

Evolving Direct Perception Models of Human Behavior in Building Systems^{*}

Alasdair Turner and Alan Penn

Bartlett School of Graduate Studies,
University College London, London WC1E 6BT, UK,
a.turner@ucl.ac.uk, a.penn@ucl.ac.uk
WWW home page: <http://www.vr.ucl.ac.uk/people/>

Abstract. Software agents that use direct (or active) perception of the environment have recently been shown to correspond well with pedestrian movement within building and urban systems. The algorithm, based on Gibson's theory of affordances, combines random selection of destination from their field of view with reassessment of the destination every few steps. However, although the agents correlate with human movement on aggregate, as individuals they progress more erratically than people do. It might seem necessary to add higher cognitive functions in order to guide them more convincingly, but here we show that it is possible to improve their behavioral response through artificial evolution of their existing navigation rules. First we show that the destination selection method approximates stochastic direction choice by length of line-of-sight. Then we use the lines of sight to provide a set of inputs to the agents, or animats, which we evolve to fit human usage patterns within a building as best possible. We demonstrate that while agents using informational change inputs fail to evolve to fit movement patterns, an input that compares sight-line lengths improves models qualitatively, but not quantitatively, which further implies that the individual guidance mechanism may be independent of the inherent spatial properties acting on direct perception.

1 Introduction

When we think of people as they navigate through the environment, it is usual to regard them as rational beings, making complex decisions based not only on the information available to them at the current time, but also on the history of their actions, and their future goals. Thus, Gibson's theory of ecological perception [1], where the occupant of the environment interacts with it directly, through the affordances it offers her or him, is usually applied to theory of insect behavior

^{*} Turner, A. and Penn, A., 2007, Evolving direct perception models of human behavior in building systems. In Waldau, N., Gattermann, P., Knoflacher, H. and Schrekenberg, M. (Eds) *Pedestrian and Evacuation Dynamics 2005* (Springer: Berlin) pp. 411-422
ISBN 978-3-540-47062-5

rather than pedestrians. When it comes to people, we tend to think of how their navigation relates to a remembered view or an internal cognitive map, or at least to the cost of the activity in which they are involved. That is, if we consider a four level hierarchical framework for movement, as proposed by Trullier et al [2], with guidance at the bottom, followed by recognition-triggered response, topological navigation and metric navigation, we tend to assume that the observed actions of people — or at least their behaviorally interesting actions within geographical space — take place on the higher three levels. Guidance is usually seen as merely a device to avoid collisions and other obstacles. Therefore, while there is considerable research into how navigation is affected by place or view recognition [3, 4], spatial topology [5, 6], or the perceived metric features of the environment [7, 8], there is little to suggest how the rules of guidance may actually lead to observed ‘navigational’ behavior.

We can regard people as simple affordance-based agents. Gibson states: ‘When no constraints are put on the visual system, we look around, walk up to something interesting and move around it so as to see it from all sides, and go from one vista to another. That is natural vision...’ p. 1 [1]. To him, people simply take whatever affordance is apparent and act accordingly: perhaps to an attractive object, or simply from vista to vista. Gibson proposes that the vista itself is simply an ambient optic array of light incident on the current occupant’s location. As we move through the optic array, so we cause perturbations within it, and experience optic flow, until we find ourselves at a new vista. The transfer from vista to vista we might call *natural movement*, as it requires no recognition of any object within it, but simply the recognition that there is environment to move through [9]. In the natural movement schema, the possibilities of the environment drive the agent through the ambient optic array from vista to vista as more possibilities open up ahead of it (see figure 1). Thus, the configuration is the determining factor of the resulting natural movement, even if different agents take different paths through it as individuals.

However, if we are to construct agents, we must consider *how* they recognize these environmental possibilities for movement. There is good evidence that insects use the optic flow itself to perceive the environment: honeybee dances seem to indicate the optic flow on a journey from A to B rather than landmark or other cues [10]. As the bee flies through a narrow gap, there will be fast flow of information to its sides. Humans possess radial receptors in their eyes, and so it might be thought that they too use optic flow to perceive the environment. However, the situation is more complicated: where stroke victims lose the use of the receptors, they still manage to navigate easily [11]. Further, experiments using children show they tend to use geometric as opposed to texture information to guide decisions — that is, other forms of depth-cue are being used to calculate action through basic mental manipulation [12].

