Environment and Planning B: Planning and Design 2002, volume 29, pages 473 – 490

DOI:10.1068/b12850

Encoding natural movement as an agent-based system: an investigation into human pedestrian behaviour in the built environment

Alasdair Turner, Alan Penn
Bartlett School of Graduate Studies, University College London, 1 – 19 Torrington Place,
London WC1E 6BT, England; e-mail: a.turner@ucl.ac.uk; a.penn@ucl.ac.uk
Received 14 August 2001; in revised form 30 November 2001

Abstract. Gibson's ecological theory of perception has received considerable attention within psychology literature, as well as in computer vision and robotics. However, few have applied Gibson's approach to agent-based models of human movement, because the ecological theory requires that individuals have a vision-based mental model of the world, and for large numbers of agents this becomes extremely expensive computationally. Thus, within current pedestrian models, path evaluation is based on calibration from observed data or on sophisticated but deterministic route-choice mechanisms; there is little open-ended behavioural modelling of human-movement patterns. One solution which allows individuals rapid concurrent access to the visual information within an environment is an 'exosomatic visual architecture', where the connections between mutually visible locations within a configuration are prestored in a lookup table. Here we demonstrate that, with the aid of an exosomatic visual architecture, it is possible to develop behavioural models in which movement rules originating from Gibson's principle of affordance are utilised. We apply large numbers of agents programmed with these rules to a built-environment example and show that, by varying parameters such as destination selection, field of view, and steps taken between decision points, it is possible to generate aggregate movement levels very similar to those found in an actual building context.

Introduction

In the literature on cognitive psychology, the issues raised by Gibson's ecological theory of perception (Gibson, 1979) have been examined and taken on board. Gibson's theory was formulated primarily in order to overturn theories laden with subjective and objective knowledge, and to replace them with a model in which the agent and its environment are conjoined by a set of *affordances* so the agent perceives the contents of the environment directly and uses the affordances within it to guide its action without reference to superior representational models. Today, Gibson's work has been contextualised and broken into further models in which recognition and representation do play a part (Neisser, 1994). However, in the domain of agent-based modelling we still appear to ignore the original concerns he voiced—in particular 'the model' is always preceded by a theoretical framework, rather than simply being a perceptual model in its own right (see, for example, Casti, 1998; Epstein and Axtell, 1996). This paper is not an attempt to overturn the body of literature which already exists in agent-based modelling, although it does constitute a plea for the use of direct perception where the approach is available, and to try to regard the environment as the provider of possibilities rather than as a place to be rationalised. As an example, consider human movement around an art gallery. There might be any number of causal factors for the routes people take. People might, for example, follow a map, or signage, take into account the direction other people are taking, a glimpse of a familiar painting, reject a route on the grounds of personal prejudice against a style, and so on. On the other hand, the possibility of exploring the walkable surface of the layout ahead (the rooms of the gallery) may simply be enough for a human to do exactly that. If this is the case, we should be able to reproduce movement levels within the gallery by using an agent-based model, with movement rules based solely on building configuration, without recourse to models involving learned paths, goals, or destinations, or any more detailed social theory framework. As such, in this paper we attempt to find out to what extent it is possible to use configuration alone to explain movement.

Of course, it is obvious that socioeconomic factors do affect human behaviour at a fundamental level. Why would someone walk down a street if not to buy a loaf of bread, to go to work, or to meet and interact with someone or something? Thus, when social science turns to agent-based modelling, it seems natural that at least some component should be based on the concept of the rational being. One can even argue that any human-based activity can be regarded as cost – benefit behaviour. For example, the putative trip to an art gallery permits esoteric gains in the form of artistic appreciation in exchange for the more straightforward cost of hours of walking. Such a cost-based rule set can be encoded in an agent-based system. Hoogendoorn et al (2001) have demonstrated a cost-minimising approach to human-pedestrian modelling in an airport layout, producing cost paths to buy a newspaper or catch a train or plane. However, is it really plausible that the human brain continually reassesses an internal cost function, or is it that the human is led by less tangible factors—her curiosity or his desire? It is also obvious that humans are physical entities: that one human moving through one place at one time must do so to the exclusion of other humans. When looking at the interaction of a group of humans, it seems natural, in turn, to represent the system as a group of interacting particles. Humans moving through a crowd would seem (at least superficially) to have many similarities with granular flow. For example, people walking along a street may be represented as particles driven in a certain direction, pushing through one another and forming lanes. Again, this can be encoded in an agent-based system: Helbing and Molnár (1997) have demonstrated lane formation of directed particles in a pipe system. Now the concern is reversed: does the corporeal human bump through a crowd of corporeal humans, or does the human guide him or herself through gaps in the crowd? What appears to be constraining both Hoogendoorn et al and Helbing and Molnár is the ability to see. Without it, their models can only be, literally, stabs in the dark at able-sighted human behaviour. We propose to add the ability to see, to investigate this ability as an ecological phenomenon leading to natural movement, and then return to socioeconomic factors and granular physics to provide the constraining framework for the models—as and when necessary.

Soley by adding the ability to see we are led to an intuitively attractive, although still incomplete, model of human-pedestrian behaviour: that the human moves in a direction that provides him or her the potential for possible further movement. Gibson calls such interaction between human and environment 'natural vision':

> "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 ..." (1979, page 1).

Thus, natural vision is a combination of visual factors affecting behaviour. We might characterise "we look around" and "we go from one vista to another" as the conditions of *natural movement*, while thinking of "we walk up to it" and "we move around it" as something else, for example, *natural interaction*. The distinction is important: for natural movement, an agent does not even require the ability to recognise 'object' as distinct from 'environment': the agent merely has to recognise that there is environment which may be explored in order to *move* (though not necessarily to *navigate*). The names we have chosen are no accident, for there already exists a theory of natural movement. Hillier et al (1993) show that the majority of human-pedestrian movement occurs along lines of sight, and that the more 'integrated' (in terms of connection to other lines of sight) a line is, the more movement exists along it. Hence, the theory of natural movement is this: human movement, ceteris paribus, is *generated* by configuration. The relationship of configuration (Hillier et al) to surface (Gibson) is direct. The configuration as provider of possibility, and walkable surface as provider of possibility, are equivalent. However, the difficulty with Hillier et al's theory is in the words 'ceteris paribus'. What are the other things that are equal? Although the theory is never stated explicitly as such, we can infer that it is the socioeconomic framework and physical constraints that must be equal,(1) that is, they are the remaining factors identified by Epstein and Axtell (1996, page 1) as essential ingredients of a complete agent-based model (economic, demographic, and cultural)—although there is no need to identify a priori what these factors might be. In order to keep a consistent model, Hillier et al used a small area around King's Cross, London, as an example system. For our sample experimental study, we shall return to the putative art gallery. An art gallery represents a good control, as it attracts a small socioeconomic cross-section (Berger, 1972, page 24), and, being both reasonably small and for the most part on the one floor, also physically homogeneous. There are, of course, still many factors which are not included in our model; however, as we shall see, these appear to be of minor importance.

In the next section we review related work from the fields of robotics, agent-based pedestrian models, microeconomics, and spatial cognition. Against this background, the basis for our experimental model, the exosomatic visual architecture is introduced and the natural-movement methodology explained. We then apply the experimental model to the Tate Britain Gallery, London, and compare the results with observations of real people movement in the gallery. We conclude with a discussion of the implications for agent-based modelling of pedestrian movement.

