit are unconstrained, but also perhaps because the agents appear to congregate in larger spaces, as their direct perception leads then towards open areas. This can lead to a stark contrast between agents and observation where there is open space such as a park. For example, in the area of the South Bank in London by the London Eye, paths recorded for people followed through the space differ strongly from the patterns of agent movement, as shown in figure 3(a) and figure 3(b).

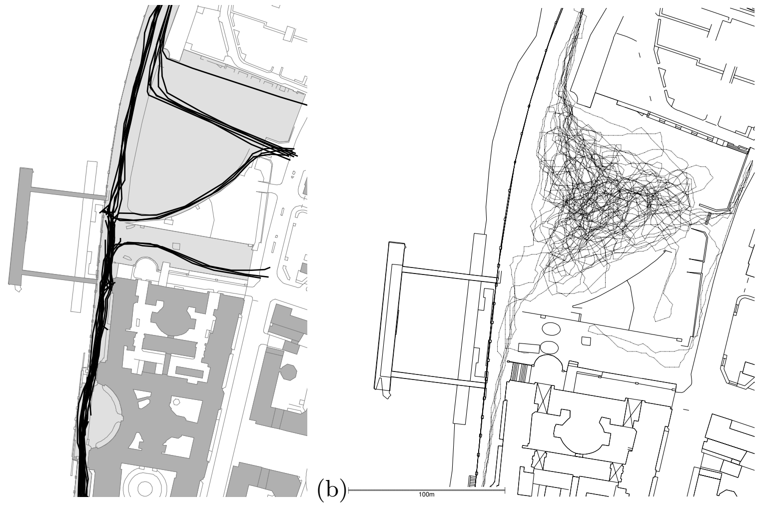

(a) (b)

Fig. 3. (a) Pedestrian movement through the South Bank, London. Observations and image by Christian Beros, used with kind permission. (b) Trails of standard direct-perception agents in the same environment.

One might suggest that perhaps this 'milling' in the centre of the space is a simple outcome of the rules employed – if agents were to continue to their destination, rather than simply taking three steps before selecting another destination (see the beginning of the last section), then surely they would walk straight across open areas? However, there are two downsides to using this strategy: firstly, that experimentally this was found to be poor by Turner and Penn

[18], but also that the agents may still choose destinations in the middle of the space, and the people observed in figure 3(a) clearly are not. Their selection is to an edge, away from the central space.

We might also posit other possible mechanisms for movement based on the properties of locations. Conroy [13] has observed that people slow and appear to make direction decisions in locations where the isovist properties are skewed, that is, the centre of mass of the isovist and the location where the isovist was generated are separated. This implies that some property of the isovist is considered before making a movement decision. However, just because people are observed making decisions at locations with high skew does not mean that all locations with high skew result in movement decisions. Furthermore, it does not suggest *how* people move from area of high skew to another similar area.

Of course, a simple wall-following procedure may lead to the pedestrian movement patterns shown in figure 3(a) (and indeed we have presented results of animat experiments elsewhere that do evolve such a strategy [29]). However, whether or not people do actually simply follow walls when navigating environments is difficult to tell from observation data alone: roads typically have pavements at their edges, and town squares often have street furniture at their centres (particularly in the UK where our data is retrieved), and elsewhere avoidance of sunshine will play a role in how people cross open squares. Therefore the exact strategy used is difficult to ascertain, although further research in this area required.

In this paper, we will consider another possible mechanism to allow crossing open areas as shown in figure 3(a). When Benedikt introduced isovists, he also proposed a measure of *occlusivity*. This measure indicates where isovists have long lengths of occluding radials, that is, a radial that marks a boundary between visible and occluded objects (figure 4). For navigational purposes these occluding radials might also be important, as they mark areas of unexplored space that may be entered by continuing in the direction of the occluding radial. A rule that guides the agent according to the locations of the occluding radial or occlusion points should cause agents to take direct routes across space, as seen in the South Bank example. Therefore, in this paper we examine the effect of agents guided by occluding radials.

There are significant memory benefits from using occlusion points to guide behaviour. Once an occlusion point has been identified, it is relatively easy to track, as it marks a discontinuity between optic flow – as one walks more background comes into view, or recedes, whereas the point itself remains constant, even if it moves across the retina. Thus, it is simple to record locations of occlusion points and discard them if necessary.