From this evidence, we might assume that somehow the brain is able to construct some form of distance information to objects such as walls. While strictly speaking we should force the agent to work with the visual input directly, the use of distance data lends itself to an elegant and simple formalism, that of




Fig. 1. As an agent moves through the environment, further vistas open up ahead of it, and it can guide itself through a configuration simply by choosing directions of further possibility, for example, by length of line-of-sight.

the *isovist*. An isovist is a polygon (or polyhedron) that represents the visible area from a location. Benedikt introduces their use to architectural configurations [13], and in particular, the idea that there exist in the environment fields of isovist properties. Thus, Benedikt generates contours of equal isovist properties (for example, area or circumference) in order to try to identify areas where rapidly changing distance information takes place. He hypothesises that such zones will form decision points for onward movement. Conroy Dalton shows that people do indeed seem to slow and make decisions at points where isovist properties change rapidly [14], and thus these informationally rich areas may be important to natural movement.

Here we attempt to test the hypothesis that changes in isovist field or distance information may be a plausible strategy to guide people through a building system. To do so, we create agents that respond to changes in the local length of line-of-sight (LoS) information by turning towards that direction if a threshold is passed (so the information is used as an affordance within a direct perception model). We compare the ability of the agents to respond to the environment in a human-like manner when given one of four different informational inputs. These are: simple length of line-of-sight in the candidate direction, the change in length of LoS between this step and the last, the length of LoS in the candidate direction as compared to the length of LoS in the current heading, and, finally, a similar comparative measure for the amount of LoS change between steps. The idea behind the two variants of LoS change inputs is that they should correspond to the agents receiving an information flow, in a similar manner to optic flow. However, we should be careful to note that the change in length, the equivalent

of the forward-facing field change between two locations in Benedikt’s scheme, is not the same as the difference in ambient optic array between two locations. While the forward-facing response of a differential length model is similar to optic flow — as the target is approached there is greater optic flow and change in distance at the center of the field of view — for the radial case, high optic flow is recorded when walls are close but at a constant distance from the occupant; by contrast, the differential distance remains constant, and only changes when we cross a threshold, so it much more resembles the sort of field changes expected in Benedikt’s model than optic flow itself.

We apply these threshold agents to an art gallery situation, after original experiments using similar agents [9]. Of course, the choice of art gallery is no accident, as it represents a situation where natural movement may well occur: most people will not know the gallery plan and thus will not use higher level navigational strategies, other than perhaps path integration; however, it also acts as a large area with many potential turnings, suitable for testing the agents. At first the agents are given random thresholds to turn, but these are then adjusted by using an evolutionary algorithm, in order to see if they can be evolved to fit the movement pattern of the people using the gallery. The current test is obviously limited, in that it essentially evolves the model to fit a single case. We do find, though, that while agents with the comparative line-of-sight rule can attain a correlation coefficient of $R^2 = 0.84$ with observed aggregate movement pattern, those with differential field inputs only achieve at best a correlation of $R^2 = 0.56$, which indicates that information flow represents an unused navigational strategy.

We begin with a background to earlier experiments using affordance-based agents, and an explanation their limitations.

2 Background

The use direct perception as a guidance strategy for agents is not new, with its theoretical breakthrough in the introduction of Brooks’s robot subsumption architecture in 1991 [15]. Brooks’s robots have no knowledge representation, and therefore must interact with the environment only in the ways permitted by their sensor and motor architecture. Equally, the use of ambient optic arrays (albeit in sonar form) dates back to architectures such as those introduced by Kuipers and Byun [16], where the robot is given a tripartite system of low-level sensorimotor control, topological mapping and geometric path recording. However, within software simulation of people movement, direct perception and direct visual control have tended to give way to higher level models of path planning at the geographical scale. One reason is possibly the high overhead of the calculation of visual information, although as we will show herein, this can be partial and still lead to good results.

