A Multi-Agent Based Fitness Function for Evolutionary Architecture

Richard Holden and Angelo Cangelosi
Interactive and Adaptive Systems, University of Plymouth, Plymouth, PL4 8AA
rholden@plymouth.ac.uk, acangelosi@plymouth.ac.uk

***Abstract*—*Previous Multi-Agent System (MAS) methods for modelling discrete pedestrian dynamics have been shown to be brittle in the face of increasing density. We present an alternative method, which incorporates the novel idea of a local-to-global linkage for path discovery and add a heuristic to control critical information. This method has been successful in reproducing high-level crowd patterns and we present a preliminary approach to the evolutionary design of building architecture using a crowd model based on this method. We present a genotype $\mapsto$ fitness mapping and results of evolutionary runs in an exit sign layout scenario. We discuss important issues and how these may affect future research directions.*

I. Introduction

Evolutionary architecture

The Evolutionary Architecture of built form, as envisioned by Frazer [1], outlines the potential benefit of nature-inspired computing for building design. Since this work, others [2], [4] have attempted to implement evolutionary techniques to this end. These techniques have emphasized an exploratative approach to architecture where aesthetically pleasing forms are attainable through evolution with a human-in-the-loop approach to evolutionary evaluation. This stems from an underlying belief that because modern materials are so robust, form now rules over function [4].

However, architectural form must naturally contain some element of function. One important building function are characteristics of its spatial configuration, which impinge on the movement of crowds of people. In order to motivate and broaden the concept of evolutionary architecture our longer-term aims are to contribute by developing fitness functions to capture complex dynamics observed in real crowds and elucidate dynamics of interactions between pedestrians and between pedestrians and the architectural environment.

Computational crowd dynamics

Various MAS-based models of crowd movement have been proposed. Definitions of *motion* in these models can be either continuous [5] or discrete [6], [7]. In terms of modelling crowds in labyrinthine spaces, models based on the continuous concept of movement are ruled out due to problematic local minima in the space. This is unfortunate because local dynamics in such models promise to capture the emergent properties of real crowds well [5].

The question then arises as to which discrete method is the most appropriate to use—i.e., which method can best be used to represent local pedestrian-pedestrian and also pedestrian-obstacle interactions in the broader context of longer-range motion towards destinations in labyrinthine space?

In Section II we briefly introduce some background definitions and properties of discrete space. We discuss problematic properties of 2-lattices, which fundamentally constrain the robustness of discrete updates in other pedestrian models. We present a general method based on an inclusive search, which produces more robust behaviour, but which still suffers from inescapable underlying statistical features of the lattice. We then present a heuristic, which exploits statistical information in the spatial network, in order to avoid catastrophic failure. This information can be used to determine how much information in the space should be processed when searching for a local to global link—between the current location and a longer-range destination.

We then demonstrate how these kinds of techniques may be embedded within an evolutionary approach to the design of spatial configuration. We do this using an example application of exit-sign layout in a real building scenario. We specify a model of pedestrian escape using the techniques outlined. We then use this model in a genotype to fitness mapping in order to show how evolution may be coupled with MAS to explore spatial forms with a built-in safety function. We discuss this preliminary work and how it might be extended.

II. Definitions and Properties of Discrete Space

2-lattice graphs

Many MAS models of crowd dynamics implicity use the 2-lattice as an underlying spatial representation. This is a graph G consisting of a vertex set V and an edge set E. The set V represents allowable points of spatial occupation by pedestrians and E links the vertices so that movement between these points may occur. We present examples of these kinds of spaces, which can be constructed from different neighbourhoods (see Fig 1). A neighbourhood is denoted by Γ. For example, a Moore neighbourhood $\Gamma^1(v)$, which has a parameter $k = 8$ to define links to neighbouring lattice vertices—this is the *immediate* neighbourhood of a given vertex v. Movement is therefore characterized by discrete steps in the lattice where pedestrians P_i make moves into $\Gamma^1(v_i)$. Where no move is made the current vertex $\Gamma^0(v) = \{v\}$ is maintained at the next time step $t = t + 1$. In this most simple case a pedestrian will move in incremental steps of length 1 and over time define a graph walk. Many discrete model specify these kinds of dynamics and it is this general behaviour, which is of interest to us.