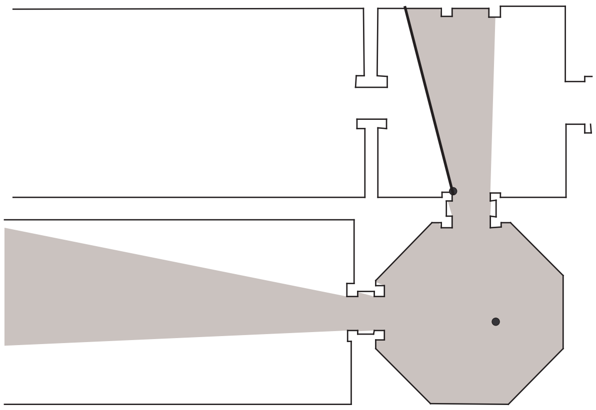

Fig. 4. An isovist showing an *occlusion point* (where the isovist meets an occluding edge), and an *occluding radial* (the portion of the isovist radial from the occluding point to its far end).

3.2 Implementation

For building environments, we construct a set of isovists covering a dense grid[3], placed every 0.75m throughout several test environments, as this gives an approximation to human step size. For urban environments, we follow [19], and use 3m grid spacing for ease of implementation. We break each isovist down into a series of 32 angular bins, and record the occluding radials in each bin. Simulation agents are then run in the system. Following [18], each agent has a field of view of 15 bins (about 170°), and moves for three steps[4] before making a new direction decision. As per [19] we simply transfer the 3-steps from building environment directly to the urban environment, regardless of scale implications.

Where we differ from our earlier agent models is that rather than selecting a point from the isovist to move towards, the agent chooses a direction based on the occluding points within its view. Several different approaches were attempted in order to try to understand how the occluding point movement decision could be best applied. In particular, there is question of resolution of the representation of the environment. If the resolution is too fine, then many occluding points occur at the edge of the system with every minor deviation. Therefore we ignore any occluding radial of length less than 1.5m (the occluding radial is considered just the open portion of the radial, as shown in figure 4). Obviously, a radial encountered at an oblique angle may still trigger the selection, but then this is probably reasonable: without further information, the agent is not to know

[3] The experiments in this paper were conducted using the Depthmap program, written by the author.

[4] In fact, this is actually implemented as a $\frac{1}{3}$ chance of changing direction in any one step, which does not equate to exactly a 'three-step' rule. It might better be conceived as a turn probability.

that this actually obscures a dead end. Once this simplification had been made, we then assayed several rules for choosing a new destination. In each case, the temporary destination was the occluding point itself. The experiments were as follows:

1. Control: standard direct perception agents as per [18].
2. Choose any occluding point at all from the available bins in the field of view.
3. For each bin in the field of view, choose the furthest occluding point, if any. Then choose any one of these at random.
4. For each group of 3 bins in the field of view (i.e., about 45° sections), choose the furthest occluding point, if any. Then choose any one of these at random.
5. For each group of 5 bins in the field of view (i.e., about 60° sections), choose the furthest occluding point, if any. Then choose any one of these at random.
6. For each bin in the field of view, choose the furthest occluding point, if any. Then choose one of these weighted according to how far away it is.
7. For each bin in the field of view, choose the furthest occluding point, if any. Then choose one of these weighted according to its angular deviation from the current course (the more, the higher the probability).
8. For each bin in the field of view, choose the furthest occluding point, if any. Then choose one of these weighted according to how far away it is multiplied by the angular deviation from the current course.

Experimental set up 1 was simply a control set up based on previous agent implmentations. Set up 2 was to test the affect of solely switching to occlusion points rather than any point within the system as a next step. Set up 3 was to try to cut down the number of occlusion points, so that directions with many occlusions did not bias directions with few, but important occlusions exist (for example to mitigate against situations where many columns exist). Experimental set ups 4 and 5 were to see if any advantage was obtained by simplifying the bin system further. Experiment 6 was used to try to influence the agents to follow better potential movement lines, rather than being distracted by elements close to the current location. Experiment 7 takes a different line to the same idea: rather than having the agent select a physically far location, to go for an angularly separated location. Finally, experimental set up 8 merely combines set ups 6 and 7.