The introduction of an exosomatic visual architecture (EVA), where the visual information is contained in the environment, rather than calculated by an individual, leads to a significant reduction in time overhead [17]. The exosomatic visual architecture might be thought of as an ambient optic array, as visible loca-

tions waiting to be sampled from the occupant’s location. We first implemented the EVA as a visibility graph of many point locations sampled on a grid [18]. This allowed the introduction of many parallel agents that choose a destination from their current field-of-view (FoV). The algorithm we used was simply to pick any destination at random from those visible on the floor ahead, and walk towards it for a number of steps, before selecting another destination, walking towards it, and so on. We applied the agents to a plan of the Tate Britain Gallery (a large art gallery in London) and adjusted the FoV and number of steps until we found the best correlation with observed room through movement for people, which occurred at a 170° FoV and with a decision after three steps, giving a correlation coefficient of $R^2 = 0.77$ [9]. Although this is extremely encouraging, there are a number of caveats. Firstly, the three step parameter is arbitrary. Would three steps hold in any other building, let alone outside? And surely a good direction might be completely overlooked if the agent were on the first or second step? Secondly, the agent paths themselves were not convincing (see figure 3 further below), involving much backtracking and almost constant turning around. Thirdly, the correlation takes no account of the intercept of the best-fit line. The best fit line in fact does not pass through the origin (as it should if zero agents were to correspond to zero people). Here we try to remedy these three problems, although our major success is with the first rather than the second two.

In our model here we will lose some of the advantages of the the original affordance-based rule, by using line-of-sight data rather than the complete visibility graph. We should not forget that optic flow is not just happening to our sides but above and below us, and there is evidence to suggest that the angle below the horizon strongly influences distance perception [19]; therefore we may well be using the ground for clues about where we are going. Indeed, obstacle avoidance behavior typically takes place over short ranges [20], so we may well be looking down rather than ahead while navigating. However, for overall ease, rather than using the entire visibility graph, it would be appropriate to use isovist boundary information instead, and it allows us to test a range of inputs unavailable with previous method, including an easy assessment of informational flow. Furthermore, as we will show in the next section, line-of-sight and destination selection strategies correlate well with each other, to the extent that they must be approximating the same underlying rule.

3 Simple Line-of-Sight Agents

In this section, we demonstrate that a random selection of a destination within the field-of-view (FoV) is equivalent to the stochastic choice of direction based on the length of line-of-sight (LoS). In order to construct an agent able to sample the FoV, we assume that we have precalculated information both about visible destinations available from the current location, and for convenience, separated them into 32 ‘bins’, see figure 3(a). We also assume that we have calculated the length of LoS in the direction of each of the 32 bins, as shown in figure 3(b).

Both these calculations are performed using UCL’s Depthmap software¹. Using 15 bins to select a destination is thus equivalent to a FoV of about 170°.




Fig. 2. Using precalculated bins of information to select (a) a destination location, and (b) a line-of-sight from the current field of view.

We released 2000 of the standard destination selection agents into a plan of the Tate Britain Gallery, London, with bin information calculated every 0.75m. The choice of grid size is based on average step length for people [9, 21]. Each agent took three steps and then reselected its destination, and remained in the gallery for 1600 steps (about 20 minutes of simulation time). After this time it was simply removed, with no requirement to return to the entrance. Although programming return is possible, it posits an attractivity of the entrance based on a desire to leave rather than a strict direct perception model of visual affordances; therefore we leave it out at this stage. After the experiment had been run, the number of agents entering each of 55 rooms within the gallery was recorded, and compared with secondary observation data collected for a two day period within the actual gallery [22]. Obtaining a correlation is somewhat difficult, as much of the movement within the gallery can simply be accounted for as diffusion from the entrance. As we leave the entrance, numbers drop off rapidly, creating a skewed data set, with many more locations at lower numbers, meaning the correlation coefficient can look artificially high. One resolution would be to use Spearman Rank Correlation; however, this loses the ability to see exactly how the model is over or underperforming. Instead, we choose a function of the data. For these experiments, we realized that a near-normal distribution of the data could be found using the cube root of the actual through movement rate². The cube root was applied to both the observation data and agent count data. This simply

¹ See <http://www.vr.ucl.ac.uk/depthmap/> for details.