Related work

Research on robots has examined action guided by vision for some time, and recently Murphy (1999) has explicitly made the link to Gibson's ecological approach. However, most robot research has concentrated on the building of map structures, with emphasis being placed on neurological evidence of a cognitive map underlying navigation (O'Keefe and Nadel, 1978). Thus robots have come to base their navigation on the hippocampus in mammals (see, for example, Bachelder and Waxman, 1994). Based on neurological evidence, Trullier et al (1997) propose a four-level hierarchical framework for movement: guidance, place-recognition triggered response, topological navigation, and metric navigation. Of the items in this hierarchy, topological navigation has received most attention in the literature starting with simplified topological maps (Laumand, 1983), followed by techniques such as robot learning of visibility graphs to record collocation of structures within the environment (Oommen et al, 1987), through to a tripartite system of low-level sensorimotor control, topological mapping, and geometric path recording to move around and record an environment (Kuipers and Byun, 1991). All of these approaches are interesting from an architectural point of view because they mirror the graphing techniques used by various researchers in the built environment. Laumand's technique may be compared with the topological mapping of units produced by Steadman (1973); Oommen et al's work has parallels in Braaksma and Cook's (1980) visibility graphs of airport terminals. Kuipers and Byun's robots use a systematic retrieval of 'distinctive places', which is akin to the use of isovist measures to identify places (Benedikt, 1979). Of course, there is not a one-to-one mapping of all

(1) Developing the theory of the movement economy, Hillier (1996) argues that the socioeconomic framework itself is affected by the configuration.

techniques, but there is a certain parallel in development. Such techniques could be applicable to a natural-vision approach to robot perception; however, the overriding aim in robot research has been cognitive map building, and so the robots have been restricted to relatively simple environments where the possibilities of natural movement are not available. When it has been considered, the examination of natural movement has concentrated on the problems of sensorimotor control; for example, keeping the robot equidistant from each wall. Thus, comparison with human-movement patterns, or large-scale experimentation with many agents, is not a feature of robot research.

Large-scale simulation has been investigated for some time by means of agent-based pedestrian models. Pedestrian models can be broadly categorised into three levels: macro, meso, and micro (Hoogendoorn and Bovy, 2000). The macrosimulation level is essentially inherited from transportation modelling, stemming from Lighthill and Whitham's (1955) continuum model; that is, solving differential flow equations. Al Gadhi et al (2001) have recently shown how such models may be applied to pedestrian situations, constructing a flow model of events occurring during the Hajj. By contrast, the mesoscopic level is ammenable to agent simulation although, again, it is primarily suited to traffic modelling. This class of system includes the cell-transmission model of Daganzo (1994) and TRANSIMS, the cellular-automata-based system (Beckman, 1997; Casti, 1998; Nagel et al, 1996), which may model 20 000 agent journeys concurrently. Both present a good example of an urban-level simulation, albeit designed for traffic rather than pedestrian flow. The systems combine set origin – destination (OD) pairs with route-choice behaviour to simulate journeys within the system. In TRANSIMS the agents make a shortest-time assumption, based on the previous run flows, and readjust to optimise a combination of trip time and trip length for the next simulated day. Work on such cellular automata (CA) models has been applied to pedestrian movement (for example, Blue and Adler, 1998). However, in pedestrian models, where there is high linkage between cells, the CA community have yet to implement effective natural-movement rules. For example, Blue and Adler's model uses unidirectional flow. Microsimulation has to date focused mainly on granular-physics models of flow, with models such as Helbing and Molnár's (1997) crowding simulation again using predetermined directional paths. These have led to observations of lifelike emergent phenomena based on simple rules such as lane forming simply by a predisposition to move to the left or right in the face of oncoming traffic (for further details, see Helbing et al, 2001). Fire evacuation models also fit into the microsimulation category, with agents released from set positions finding their way to the nearest exit again by predetermined paths (for example, Galea et al, 1996). Again, in fire-emergency examples, people will be moving towards an exit so this is a valid routing mechanism for this particular scenario, but perhaps not for the more general case where a choice of destinations is offered.

Certainly, there is interest in human routing, particularly from a microeconomic perspective. Borgers and Timmermans (1986) initiated research into route-choice behaviour in humans, implementing submodels for destination choice, route choice, and impulse stops in place of the more usual gravitational attractor models used in macroscopic and mesoscopic simulations. Originally, assumptions about shortest paths between individual destinations were made; more recently, these have been replaced by sets of local optimisations, so that the human may be assumed to take the shortest path from shop to shop after arriving at a shopping area (Kurose et al, 2001), in order to account for the fact that humans do not always have complete global knowledge of a system. This kind of behaviour has also been taken into account in Hoogendoorn et al's (2001) cost-based model of human-pedestrian behaviour, where a number of different strategies are employed and a cost-surface calculated. However, cost-surface methods rapidly become more complicated if combinations of destinations and goals are incorporated; for example, buying a paper then catching a train involves a dual cost-surface for the agent states of 'bought paper' and 'unbought paper', if the route to be traced must return on itself. However, more importantly, neither Kurose et al's nor Hoogendoorn et al's methods have been validated against empirical data, although there is good evidence to suggest that local route optimization is considered by shoppers (Gärling and Gärling, 1988). Of course, any full microscopic model does require that the goals of individual agents are programmed in, but it should be noted that it has been discovered that even when OD pairs are well specified, and a participant is well acquainted with an area, a shortest-path route is not always followed (Golledge, 1995). Hence, even assumptions about local shortest-path planning cannot be made. We might add to this other objections about using OD pairs: that at the start of the trip the destination may not have been decided (Batty, 2001). And objections concerning the use of route-choice behaviour: that appropriate junctions have to be identified in advance, and that the behaviour at these junctions must be calibrated against real movement levels discovered at real junctions (thus reducing the model to a series of predetermined flows between junctions).

However, there is also interest in microscopic human-movement patterns from a spatial-cognition perspective. Golledge (1995) asked participants to walk to destinations within a college campus. He discovered that people do not always take the shortest route: in particular, that there are differences in taking a route from O to D and a route from D to O:

> "perceptions of the configuration of the environment itself ... may influence route choice. Thus, a route that seems shorter or quicker or straighter from one end may not be so perceived from the other end ..." (page 221).

That a configuration significantly affects movement at a microscopic level has also been discovered by researchers interested in space syntax (Hillier and Hanson, 1984). Peponis et al (1990) and, more recently, ul Haq (2001), make empirical studies of people movement in hospital environments. Peponis et al attempt a categorisation of a number of 'rules' of navigation:

(1) Avoid backtracking.
(2) If all else is equal, continue in the same direction.
(3) Divert from the current heading when a new view allows you to see more space and/or activity (that is, other people).