4 Results and Discussion

For each experimental set up, a quantitative assessment was made against two test cases for which we have pedestrian data: a small model of an urban area (Barnsbury in North London, shown in figure 1) and a large public building space, the Tate Britain Gallery (see ahead to figure 7). For the Tate, agents were released from the entrance allowed to take 1800 steps in the system (as per [18]). Counts of the numbers of agents moving through the rooms of the gallery were then compared with data collected for people moving through the gallery (obtained from [30]) for 54 rooms on the ground floor of the gallery. Similarly, for the area of Barnsbury, we compare aggregate gate counts of pedestrians for the whole day published by Penn and Dalton [31] with gate counts of agents at 106 gates. We apply a fairly arbitrary 1000 steps in the Barnsbury case, and release the agents from any location. In addition to these two quantitative experiments, we also looked at the qualitative output for the South Bank area, to see how the agents performed against the observed pattern when released from entrance points to the system (as the original observations of people were conducted). The quantitative results are shown in table 1, and the qualitative results in figure 5. Table 1 shows the linear regression R^2 correlation coefficient between the log of agent counts and the log of room or gate counts. The log is applied in order to distribute the data roughly according to a normal distribution.

Table 1. Correlation coefficients for agent versus observed movement for the Barnsbury and Tate Britain Gallery experiments.

Experiment	Barnsbury R^2 ($n = 106$)	Tate R^2 ($n = 54$)
1	0.55	0.76
2	0.60	0.59
3	0.42	0.55
4	0.24	0.45
5	0.28	0.51
6	0.47	0.57
7	0.67	0.60
8	0.68	0.34

From figure 5 it is apparent that none of the methods catch exactly the type of movement observed in people (refer back to figure 3(a)). In all cases, apart from perhaps experiment 8 (where occlusion point choice is weighted by distance and angle), the direct North to South path is lost. Even in experiment 8, the agents seem to decide to change direction directly across the open space, presumably due to the decision rule suddenly selecting a different direction after three steps. Certainly all do represent some improvement on the standard agent case. With the quantitative study the results are more mixed. Experiments 4 and 5 (where the bins are grouped) appear to do particularly badly relative to all the others. In all cases, there is a drop when the agents are placed in the Tate, with experiment 8 faring particularly badly, in contrast both to the qualitative results and Barnsbury, where it appears to fare best. Also doing well on Barnsbury is experiment 2 (a random choice of any occlusion point within the field of view). This is perhaps easy to explain, due to the fact that main streets will tend to have more intersections, and hence more occlusion points along their length, but the same explanation does not work with experiments 7 and 8, where angle of turn plays a role. It simply seems that these two experiments align better