² This in turn is more satisfactory than the logarithmic scale used in earlier papers as some rooms have very low through movement rates, which throw the logarithmic scale off at the low end.

repeats our earlier experiments and gives a correlation coefficient of $R^2 = 0.79$ (see table 1)³.

Using this experiment as a control, we then programmed agents to choose a direction based on a weighted function of LoS. The probability to choose any of the 15 available bin directions was simply its length over the total length of the 15 available bins. That is, $p_i = L_i / \sum_i L_i$, where p_i is the probability of choosing direction i , L_i the length of LoS in direction i , and i is one of 15 bins distributed evenly either side of the current heading, as shown in figure 3. Surprisingly, the results are very similar for the LoS agents. So similar, that we compared the LoS agents with the destination selection agents, and found near one to one correlation⁴, as shown in table 1.

Table 1. Correlation coefficients (R^2) of destination-selection agents against line-of-sight agents and observed movement

Correlation	Destination selection		Line-of-sight	
	0-step	3-step	0-step	3-step
Observed Movement	0.76	0.79	0.73	0.78
Destination selection	0-step	–	0.93	0.88
	3-step	–	0.86	0.91
Line-of-sight	0-step		–	0.93
	3-step			–

Although the correlation with observed movement is good both for destination selection and LoS agents, when we compare movement traces of actual people moving through the Tate Britain Gallery and the agents, there is a marked difference. Figure 3 shows the trails of 20 3-step LoS agents released into the gallery for 10 simulated minutes, and figure 3 the corresponding patterns of observed people movement, collected by following people surreptitiously from the entrance [22]. It is noticeable that the paths of the agents are far less linear than the people’s paths. This turning would initially seem to be due to the wide FoV, and the fact that the agent may take stochastically pick a wide angle, combined with the agent having no memory of its original goal. However, as we will see, neither the FoV nor the lack of memory account for this change, rather, it appears to be the lack of a default onwards forwards movement.

There is a further problem with the rules of these basic agents: it is that the ‘three steps’ parameter is effectively an act of calibration. It seems unlikely that

³ The number is slightly different to earlier reported results due to the use of the cube root rather than logarithmic scale.

⁴ The reader may wonder if the correlation might improve if we were to use a function of the LoS length to weight the choice of direction, rather than simply the LoS length itself. The answer appears to be no: in experiments weighting the direction by the square root, square and cube of LoS length, all showed worse correlation than those presented here.




Fig. 3. Trails of 20 3-step LoS agents for the first 10 minutes of their ‘visit’ to the Tate Britain Gallery.

people really do count three steps and then make a decision. If we presuppose a neural model, it makes more sense that some threshold is passed, and the person turns accordingly, should she or he actually navigates through natural movement. However, if we trim back the number of steps taken in our model so the agent continually reassesses the situation every step (we call this a 0-step rule), the correlation with actual movement drops considerably (refer back to table 1). In addition, as might be expected, the paths become still more winding still, and all the less human-like.

In order to counter these two problems, we introduce agents that continue in a straight line unless otherwise perturbed through an affordance threshold being reached. But this raises the question: how should the threshold be set? In the next section, we show how a genetic algorithm can be used to evolve to a solution.

4 Evolving line-of-sight agents

In this section, we demonstrate that superior agent behaviors can be evolved to better fit human movement patterns, but that it is the visual information itself that is important, rather than the visual information flow, to reach this goal. In order to give the agents more appropriate response to the environment, we set them to walk forward continuously. We then gave each one a set of eight




Fig. 4. Trails of 19 real people followed for the first 10 minutes of their visit to the Tate Britain Gallery, reproduced from Hillier et al. (1996), p. 15 [22].

paired receptors (see figure 4), one pair at each of 1 bin, 3 bins, 5 bins and 7 bins away from the current heading. The input received by the receptor itself was changed according to the experiment conducted, but in each case the basic rules of response were kept constant as follows:




Fig. 5. Inputs to the evolutionary agents are through eight evenly-spaced inputs reflecting properties of the ambient optic array at its current location and heading.