Only after these do they state rules requiring further cognitive ability (such as searching unexplored regions first). Conroy Dalton (2001) makes a further inroad into the importance of configuration. In experiments in a virtual environment, by carefully preparing an environment where it is possible to move only to similar-sized spaces, she found that there is also a tendency to reduce the total angle turned. If we are to produce a good model of natural movement it should take into account these findings: that is, the agents should follow Peponis et al's rules of navigation and, ultimately, follow Conroy Dalton's angular minimisation. Such a natural-movement system, in which agents are allowed to move freely according to their vision, has been considered by many to be too computationally expensive to implement. Thomas and Donikian (2000) state that "... a complete mental model based on vision and image processing cannot be constructed in real time using purely geometrical information" (page C71). To overcome this problem, they limit environmental perception significantly in their Virtual Urban Environment Modelling System (VUEMS). VUEMS guides agents by the use of a combination of visibility graphs *linking* points of interest, and Voronoï diagrams providing routes *around* those points of interest. Thus VUEMS begins to address the interaction between human and object described by Gibson. Other researchers have attempted natural movement, although they too have been restricted by complexity. In particular, Terzopoulos's (1999) fish have vision, but the environment is very simple, and the movement rules are based on Reynold's (1987) Boids-type behaviours.

There have also been configurational approaches to agent-movement rules through consideration of space syntax. For example, Penn and Dalton (1994) encode space as a set of axial lines (the longest lines of sight covering all the nontrivial loops in the configuration). Although applying these rules to a much-simplified system, Penn and Dalton found that the best rules for movement are very similar to those described by Peponis et al (1990) and Conroy Dalton (2001). They assign random OD pairs in the system and have 'rats' search for these goals, find the best combination of rules (that is, the best correlator with observed human movement in their chosen region) as following the longest line towards a destination, with a secondary factor of taking the minimum angle towards the destination. The heuristic they discover is strikingly similar to Gärling and Gärling's (1988) earlier findings, that shoppers optimise path length on routes by taking on the longest segment first. In order to achieve model natural movement, though, the assignation of OD pairs must be dropped, and agents allowed to roam freely according to their desires. Again considering the space-syntax approach, Mottram et al (1999) apply agents that decide on which direction to go based on the length of line of sight from their current position. However, in this experiment agents cast rays to sample the environment and thus the model suffers from the processing overheads Thomas and Donikian identify, which restrict the model to a few concurrent agents. To extend this to many simultaneous agents in an arbitrary space, a vision-based system requires a new approach to environmental perception. Such an approach is provided by an exosomatic visual architecture.

Method

In this section we examine the implementation of a natural movement. First, as a precursor for vision, the exosomatic visual architecture is introduced. Second, we explore how this can be used as a basis for encoding affordance-based rules. Third, details and refinements of the implementation are described for our experimental model.

Enabling vision

In Penn and Turner (forthcoming) we introduce an exosomatic visual architecture (EVA) for agent guidance. Underlying the EVA is a dense-grid visibility graph. The visibility graph is computed by overlaying a two-dimensional grid (at some arbitrary resolution) over a layout in plan view, and calculating which points within the grid are able to see which other points. The set of visible locations for each point are stored, and thus the visibility graph can be used to calculate the approximate viewable area, or isovist, from each point on the grid (see Turner et al, 2001, for details). If the set of locations visible from each point is further subdivided into bins, the approximate isovist can be split into angular segments. For the purpose of our experiment, we split the approximate isovist into 32 bins, allowing us to select a set of viewable locations from a point divided between 11.25° angular segments, as shown in figure 1.

From a mathematical point of view, if the graph is considered to be the combination of the sets of vertices and edges, $G(V, E)$, then the set of edges E may be broken down into a set of nonoverlapping bins $E = B_1 \cup B_2 \cup \ldots B_{32}$. Each bin B_b contains ordered pairs of vertices $[v_\alpha \rightarrow v_\beta]$, such that v_α can see v_β, and the heading from v_α to v_β rounds to an integer b (where b runs from 1 to 32, representing $0°$ to $360°$ in $11.25°$ segments). Thus, for an agent at a location $v(x, y)$, an approximate visual field can be

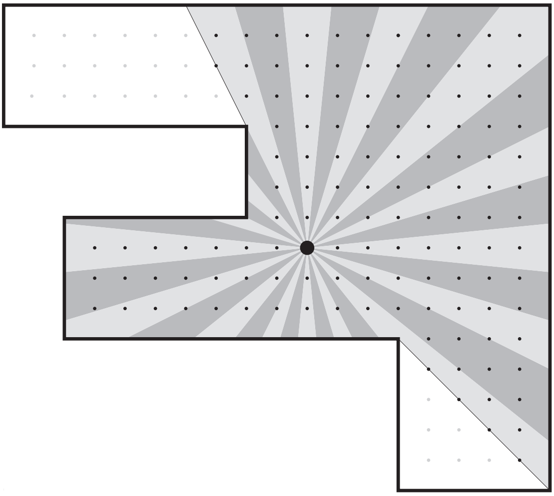

Figure 1. Visible locations on the visibility graph grid are split between 32 angular bins.

established by rounding the agent location to the nearest grid point v_a, rounding the agent heading to the nearest bin b, and taking a subset of v bins around this bin to form the visual field F:

$$F_{a,b}(v) = \{v_f : [v_a \rightarrow v_f] \in \bigcup_{b-(v-1)/2}^{b+(v-1)/2} B_b, \text{ and } v_a, v_f \in V\}, \qquad (1)$$

where v is a parameter giving the angle of the field of view in bins, which may be converted later back to degrees. Figure 2 shows an agent at such an arbitrary location and heading, and the associated field of view. Because the information for perception of the geometry lies outside the individual agents, we call the architecture *exosomatic*. The architecture also allows further enhancements, such as rapid access to which other agents are currently visible, or which paintings on the wall can be observed from the current position, although at the time of submission we are still completing these sections of the model.

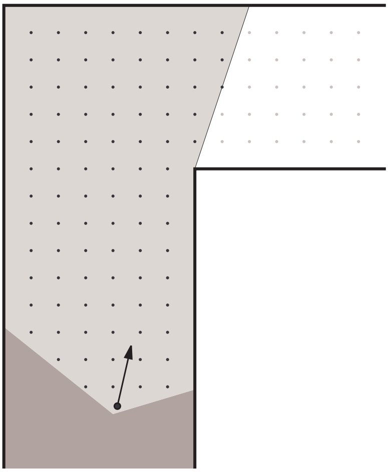

Figure 2. The visibility graph provides an agent (angle of view $v = 11$ bins) with a lookup table for possible next destinations. The current location is rounded to the nearest grid point on the visibility graph, and current heading is rounded to the nearest bin.

Encoding affordance

The space syntax-guided research of Peponis et al (1990), Hillier et al (1993), and Penn and Dalton (1994) has emphasised that it is the length of line of sight that is important to natural movement and search tasks. The axial line, as introduced by Hillier and Hanson (1984), is considered as the primary guiding mechanism of human-pedestrian behaviour. By contrast, Gibson (1979) asserts that it is the available walkable surface which affords movement. The two are clearly extremely closely connected: a long line of sight also presents a large surface area, and vice versa. However, Gibson's theory is perhaps the more concise, as the concept of 'line' is unnecessary to his notion of movement.(2) Of course, the environmental approach to perception is more complex than a mere appraisal of available surface—for example, there are objects within the environment—but at its most basic, the environment is considered as a walkable surface broken by other surfaces. Walkable surface provides affordance for further walkable surface (where we might replace walkable with 'wheelchairable', as appropriate, although notice that the two may well not be equivalent, and we might expect the natural movement of wheelchair users to be different from that of pedestrians). Stated as a hypothesis:

Hypothesis: *When engaging in natural movement, a human will simply guide him or herself by moving towards further available walkable surface. The existance of walkable surface will be determined via the most easily accessed sense, typically his or her visual field.*