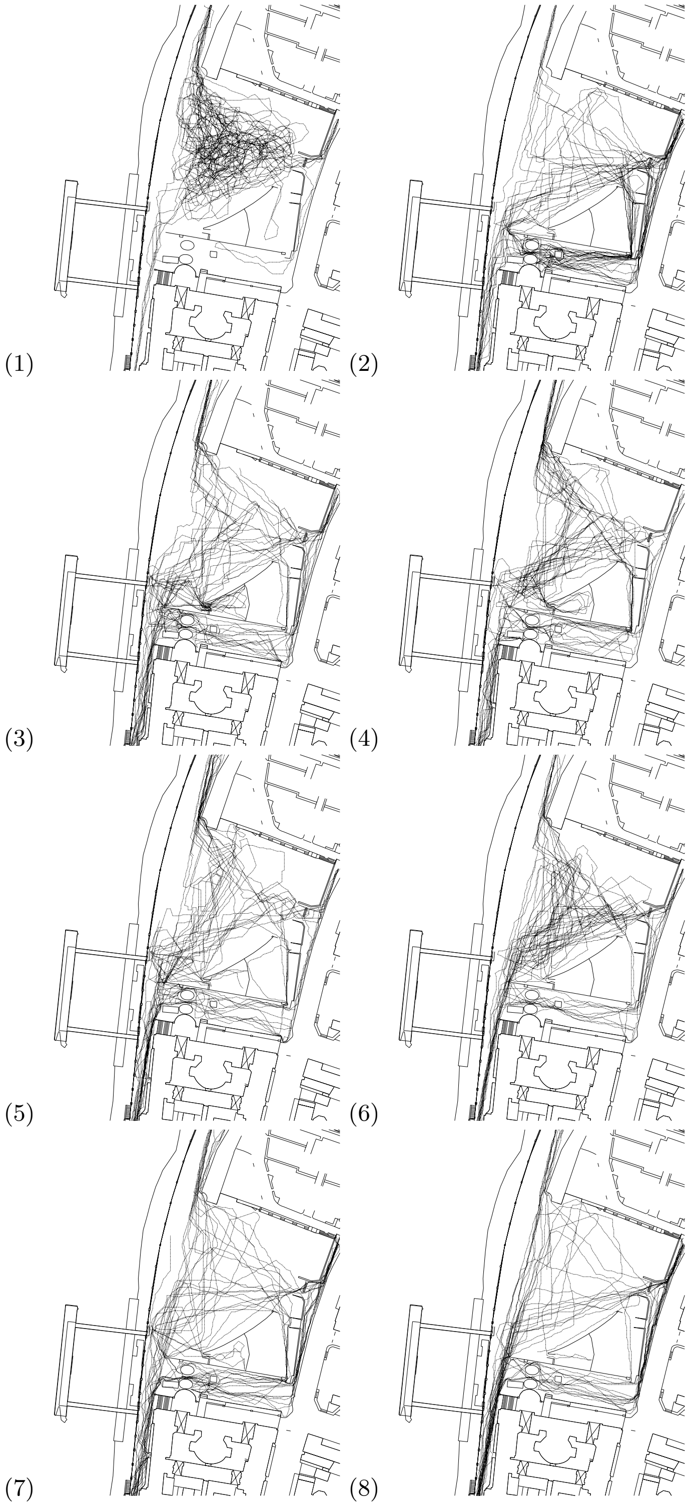

Fig. 5. Results from each of the 8 experimental set ups when applied to the South Bank area in London.

with the way people move through the system. They correlate with pedestrian movement considerably better than standard agents in Barnsbury, but worse in the Tate gallery; so what is going on? Figure 6 shows cummulative movement

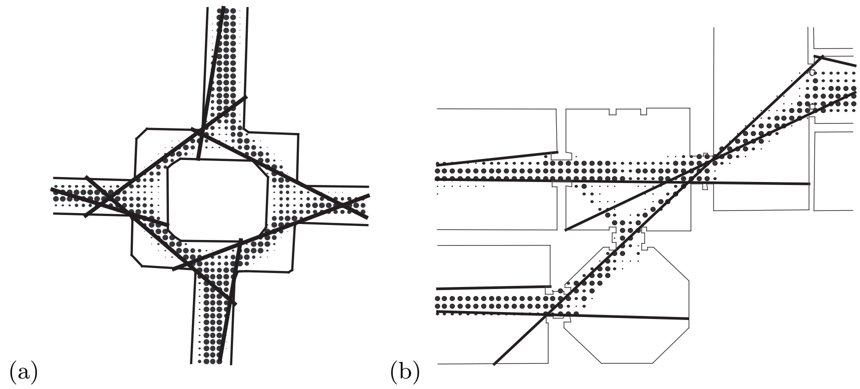

Fig. 6. Cummulative numbers of experiment 8 occlusion-driven agents (larger dots represent greater numbers of agents) overlaid on axial maps (bold lines) in (a) a section of Barnsbury and (b) part of the Tate Britain Gallery.

patterns of the experimental set up 8 agents as compared to the axial lines for the same areas. It appears that the movement of the agents aligns itself to a great degree with the axial lines. The same is true of the whole system, but difficult to appreciate without zooming into the detail. We should not be surprised that the agents do align with axial lines, as their movement is defined in a very similar way to the way in which axial lines are drawn. One of the rules for drawing axial lines is to join two reflex corners, and extend them. For the agents, reflex corners are occlusion points (corners around which the agents cannot immediately see). Therefore, if an agent moves by choosing one occlusion point, then moving to the next, and so on, then it is actually joining reflex corners akin to the axial map, regardless of the manner in which it chooses those occlusion points. The fact that the later methods depend on selecting the furthest occlusion point reinforces another of the byproducts of axial mapping: that longer lines are prioritised. So the situation leads to the agents sampling the axial lines of the system. In the Barnsbury case, this obviously leads to movement along the lines of actual pedestrians. However, in the Tate it appears to be flawed compared to the standard direct perception agents. We can show why this is the case with another pair of figures.