- Each pair of receptors is checked in a set order, determined by a genetic algorithm gene, for example, 3, 1, 2, 4, which represents look at the third pair of receptors first (5 bins away from the current heading), then the first pair (1 bin away from the current heading), second pair and finally the fourth pair.
- For the pair of receptors being checked, both its inputs are compared to the a genetically determined threshold.
- If one or other of the receptors is above the threshold, then it turns to the receptor according to a genetically determined probability. If both receptors are above the threshold, it turns to either one with the same probability.
- If the input actually triggers the turn, the remaining receptor pairs in the list are not checked and the agent makes a step.
- If no receptor pair is chosen, then the agent makes a step in the current forward facing direction.

The aim of this set of rules is to allow the agent its own evolved choice of FoV to apply, and its own choice of which directions are most important to it. The scheme is far more simple than neural models used for other animat approaches such as evolutionary robotics, but we wanted to get to the heart of how the agent was responding to the environment, and this method allows open analysis of what control method has actually evolved.

The method of evolution was a simple genetic algorithm with nine genes (one for the rule order, and eight for threshold and probability for each pair of genes). Chiasma points were placed between the genes, and the genetic algorithm used rank selection, uniform crossover and a 5% mutation rate to evolve the agents, see [23] for details of genetic algorithm operators. The parameters for the algorithm were a population size of 300 agents, evolved for a period of 3000 generations.

500 agents of each phenotype to be tested were released into the Tate Britain Gallery plan from the entrance and allowed to explore for 1600 timesteps as before. At the end of the run, the fitness was assessed as simply the correlation coefficient (R^2) between agent and observed room through movement for the 55 rooms and corridor spaces of the gallery. The experiments were run on a PC with an AMD Athlon 2400+ chip, and each took about 2 hours to complete.

Four different types of receptor were applied. Firstly, a straight forward threshold based on distance. Each threshold was initially set to a random value between 0 and 100m (about 2/3 of the length of the central axis). Secondly, in an attempt to see if informational flow could guide the agents better, the threshold was based on length of LoS in the receptor direction divided by its length on the previous step (both in the same direction). This threshold was given a value between 0 times as far and 5 times as far as before. Thirdly, we attempted to see if a comparison with the forward direction had an effect. This threshold was defined as the length of LoS in the receptor direction divided by the length of LoS in the forward direction. This threshold was given a value between 0 times and 10 times as far as the forward direction. Finally, we implemented a threshold based on comparison of the informational flow in receptor and forward direction, dividing one by the other as before. The values used were again between 0 and 10 times the informational flow in the receptor direction.

Table 2. Fitness values and corresponding rulesets for best performing evolved agents

	Directional		Comparative	
	Line length	Differential length	Line length	Differential length
Fitness	0.81	0.54	0.84	0.56
Rule Order	3,1,2,4	1,3,4,2	2,4,3,1	4,3,1,2
Rule 1 (Bin 1 $\approx 10^\circ$)				
Threshold	16.5m	2.40	0.22	0.77
Probability	0.94	0.01	0.40	0.00
Rule 2 (Bin 3 $\approx 35^\circ$)				
Threshold	8.5m	4.09	9.13	2.09
Probability	0.67	0.70	0.04	0.46
Rule 3 (Bin 5 $\approx 55^\circ$)				
Threshold	20.1m	2.39	8.58	4.25
Probability	0.07	0.18	0.27	0.85
Rule 4 (Bin 7 $\approx 80^\circ$)				
Threshold	6.3m	3.92	1.66	3.79
Probability	0.86	0.83	0.27	0.03