Such a formulation also allows a concise agent implementation in an exosomatic visual architecture: the agent merely has to choose a location to walk through a stochastic process for it to be engaging in natural movement. We use a correspondingly simple agent-decision process:

```
Loop
    Pick a visibility graph vertex from the field of view $F_{a,b}(v)$ by selecting any vertex from $F$ with equal probability.
    Take, on average, $n$ steps towards that vertex, based on a Poisson distribution.
End loop
```

Hence n (the number of steps) and v (the angle of view) define the parameter space for our agent model. Space or, more correctly, configuration is the guiding mechanism, as shown in figure 3. Note that what appears to be a dichotomy must be resolved: Gibson clearly denies that the concept of space in the Cartesian sense is of relevance to animal perception, and yet we appear to propose to guide our agents by an assessment of Cartesian space, by allowing them to choose a location on a Cartesian surface. However, our agents are *not* in fact 'seeing' space, but are basing their decisions on the availability of a destination—a point on a surface within the environment—which *affords* them the possibility of a further destination, and so on. That this is the case can only be demonstrated with an example. Figure 4 shows the path of an agent actually recorded during our experimentation; this sort of path interested us as it recurred time and time again. The agent seemingly interacts with an art installation in the environment: going up to it, walking around it, and through it. Actually, the agent sees the *possibility* to walk up to an object, walk around it, and through it, and *so it does.* That this happens is remarkable, and an interesting emergent phenomenon (we did not code an agent with an appreciation of artistic installation—although whether this is strictly emergent is, of course, open to debate; for example, as we note later, the single-agent case is simply a Markov process).

---

(2) The converse is not true: a line of sight presupposes a space, or surface, through which the line is drawn. Penn (2001) demonstrates how the axial line may in fact be retrievable from spatial exploration.

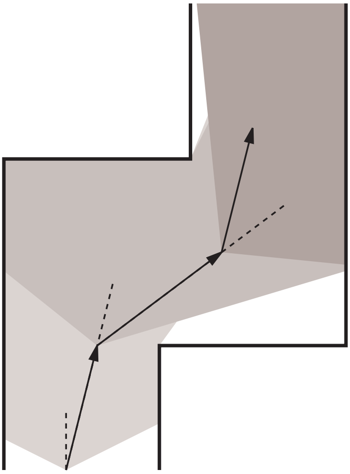

Figure 3. Theory: selecting destinations from the visibility graph through a stochastic process should draw the agent through a configuration.

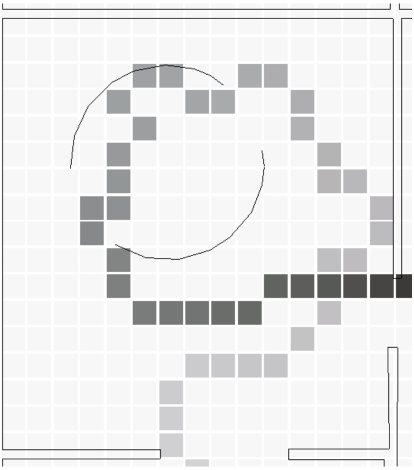

Figure 4. Natural movement in action. In an experiment, the configuration of an objet d'art draws an agent through it. Grid squares $0.75 \text{ m} \times 0.75 \text{ m}$ are coloured darker the more recently the agent has moved through them.

It should be noted that there are genuine problems with the use of line-of-sight or walkable surface as affordance. Our model uses *infinite* sight, and therefore an infinitely long corridor with respect to side corridors would drive all movement continuously along that corridor, whereas we might expect a human to take an exit some way along the corridor. The opposite of a long thin corridor—a large open space with no configurational clues—also poses a problem to the agent: it will be 'confused' by the area and unable to walk consistently in any direction (this is less of a problem than the long thin corridor: a human placed in a snowstorm without visual clues will also suffer confusion and lack of direction). For the time being, we will ignore this problem as neither infintely long corridors nor featureless unbounded space occur in practice. As a future adaptation we might consider evidence that human perception of distance by angle can foreshorten long distances (Ooi et al, 2001), which may resolve situations in which agents are biased by large expanses of space in the distance.

EVA system

So far our discussion has been concentrated on arbitrary resolution systems, and we have ignored the physicality of the agents. For the purposes of producing an experimental model, we apply a visibility graph with a 0.75×0.75 m resolution grid. The value is derived from the average step length of mature humans—approximately 0.77 m (Sutherland et al, 1994), so that possible destinations are rounded to places within one step of each other. The use of step length reflects our own statement that a dense-grid visibility graph should be constructed at a resolution appropriate to its use, and in this case mapping the space that is humanly accessible is sufficient (Turner et al, 2001, page 106). In addition, Sutherland et al also show that walking pace on a level surface is in fact very consistent in humans ($\sim 1.5 \text{ ms}^{-1}$), and so we can safely ignore, at least within an art gallery, concerns about different agent speeds that may arise (such as those expressed by Kerridge et al, 2001). Thus, each agent in the model moves at a constant simulated speed of 1.5 ms^{-1}, that is, taking one grid-square-sized step every half second, although, as the agent may take any heading, the precise location of each agent is in continuous rather than discrete coordinates.

The agents are also given a physical presence in that, rounding the position to the nearest grid location, no two agents can exist in the same location. As the agents are currently unable to see each other, this may (and does) lead to deadlock situations when agents are travelling in opposite directions. Thus, we introduce an antideadlock rule: a side-step rule whereby the agent attempts to take a sideways step to avoid the obstacle (to which side the step is taken is chosen at random). Although this rule works most of the time when agents collide, in crowded situations agents may (and do) achieve gridlock when agents slot into holes left by the other agents—a situation dubbed 'freezing by heating' by Helbing and Molnár (1997), who demonstrate the effect using granular physical agents. Therefore, we add a further rule: if, because of blockage, an agent cannot apply a side step it gives up and chooses at random any location from its entire 360° field of view as a new destination.$^{(3)}$ The full agent-decision process is shown in figure 5.