Figure 7 shows details of agent trails for previous agent implementations (figure 7(a)) and occlusion driven agents (figure 7(b)). It is obvious that the occlusion driven agents are drawn towards the top right-hand corner of the gallery. This is because there are more axial lines in the top right-hand corner.

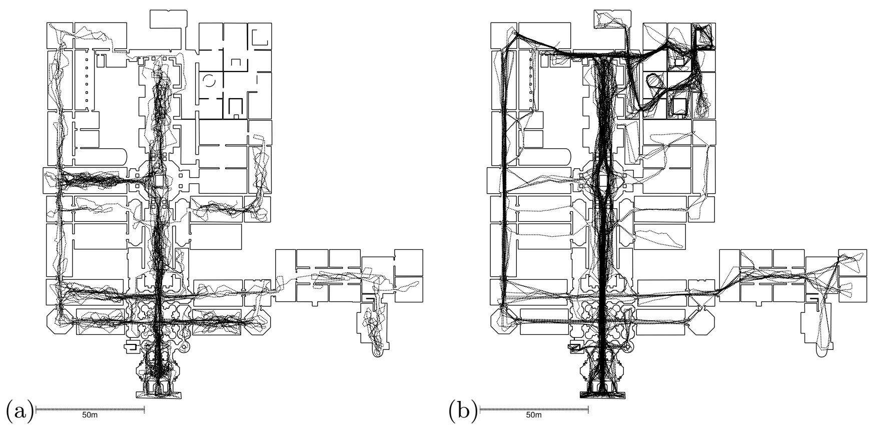

Fig. 7. Trails from (a) experiment 1 standard agents and (b) experiment 8 occlusion-driven agents in the Tate Britian Gallery.

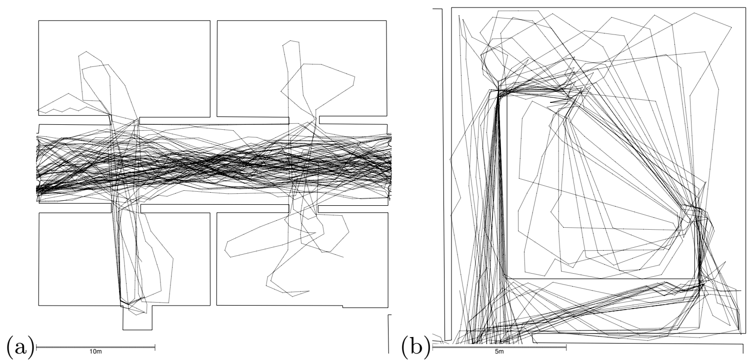

Fig. 8. Detail of trails from experiment 8 occlusion-driven agents in the Tate Gallery. (a) Rooms without internal configuration give no incentive to enter (note how a little extra configuration aids the room in the bottom left). (b) Rooms with internal configuration can over-incentivise entrance.

Figure 8 shows details of agent trails in two areas of the Tate gallery. The rooms in figure 8(a) are actually fairly well visited, but very few agents enter, precisely because there is no configuration within the rooms. The agents simply 'hover' at the entrances and leave, as the lack of further occlusion points provides no incentive to enter. It might be argued that allowing the agents to continue along the occluding radial may lead to their entering, in much the same way as axial lines are extended into spaces. However, this seems to lack the elegance of the occlusion point drive, and misses the point of why people may enter a room: if the room were truly empty then it is indeed likely that people would pass it by. We hypothesise that it is because these rooms have paintings they are entered. In any case, there is a second problem that cannot be extended by drawing axial lines. The room in figure 8(b), from the top right-hand corner, is barely entered in reality. However, the internal configuration of a temporary exhibition leads the agents into it (modelled with axial lines, the same would happen). This would probably correspond to the wishes of the exhibition designers, but in fact, the space (darkened to play film), attracts barely no-one. If these agents do so badly (in fact it is not 'so badly' for experimental set up 7, but comparatively badly), then how did the original space syntax analysis of the Tate Britain Gallery compare? The answer is, for a standard axial model, somewhat badly. In [30], the best model of pedestrian behaviour was in fact a combined model, that took axial lines but weighted them with convex spaces (essentially, the open room spaces). This is interesting, as the convex spaces, might be thought of as open space weights: exactly the attractors we were trying to exclude from the original first-order direct perception agents. It is no surprise then, that the original agents actually perform better than the occlusion-point driven agents in the Tate gallery. It also leads us to the conclusion that occlusion-driven is not the end of the story, and in fact, neither are axial lines. People may be behaving according to the axial structure of the city, but they are not behaving according to the axial structure of the building.