Table 2 shows the result of evolving the four types of agent. Note that the fitness is only a rough correlation of the agent movement with people movement, given that 500 agents were used for the sampling rather than the 2000 used to test the correlation for the basic agents. However, there is a clear difference between the performance of informational flow agents, which perform worse,

and the length of LoS agents. Certain things are noticeable about the agents: the forward bins are an obvious hinderence to the informational flow agents, as they evolve to have almost 0 probability of firing them. Indeed, all agents had one pair of receptors that they effectively turned off. The best performing agents were the comparative LoS agents, with a fitness value of 0.84. Interestingly, these seem to apply their first rule with probability of around 1/3, and are almost certain to go above the threshold of the forward facing bins, which is set to trigger movement if above just 0.22 times the length of the current heading. Perhaps this minor adjustment is useful to further exploration, and could the approximate 1/3 probability be the 3-step rule revisiting us once again? Indeed, how do these agents compare with the basic agents?

At first glance, there really does not seem to be much difference between the comparative LoS agents' behavioral representation of people movement and that of basic agents. Figure 4 shows scatter plots of room through movement for the basic LoS agents and the evolved comparative LoS agents. In each case the scales have been adjusted to using cube roots, and set on a scale with a maximum of one. Although it the evolved version is slightly straighter (as measured by linear regression, it is not obvious from the diagram), the noticeable thing is that the outliers are all in approximately the same places. That is, it is striking that the underlying data for these two types of agents is almost identical. However, the difference becomes obvious when we look at figure 4. The evolved agents have much more realistic qualitative patterns of movement, although we have not compared the paths themselves for quantities such as mean angle of turn or frequency of major turns.

One further issue that we have not considered fully is the matter of the intercept of the best fit line. The adjusted scales in figure 4 do not show it properly, but the best fit line does not go through the origin. In fact, it is significantly out: $0.323 \times \max y$, where $\max y$ is the cube root adjusted maximum observed people count⁵. We tried to evolve the intercept difference out, with some success, by introducing a second component to the fitness function based on the distance of the intercept, c :

$$F_c = \frac{1}{|c|/\bar{y} + 1} \quad (1)$$

This function has the useful properties of a range of 0 (for intercept infinitely out) to 1 (for perfect intercept with the origin), and thus could easily be added to our correlation coefficient based measure. Applying this fitness function to the evolution of comparative LoS agents did improve the intercept variation (down to $0.253 \times \max y$), but the result was to take the correlation back down to the same as standard LoS agents. It appears, simply, that line-of-sight agents are under-representing people in areas of high density, and over representing them in areas of low people room through movement.

⁵ Note that if the cube root is removed, the intercept as a proportion of y is naturally much better, but we were interested in trying to remove this systematic error.




Fig. 6. Scatter plots of adjusted room through movement rates for 3-step LoS and evolved comparative LoS agents against observation data for the Tate Britain Gallery.




Fig. 7. Trails of 20 evolved comparative LoS agents for the first 10 minutes of their ‘visit’ to the Tate Britain Gallery.

5 Conclusion

In this paper, we have shown that we are able to reproduce aggregate pedestrian movement within buildings using extremely simple line-of-sight (LoS) based rules. Agents that use movement rules based on length of line of sight to assess the affordance of that direction for further movement correlate with $R^2 = 0.78$ with observed numbers of people walking through rooms in an art gallery environment. However, we have also shown that the individual paths of these agents do not tally with those of people, adding many more turns than would be expected. Furthermore, the agents are calibrated to correlate well with people numbers by forcing them to make a choice of destination only once every three steps. In general situations, this would lead them to miss turnings based entirely on where they are in their walk cycle.

In order to overcome these problems, we programmed agents set to turn based on thresholds of informational content in the environment. We invoked Benedikt’s isovist fields to suggest that picking a destination may be based on information or informational flow as we move through the environment. We evolved agents using these inputs, and found that informational flow did not lead to good correspondence with people movement as might have been expected. However, making a decision based on the length of LoS ahead compared to the length of LoS in a candidate direction did lead to marginally improved movement models, with $R^2 = 0.84$.