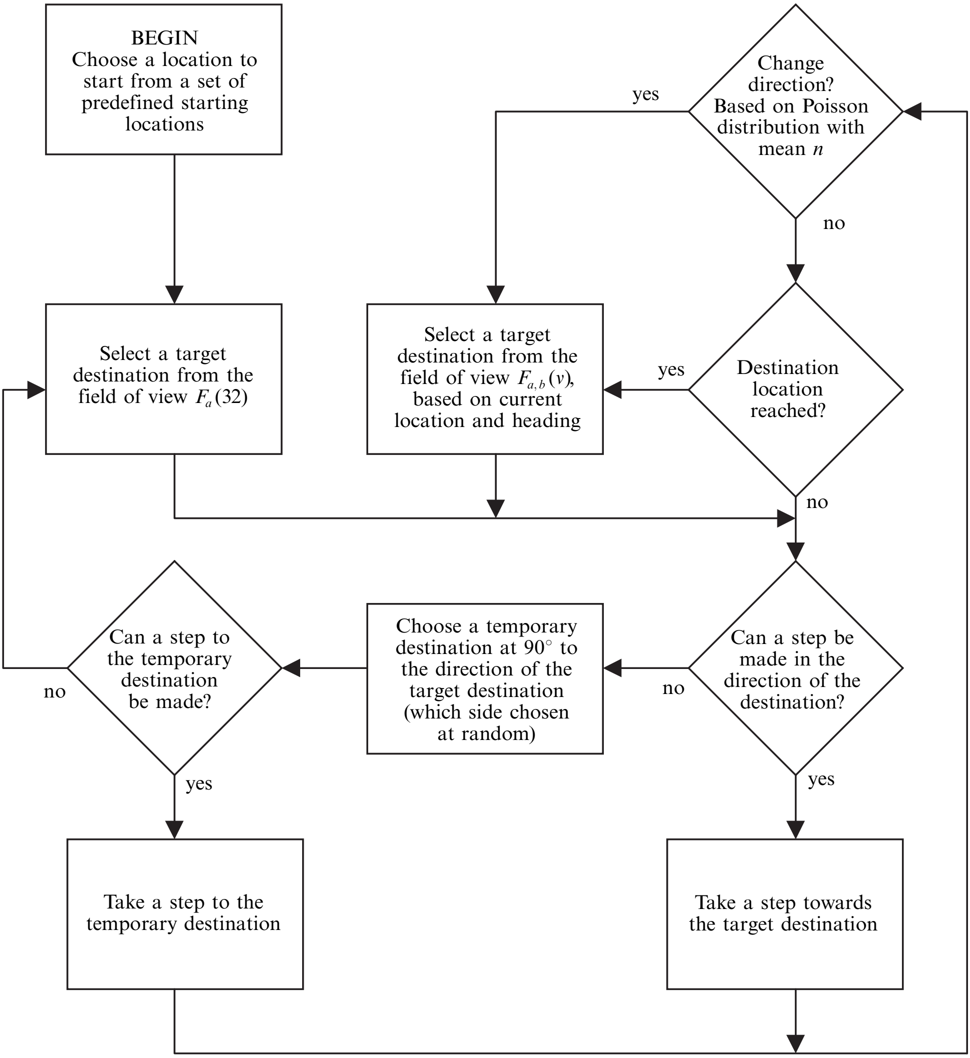

Figure 5. The complete EVA (exosomatic visual architecture) agent-decision process.

---

(3) Note that applying this second rule in isolation also led to gridlock in our experiments.

Application

In this section we describe the application of the EVA-based system to the Tate Britain Gallery, Millbank, London. Our choice of this building is based on several factors. First, detailed human observation data are available for the gallery (Hillier et al, 1996), and there is established use of the data in the literature (for example, Batty et al, 1998) which permits comparison with other methodologies. The Tate thus represents an 'Iris' dataset (Fisher, 1936) for human-movement studies. Further information is also available from the observational study, for example showing that people move neither according to a plan given to visitors, the mechanism previously thought to be operating by the gallery's managers, nor seemingly randomly (we prove that the movement is not random).

The agent-decision process was tested across the search space of the two variables—number of steps n and field of view (v bins, converted to an angle α)—by releasing agents according to a Poisson distribution from the main entrance of the gallery at the same rate as people enter the gallery. Each agent is removed from the system after 1800 moves, whereas a real human returns to the entrance to leave, pauses to observe paintings, and so on. Hence, the experimental model lacks some credence (in addition, the agents are not removed if they return to the entrance and are allowed to reenter, which alters entry rate to some extent). In addition, as previously stated, for simplicity agent walking pace is not distributed but kept at the mean walking pace of adult humans. Clearly, these differences between experimental method and reality will merely affect the level of relationship we observe, but the consistency of the rule set will be maintained.

As controls, two extra types of agents were also applied to the layout in separate experiments. First, unsighted agents, which take an average of n steps, and then make a turn of $\pm\alpha/2$ from their current heading. Second, particle agents, which take an average of n steps before taking a new random direction; they also take a new random direction on collision with a wall or another agent, and thus, in the infinite-step case, converge to Brownian behaviour. For each type of agent, experiments were run for 10 000 agent moves, releasing one agent into the system from the entrance every six agent moves. Trails of the agent paths were recorded, and a typical trail image for sighted agents is shown in figure 6 (see over), next to the actual traces obtained by following people for 10 minutes after their entry to the gallery. Although these are interesting to compare, it is the movement levels per room that are of primary importance, as they can give a quantitative assessment of how well our system actually reproduces human-pedestrian behaviour.

The correlation with observed movement was calculated by taking the linear R^2 correlation coefficient of log–log data of the observed room movement (from Hillier et al, 1996) against the simulated room movement (over 10 000 agent moves). Logarithmic scales were used to avoid the creation of strongly influential data points around the entrance hall and surroundings because of near-exponential thinning from the entrance (when log scales are applied, the data are approximately normally distributed, as is required for linear regression comparison). In both real and simulated cases, the room movement rates were calculated by taking entrance and exit counts for the 54 rooms and corridors on the ground floor of the Tate Britain Gallery, Millbank. In addition, two extra measures were calculated: *total coverage* is the number of rooms which were visited by at least one agent during the experiment, and thus gives the number of data points available from which to calculate the correlation coefficient; and *per agent cummulative isovist* the mean fraction of building area that could have been viewed by an agent during its visit had it had 360° vision (that is, the cummulative isovist area for the agent during its 1800-step lifespan, expressed as a fraction of the

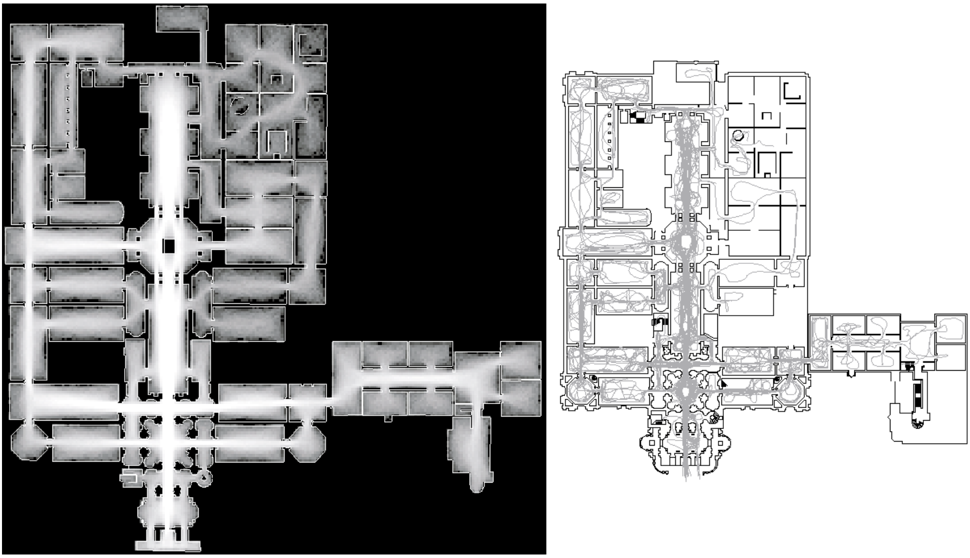
(a) (b)

Figure 6. (a) Trails left by agents walking through the Tate Britain Gallery, Millbank. As each agent steps on a grid square it increments a counter. Black areas have low counts and white areas have high counts. (b) Actual movement traces for 19 people followed for the first ten minutes of their visit to the gallery (reproduced from Hillier et al, 1996, page 15).

total floor area). The cummulative isovist gives an idea of how optimized an agent is in terms of explorative ability. The use of 360° isovists rather than actual area observed, is necessary as in general, the amount of the gallery which the agent sees is based on its angle of view α.