5 Conclusion

This paper has sketched out how agents may be programmed to use occlusion points to guide their movement. Qualitatively, we have shown that, in open areas, the paths generated by these agents correspond more closely to observed pedestrian movement than agents driven by direct perception, although there is still considerable room for improvement. Furthermore, quantitatively, the occlusion driven agents were shown to correlate well with observed movement in an urban system (up to $R^2 = 0.68$), but this was at the cost of previously found good correlation with movement of visitors to an art gallery space.

It was noted that the occluding points at the start of occluding radials are invariant as the agents move from location to location. These reflex vertices in the plan might be considered as the basis for a line map, joining the edges tranversed by the agents. Such a map has a very similar definition to an axial map from the domain of space syntax [1]. Comparison of axial maps with the cummulative traces left by our agents show very similar patterns.

It was argued that the trails left by agents represent some externalised cognitive understanding of the system; in fact, with reference to O'Keefe [26] and O'Regan and Noë [27], it was claimed that cognitive understanding could only arise from the combination of the agent and the system. Further, the mathematical point was made that for any given movement rule, there is a direct mapping of the steady state movement pattern with the transition rules of agent movement. Therefore it was argued that internal property of the agent (the movement rule) is exhibited as a direct manifestation in the environment. The implication is that the cognitive function of the agent also relates to the environment in this way, and must then pick up elements from the movement-environment 'co-map' in order to direct its movement. In the occlusion point driven agent, the instransient features are the occlusion points themselves, that lead to the axial map. Therefore any further cognitive function of the agent ought to rely on a map isomorphic in some way to the axial map.

As we have said, this is not the whole story. The occlusion agents miss something of the function of human pedestrian agents. However, we suggested that extra features are just that: extra features where there is no configuration to drive the agents further. If this is the case, then the basis of cognitive movement may currently be lodged in the environment, as an exosomatic axial map.

6 Acknowledgements

With many thanks to Christian Beros for his permission to use the South Bank data presented in this paper. This paper expands on the short paper "Isovists, Occlusions and the Exosomatic Visual Architecture" to be presented at The Cognitive Approach to Modeling Environments Workshop, GIScience 2006.

The Depthmap program used for this paper can be obtained from Space Syntax Limited, and is free for academic use. Please see `http://www.vr.ucl.ac.uk/depthmap` for details.

References

[1] Hillier, B., Hanson, J.: The Social Logic of Space. Cambridge University Press, Cambridge, UK (1984)

[2] Peponis, J., Wineman, J., Bafna, S., Rashid, M., Kim, S.H.: On the generation of linear representations of spatial configuration. Environment and Planning B: Planning and Design **25** (1998) 559–576

[3] Turner, A., Penn, A., Hillier, B.: An algorithmic definition of the axial map. Environment and Planning B: Planning and Design **32** (2005) 425–444

[4] Peponis, J., Hadjinikolaov, E., Livieratos, C., Fatouros, D.A.: The spatial core of urban culture. Ekistics **56** (1989) 43–55

[5] Hillier, B., Penn, A., Hanson, J., Grajewski, T., Xu, J.: Natural movement: or configuration and attraction in urban pedestrian movement. Environment and Planning B: Planning and Design **20** (1993) 29–66