When we came to compare these evolved agents with the basic agents, we found that there were some striking similarities between them: the 3-step rule seemed to have been replicated in a 1/3 probability level of turn in the evolved agents, and the distribution of room through movement counts was almost identical across the two. This could indicate that there is an underlying constant of natural movement we are uncovering, or, of course, it could simply be that we are overfitting both types of agents to the example of the chosen art gallery observed through movement rates. In further research, we hope to evolve the agents in one environment and apply them against another, which should inform us better about whether or not this is a fundamental phenomenon of human behavior or not.

There was, however, one major difference between the two sorts of agents. Although *where* they go is very similar, *how* they go about it is quite different. The basic agents turn continually, progressing in almost circular motion of onward circles; the evolved agents, by contrast, advance in much straighter lines, similar, at least qualitatively, to those of people observed within the gallery context.

References

1. Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston, MA (1979)
2. Trullier, O., Wiener, S.I., Berthoz, A., Meyer, J.A.: Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology* **51** (1997) 483–544
3. Darken, R.P., Sibert, J.L.: Wayfinding strategies and behaviors in large virtual worlds. In: CHI. (1996) 142–149
4. Mallot, H.A., Gillner, S., Steck, S.D., Franz, M.O.: Recognition-triggered response and the view-graph approach to spatial cognition. In Freksa, C., Mark, D.M., eds.: COSIT. Volume 1661 of Lecture Notes in Computer Science., Springer (1999) 367–380
5. Montello, D.R.: Spatial orientation and the angularity of urban routes. *Environment and Behavior* **23** (1991) 47–69
6. Conroy Dalton, R.: The secret is to follow your nose: Route path selection and angularity. *Environment and Behavior* **35** (2003) 107–131
7. Golledge, R.G.: Path selection and route preference in human navigation: A progress report. In Frank, A.U., Campari, I., Formentini, U., eds.: A Theoretical Basis for GIS. Volume 639 of Lecture Notes in Computer Science., Springer (1995) 207–222
8. Montello, D.R.: The perception and cognition of environmental distance: Direct sources of information. In Hirtle, S.C., Frank, A.U., eds.: COSIT. Volume 1329 of Lecture Notes in Computer Science., Springer (1997) 297–311
9. 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
10. Esch, H.E., Zhang, S., Srinivasan, M.V., Tautz, J.: Honeybee dances communicate distances measured by optic flow. *Nature* **411** (2001) 581–583
11. Vaina, L., Rushton, S.: What neurological patients tell us about the use of optic flow. *International Review of Neurobiology* **44** (2000) 293–313

12. Hermer, L., Spelke, E.S.: A geometric process for spatial reorientation in young children. *Nature* **370** (1994) 57–59
13. Benedikt, M.L.: To take hold of space: Isovists and isovist fields. *Environment and Planning B: Planning and Design* **6** (1979) 47–65
14. Conroy, R.A.: Spatial Navigation in Immersive Virtual Environments. PhD thesis, Bartlett School of Graduate Studies, UCL, London (2001)
15. Brooks, R.A.: Intelligence without representation. *Artificial Intelligence* **47** (1991) 139–159
16. Kuipers, B.J., Byun, Y.T.: A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems* **8** (1991) 47–63
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., Doxa, M., O’Sullivan, D., Penn, A.: From isovists to visibility graphs: a methodology for the analysis of architectural space. *Environment and Planning B: Planning and Design* **28** (2001) 103–121
19. Ooi, T.L., Wu, B., He, Z.J.: Distance determined by the angular declination below the horizon. *Nature* **414** (2001) 197–200
20. Kerridge, J., Hine, J., Wigan, M.: Agent-based modelling of pedestrian movements: The questions that need to be asked and answered. *Environment and Planning B: Planning and Design* **28** (2001) 327–341
21. Sutherland, D.H., Kaufman, K.R., Moitza, J.R.: Kinematics of normal human walking. In Rose, J., Gamble, J.G., eds.: *Human Walking*. Williams & Wilkins, Baltimore (1994) 23–44
22. 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)
23. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, London, UK (1989)