The key findings are summarised in table 1; notably, that the best combination of parameters discovered, a mean of 3 steps with a field of view of 170°, gives a correlation coefficient $R^2 = 0.76$ for observed human-movement levels against simulated movement (the standard error was calculated from the results of five separate experiments, the results of a typical experiment are shown in figure 7). Of course, the question arises as to whether any movement rule at all would give this sort of correlation. The answer is no, the use of particle diffusion (particle, ∞ steps in table 1) from the entrance gives $R^2 = 0.46$. Obviously, to some degree humans are simply diffusing from the entrance. However, this figure is nowhere near as convincing as

Table 1. The correlation of observed movement against agent-movement rates, per agent cummulative isovist and total coverage (out of 54 rooms) for different types of agents. Sighted agents have $\alpha = 170°$ angle of view.

Type	Steps n	Correlation R^2	Per agent cummulative isovist	Total coverage (rooms)
Particle	0	0.23 ± 0.01	0.073 ± 0.001	5.5 ± 0.5
	3	0.32 ± 0.01	0.116 ± 0.002	19.5 ± 0.5
	∞	0.46 ± 0.02	0.199 ± 0.004	48.0 ± 0.5
Unsighted	0	0.41 ± 0.01	0.164 ± 0.004	33.5 ± 2.0
	3	0.41 ± 0.02	0.227 ± 0.003	52.0 ± 0.5
	∞	0.33 ± 0.01	0.224 ± 0.004	54.0 ± 0.0
Sighted	0	0.63 ± 0.02	0.370 ± 0.002	50.0 ± 2.0
	3	0.76 ± 0.01	0.488 ± 0.002	54.0 ± 0.0
	∞	0.45 ± 0.01	0.529 ± 0.004	54.0 ± 0.0

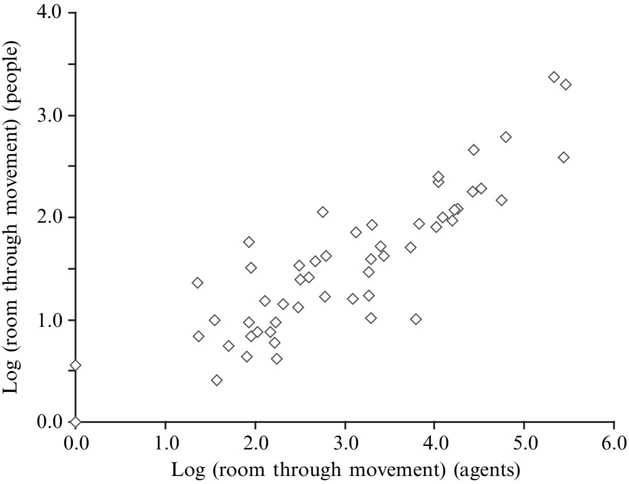

Figure 7. Observed room movement rates against agent room movement rates, plotted on logarithmic scales, for 3-step sighted agents with $\alpha \simeq 170^\circ$ angle of view. Although the correlation is good, the agent movement is generally higher throughout the building, indicating that further calibration of our model is necessary.

that for agents with vision. The next pertinent question is, would the agents do just as well if their sight were removed? If we leave the other rules the same, so that agents no longer choose a visible location at random, but a direction at random, the answer is a clear no: $R^2 = 0.41$ (unsighted, 3 steps in table 1). So it appears that the hypothesis is confirmed. Because the EVA system rule-base weights higher viewable area in a certain direction with higher probability of that direction being chosen, EVA agents have a higher chance of moving in the direction of higher walkable surface, and those agents are the most successful at reproducing aggregate human-pedestrian behaviour.

It is also interesting to look at the results across the parameter space. Figure 8 (over) shows the experimental results of varying the parameters field of view α and number of steps n. In particular, note that 0 steps (that is, choose a new location every step) and higher numbers of steps (that is, continue towards your chosen location until you reach it—see table 1), do not correlate nearly as well. The per agent cummulative isovist (fraction of the building possibly viewed) drops as the step is count reduced, and hence the 0-step rule is a poor heuristic for full exploration of a space. We might expect this result if humans are optimised explorers. However, at large numbers of steps the amount of building observed does not drop significantly, so we can only assume that there is some other cost involved with this heuristic.

If a single agent were released within the system then the simulation would be a Markov process, as the probability function of next moves depends entirely on the current state (position, direction, steps taken in this direction) of the agent. It is not clear that it remains a Markov process for a multiplicity of agents, as the interaction through collisions between agents allows a degree of positive feedback into the system. However, it seems likely that for our implementation, because of the marked consistency of the results, we are seeing an essentially Markov process, which is perhaps perturbed by agent collisions. Of course, we can test to see at what stage the collision behaviour becomes a significant factor by loading the system with different population densities. However, in our experiments, loading to extreme merely blocked the entrance, leaving the rest of the system to maintain its original movement levels—thus, even entry rates of 20 agents per agent move (over 100 times more than the actual Tate entry rates) gave a correlation with human movement of $R^2 = 0.72$.

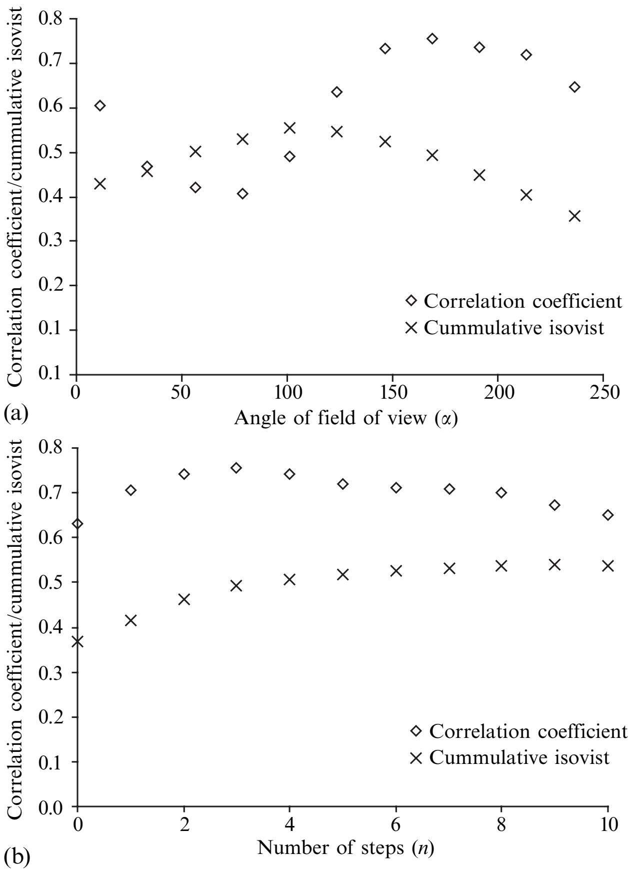

Figure 8. Correlation with observed movement and per agent cummulative isovist against angle of field of view for $n = 3$ step agents. The peak correlation is with a field of view of ~170° (our system does not permit exactly 180° field of view). There is a curious upturn at the lower end of the field-of-view experiments, which might suggest that two operations are at work in the human—one tending to turn out at near 90°, the other tending to continue in a straight line. Above ~240° the total coverage drops rapidly to zero, making direct correlation impossible. (b) Correlation with observed movement and per agent cummulative isovist against number of steps for agents with $\alpha \simeq 170°$ field of view. Both the correlation and cummulative isovist tail off rapidly towards the 0-step case.