[6] Read, S.: Space syntax and the Dutch city. Environment and Planning B: Planning and Design **26** (1999) 251–264
[7] Hillier, B.: Space is the Machine. Cambridge University Press, Cambridge, UK (1996)
[8] Jiang, B.: A space syntax approach to spatial cognition in urban environments. In Hirtle, S., MacEachren, A.M., eds.: Cognitive Models of Dynamic Geographic Phenomena and Their Representations, Pittsburgh, PA (1998) 48–58
[9] Penn, A.: Space syntax and spatial cognition: or why the axial line? Environment and Behavior **35** (2003) 30–65
[10] Kuipers, B., Tecuci, D., Stankiewicz, B.: The skeleton in the cognitive map: A computational and empirical exploration. Environment and Behavior **35** (2003) 80–106
[11] Kim, Y.O., Penn, A.: Linking the spatial syntax of cognitive maps to the spatial sytnax of the environment. Environment and Behavior **36** (2004) 483–504
[12] Dalton, N.: Fractional configurational analysis and a solution to the Manhattan problem. In Peponis, J., Wineman, J., Bafna, S., eds.: Proceedings of the 3rd International Symposium on Space Syntax, Atlanta, Georgia, Georgia Institute of Technology (2001) 26.1–26.13
[13] Conroy, R.A.: Spatial Navigation in Immersive Virtual Environments. PhD thesis, Bartlett School of Graduate Studies, UCL, London (2001)
[14] Hillier, B., Iida, S.: Network effects and psychological effects: A theory of urban movement. In van Nes, A., ed.: Proceedings of the 5th International Symposium on Space Syntax. Volume I., Delft, Netherlands, TU Delft (2005) 553–564
[15] Benedikt, M.L.: To take hold of space: Isovists and isovist fields. Environment and Planning B: Planning and Design **6** (1979) 47–65
[16] Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston, MA (1979)
[17] Penn, A., Turner, A.: Space syntax based agent models. In Schreckenberg, M., Sharma, S., eds.: Pedestrian and Evacuation Dynamics. Springer-Verlag, Heidelberg, Germany (2002) 99–114
[18] Turner, A., Penn, A.: Encoding natural movement as an agent-based system: an investigation into human pedestrian behaviour in the built environment. Environment and Planning B: Planning and Design **29** (2002) 473–490
[19] Turner, A.: Analysing the visual dynamics of spatial morphology. Environment and Planning B: Planning and Design **30** (2003) 657–676
[20] Brooks, R.A.: Intelligence without representation. Artificial Intelligence **47** (1991) 139–159
[21] Maturana, H.R., Varela, F.J.: The Tree of Knowledge: The Biological Roots of Human Understanding. Shambhala Publications, Boston, MA (1987)
[22] Wheeler, M.: From robots to Rothko: The bringing forth of worlds. In Boden, M.A., ed.: The Philosophy of Artificial Life. Oxford University Press, Oxford, UK (1996) 209–236
[23] Hutchins, E.: Cognition in the Wild. MIT Press, Cambridge, MA (1995)
[24] Clancy, W.J.: Situated Cognition. Cambridge University Press, Cambridge, UK (1997)
[25] O'Keefe, J., Nadel, L.: The Hippocampus as a Cognitive Map. Clarendon, Oxford, UK (1978)
[26] O'Keefe, J.: Kant and the sea-horse: An essay in the neurophilosophy of space. In Eilan, N., McCarthy, R., Brewer, B., eds.: Spatial Representation. Blackwell, Oxford (1993) 43–64

[27] O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences **24** (2001) 939–1031

[28] Wagner, R.: How to predict movements of random agents. Personal Communication (2003)

[29] Turner, A.: First steps towards evolution of pedestrian animats. In: Models of Machine Learning in Creative Design Workshop, DCC '06, Eindhoven, Netherlands, TU Eindhoven (2006) Forthcoming

[30] Hillier, B., Major, M.D., Desyllas, J., Karimi, K., Campos, B., Stoner, T.: Tate Gallery, Millbank: A study of the existing layout and new masterplan proposal. Technical report, Bartlett School of Graduate Studies, UCL, London (1996)

[31] Penn, A., Dalton, N.: The architecture of society: Stochastic simulation of urban movement. In Gilbert, N., Doran, J., eds.: Simulating Societies: The Computer Simulation of Social Phenomena. UCL Press, London (1994) 85–125