Conclusion

Our question, perhaps trivial, has been: if configuration is important, how exactly does it affect the way people move around the world? We proposed to answer that question with an agent simulation in which the only movement strategy possible is dependent on the configuration of a space. We showed that it is possible to emulate human-movement patterns within a building environment by encoding Gibson's principle of affordance in the context of natural movement. We found that the actual level of aggregate movement observed within the Tate Gallery, Millbank, can be reproduced with a correlation coefficient of $R^2 = 0.76$ if we apply only elementary guidance rules. These rules were that destinations may only be chosen from a ~170° visual field from the current heading, and that the destination is reassessed every three steps. Whether or not this result is reproducible for other buildings has not been resolved, although the consistency of the results within this single building is encouraging. Large-scale experiments with many hundreds of visual agents have been made possible by the use of an exosomatic visual architecture (EVA), in which sets of locations visible from each point within the configuration are held in a look-up table accessed at runtime.

The strength of correlation found would seem to suggest that it is imperative to include a natural-movement rule set for any complete model of pedestrian movement at a mesoscopic or microscopic level. Interestingly, this rule set does not have to be complicated by complex theories involving higher representational models. In further work presented elsewhere (Penn and Turner, 2002), we have compared agents constrained by the movement rules applied here with agents which can identify junctions (based on visibility-graph measures). No such superior knowledge of configuration was found to be necessary and, as shown herein, a stochastic process based on assessing available walkable surface through vision does provide agent behaviour that corresponds to human behaviour. However, there are caveats to the exclusive use of direct visual perception that must be considered. Throughout, we have insisted that natural movement be visually guided (indeed, we have shown that removing vision severly impairs the ability of our agents both to explore and to reproduce human behaviour in our sample environment), whereas it is obvious that blind and visually impaired people do explore and navigate around environments and, it appears, employ cognitive maps in a similar manner to able-sighted people (Golledge et al, 1996). Further, we have formulated natural movement based on pedestrian 'steps', and this would differ if we were to consider other modes of transport such as wheelchair, bicycle, or car. Therefore, we should be clear that we are not trying to provide a complete model of all human-pedestrian behaviour. And the results, though good, also show that a direct perceptual system does not suffice on its own. In particular, we stated that the model should incorporate a socioeconomic framework and physical constraints where necessary. Of these, we have included modest physical constraints: the agents cannot co-occupy spaces; and a couple of rules resulting from Helbing and Molnár's (1997) granular-physics model were incorporated in order to avoid gridlock. Perhaps more importantly, the mental model also needs to be completed. Our primary purpose has been to show that it is possible to generate much of human-pedestrian behaviour from a limited rule set, but only so much behaviour can possibly be accounted for from our stochastic methods. We have concentrated on exploratory behaviour in public buildings where such behaviour would be expected. A notable omission is the fact that our agents cannot even return to an entrance. For such behaviour, two cognitive elements have to be added: the notion that there is an entrance to be returned to; and the ability to plan a path to that entrance, by whatever means—for example, place recognition, or topological or metric navigation, as detailed by Trullier et al (1997). It would seem sensible to consider such options as path integration, although it should be noted that our implementation is directly suited to a cognitive-map approach. For example, Mallot et al (1999) describe a view-graph approach, in which an agent may represent key locations on a visibility graph with remembered views, thus covering both place recognition and topological navigation. Again, for a complete description of human behaviour, such a formalism would have to be extended so that it included place recognition in blind and visually impaired people, or those who may use different navigational strategies. Although even after these shortcomings, direct perception has still not been exhausted; it will be interesting to find out that effect of the addition of paintings on agent behaviour. In some senses, of course, we have applied sleight of hand: our correlation is with overall movement patterns, and does not consider individual agent paths. Kerridge et al (2001) discuss what is required of an agent-based model of pedestrian behaviour, in particular that agents should correspond with empirical data both quantitatively *and* qualitatively, where the qualitative observations are of individual movement patterns. How qualitative performance should be graded is still an issue, but our model must also be tested against such criteria as and when they are developed.

When Reynolds (1987) introduced Boids, he revolutionised agent simulation of previously noncomputable flocks and herds by utilising the fact that agents can quickly access information about other agents without needing a full-vision architecture. In an analogous method to the way in which Boids interrogate each other, we have proposed that agents interrogate the environment by sampling visual information prestored as a look-up table. This exosomatic visual architecture permits simplified agents guided by visual affordances, and this study has shown that it is possible to derive simple laws of 'natural movement' which correlate well with observed human behaviour. Further study is required to test whether the results hold in other buildings. However, our investigation has shown that, by applying agent modelling with direct-perception rule sets, we can extend our understanding of how people behave in building environments.

Acknowledgements. The authors would like to thank Chiron Mottam and Ruth Conroy Dalton, as well as Dr Alex Scott (mathematical insights), Professor Bill Hillier, and Space Syntax Limited (datasets and analysis), and the particularly incisive reviewers (in the sense of sharp-witted, rather than cutting). The work presented here was partially funded by EPSRC Platform grant GR/N21376/01.

References

al Gadhi S A H, Mahmassani H S, Herman R, 2001, "A speed – concentration relation for bi-directional crowd movements with strong interaction", in *Pedestrian and Evacuation Dynamics* Eds M Schreckenberg, S Sharma (Springer, Heidelberg) pp 3 – 20

Bachelder I A, Waxman A M, 1994, "Mobile robot visual mapping and localization: a view-based neurocomputational architecture that emulates hippocampal place learning" *Neural Networks* **7** 1083 – 1099

Batty M, 2001, "Agent-based pedestrian modeling" *Environment and Planning B: Planning and Design* **28** 321 – 326

Batty M, Conroy R, Hillier B, Jiang B, Desyllas J, Mottram C, Penn A, Smith A, Turner A, 1998, "The virtual Tate", WP 5, Centre for Advanced Spatial Analysis, University College London, London

Beckman R J (Ed.), 1997, "The Dallas – Ft. Worth case study", unclassified report LA-UR-97-4502, Los Alamos National Laboratory, Los Alamos, NM

Benedikt M L, 1979, "To take hold of space: isovists and isovist fields" *Environment and Planning B: Planning and Design* **6** 47 – 65

Berger J, 1972 *Ways of Seeing* (BBC & Penguin Books, London)

Blue V J, Adler J L, 1998, "Emergent fundamental pedestrian flows from cellular automata microsimulation" *Transportation Research Record* number 1644, 29 – 36

Borgers A, Timmermans H, 1986, "A model of pedestrian route choice and demand for retail facilities within inner-city shopping areas" *Geographical Analysis* **18** 115 – 128

Braaksma J P, Cook W J, 1980, "Human orientation in transportation terminals" *Transportation Engineering Journal* **106**(TE2) 189 – 203

Casti J L, 1998 *Would-Be Worlds: How Simulation Is Changing the Frontiers of Science* (John Wiley, New York)

Conroy Dalton R A, 2001, "The secret is to follow your nose", in *Proceedings of the 3rd International Symposium on Space Syntax* Georgia Institute of Technology, Atlanta, Georgia, pp 47.1 – 47.14

Daganzo C F, 1994, "Cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory" *Transportation Research Part B: Methodological* **28** 269 – 287

Epstein J M, Axtell R, 1996 *Growing Artificial Societies, Social Science from the Bottom Up* (Brookings Institution, Washington, DC)

Fisher R A, 1936, "The use of multiple measurements in taxonomic problems" *Annals of Eugenics* **7** 179 – 188

Galea E R, Owen M, Lawrence P J, 1996, "Computer modelling of human behaviour in aircraft fire accidents" *Toxicology* **115** 63 – 78

Gärling T, Gärling E, 1988, "Distance minimization in downtown pedestrian shopping" *Environment and Planning A* **20** 547 – 554

Gibson J J, 1979 *The Ecological Approach to Visual Perception* (Houghton Mifflin, Boston, MA)

Golledge R G, 1995, "Path selection and route preference in human navigation: a progress report", in *Spatial Information Theory—Lecture Notes in Computer Science* number 988, pp 207 – 222

Golledge R G, Klatzky R L, Loomis J M, 1996, "Cognitive mapping and wayfinding by adults without vision", in *The Construction of Cognitive Maps* Ed. J Portugali (Kluwer Academic, Dordrecht) pp 215 – 246

Helbing D, Molnár P, 1997, "Self-organization phenomena in pedestrian crowds", in *Self Organization of Complex Structures: From Individual to Collective Dynamics* Eds F Schweitzer, H Haken (Gordon and Breach, Amsterdam) pp 569 – 577

Helbing D, Molnár P, Farkas I J, Bolay K, 2001, "Self-organizing pedestrian movement" *Environment and Planning B: Planning and Design* **28** 361 – 383

Hillier B, 1996 *Space is the Machine* (Cambridge University Press, Cambridge)

Hillier B, Hanson J, 1984 *The Social Logic of Space* (Cambridge University Press, Cambridge)

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

Hilier B, Major M D, Desyllas J, Karimi K, Campos B, Stoner T, 1996, "Tate Gallery, Millbank: a study of the existing layout and new masterplan proposal", technical report, Bartlett School of Graduate Studies, University College London, London

Hoogendoorn S P, Bovy P H L, 2000, "Continuum modeling of multiclass traffic flow" *Transportation Research Part B: Methodological* **34** 123 – 146

Hoogendoorn S P, Bovy P H L, Daamen W, 2001, "Microscopic pedestrian wayfinding and dynamics modelling", in *Pedestrian and Evacuation Dynamics* Eds M Schreckenberg, S Sharma (Springer, Heidelberg) pp 123 – 154

Kerridge J, Hine J, Wigan M, 2001, "Agent-based modelling of pedestrian movements: the questions that need to be asked and answered" *Environment and Planning B: Planning and Design* **28** 327 – 341

Kuipers B J, Byun Y T, 1991, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations" *Journal of Robotics and Autonomous Systems* **8** 47 – 63

Kurose S, Borgers A W J, Timmermans H J P, 2001, "Classifying pedestrian shopping behaviour according to implied heuristic choice rules" *Environment and Planning B: Planning and Design* **28** 405 – 418

Laumand J, 1983, "Model structuring and concept recognition: two aspects of learning for a mobile robot", in *Proceedings of the International Joint Conference on Artificial Intelligence* (William Kaufman, Los Altos, CA) pp 839 – 841

Lighthill M H, Whitham G B, 1955, "On kinematic waves II: a theory of traffic flow on long, crowded roads" *Processings of The Royal Society of London, Series A* **229** 317 – 345

Mallot H A, Gillner S, Steck S D, Franz M O, 1999, "Recognition-triggered response and the view-graph approach to spatial cognition", in *COSIT '99—Lecture Notes in Computer Science, Number 1661* (Springer, Berlin) pp 367 – 380

Mottram C, Conroy R, Turner A, Penn A, 1999, "Virtual beings: emergence of population level movement behaviours from individual rule-sets", in *Proceedings of the 2nd International Symposium on Space Syntax, Volume 3* Universidad de Brasil, Brasilia, Brazil

Murphy R R, 1999, "Case studies of applying Gibson's ecological approach to mobile robots" *IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans* **29** 105 – 111

Nagel K, Barrett C L, Rickert M, 1996, "Parallel traffic microsimulation by cellular automata and application for large-scale transportation modeling", unclassified report LA-UR-97-50, Los Alamos National Laboratory, Los Alamos, NM

Neisser U, 1994, "Multiple systems—a new approach to cognitive theory" *European Journal of Cognitive Psychology* **6** 225 – 241

O'Keefe J, Nadel L, 1978 *The Hippocampus as a Cognitive Map* (Clarendon, Oxford)

Ooi T L, Wu B, He Z J, 2001, "Distance determined by the angular declination below the horizon" *Nature* **414** 197 – 200

Oommen B J, Iyengar S S, Rao N S V, Kashyap R L, 1987, "Robot navigation in unknown terrains using learned visibility graphs. 1: the disjoint convex-obstacle case" *IEEE Journal of Robotics and Automation* **3** 672 – 681

Penn A, 2001, "Space syntax and spatial cognition, or why the axial line?", in *Proceedings of the 3rd International Symposium on Space Syntax* Georgia Institute of Technology, Atlanta, Georgia, pp 11.1 – 11.16

Penn A, Dalton N, 1994, "The architecture of society: stochastic simulation of urban movement", in *Simulating Societies: The Computer Simulation of Social Phenomena* Eds N Gilbert, J Doran (UCL Press, London) pp 85 – 125

Penn A, Turner A, 2002, "Space syntax based agent models", in *Pedestrian and Evacuation Dynamics* Eds M Schreckenberg, S Sharma (Springer, Heidelberg) pp 99 – 114

Penn A, Turner A, forthcoming, "A system and method for intelligent modelling of public spaces", Patent Application GB0020850.4, UK Patent Office (filed)

Peponis J, Zimring C, Choi Y K, 1990, "Finding the building in wayfinding" *Environment and Behavior* **22** 555 – 590

Reynolds C W, 1987, "Flocks, herds, and schools: a distributed behavioral model" *Computer Graphics* **21**(4) 25 – 34

Steadman P, 1973, "Graph theoretic representation of architectural arrangement" *Architectural Research and Teaching* **2**(3) 161 – 172

Sutherland D H, Kaufman K R, Moitoza J R, 1994, "Kinematics of normal human walking", in *Human Walking* Eds J Rose, J G Gamble (Williams and Wilkins, Baltimore, MD) pp 23 – 44

Terzopoulos D, 1999, "Visual modelling for computer animation: graphics with a vision" *Computer Graphics* **33** 42 – 45

Thomas G, Donikian S, 2000, "Modelling virtual cities dedicated to behavioural animation" *Computer Graphics Forum* **19**(3) C71 – C80

Trullier O, Wiener S I, Berthoz A, Meyer J A, 1997, "Biologically based artificial navigation systems: review and prospects" *Progress in Neurobiology* **51** 483 – 544

Turner A, Doxa M, O'Sullivan D, Penn A, 2001, "From isovists to visibility graphs: a methodology for the analysis of architectural space" *Environment and Planning B: Planning and Design* **28** 103 – 121

ul Haq S, 2001, "Just down the road a piece: the development of topological knowledge in building layouts", in *Proceedings of the 3rd International Symposium on Space Syntax* Georgia Institute of Technology, Atlanta, Georgia, pp 48.1 – 48.17