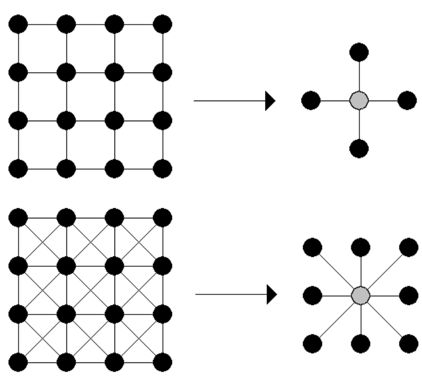

Fig. 1. Example 2-lattices with Moore and von Neumann neighbourhoods.

Statistical properties of graphs

When designing graph-based models of pedestrian movement it is crucially important to appreciate well known statistical properties in boolean 2-lattices. The reasons boolean properties are important is because a pedestrian model is an instance of a boolean lattice—pedestrians either occupy a cell or they don't. Interpreted this way, given some probability ρ, a given pedestrian P_i can occupy v_i. Then it is a general property with such definitions that percolation across Γ^i will occur at a critical point ρ_c. A more intuitive picture is to imagine that at certain pedestrian densities, defined by rho, large clusters of pedestrians may link to form a giant cluster. Such kinds of giant clusters appear in von Neumann lattices $\rho_c = 0.59$ and have been used to provide analogies to many complex systems [3]. This behaviour also has important implications for search algorithms used in discrete models of pedestrian dynamics.

Implications

Some form of search in a discrete model will inform the possible best moves to make from the perspective of a pedestrian P. For example, in a simple model of pedestrian lane formation it has been shown that simple rule-based decisions can be made on the basis of local searches for locally optimal moves, which depend on the flow of other pedestrians in a local neighbourhood [7]. At certain densities ($\rho \approx 0.2$) these models have successfully produced higher-level patterns of lane formation. However, other work has also shown that by increasing $\rho > 0.2$ catastrophic failure occurs where lane formations suddenly collapse at critical points in the underlying 2-lattice [8]. The implication is that at certain densities local search is crippled and graph percolation invades the dynamics. We have argued that a higher-level and more inclusive search is required, based on incremental neighbourhood search [9].

III. Robust Search

Simple neighbourhood search

The neighbourhood $\Gamma(S)$ of a connected subgraph S consists of vertices adjacent to, but not including $v_i \in S$. Where $S = \Gamma(v)$, $\Gamma(S) = \Gamma(\Gamma(v)) = \Gamma^2(v)$, the 2^{nd} neighbourhood of v and generally $\Gamma(\Gamma^{i-1}(v)) = \Gamma^i(v)$ defines the i^{th} neighbourhood of v. In a 2-lattice constructed from Moore neighbourhoods $|\Gamma^i(v)|$, from the centre vertex of the lattice, where $i = \{2, 3, 4, \ldots n\}$, increases according to $|\Gamma^{i-1}(v)| + 8$, where $|\Gamma^1| = k = 8$. The resulting sequence is linear.

Previous work has shown that an inclusive search algorithm, implemented as a Pulse Coupled Neural network can solve complex maze tasks where inputs to the network and neuron parameters are set appropriately [10]. In a related approach we have used the following equations to guide a pedestrian from a vertex origin v_o to a vertex destination v_d:

$$F_{ij}(t) = S_{ij} \tag{1}$$

$$L_{ij}(t) = e^{-\delta_L \Delta t} L_{ij}(t-1) + \sum_{kl} M_{ijkl} Y_{kl(t-1)} \tag{2}$$

$$U_{ij}(t) = F_{ij} + L_{ij} \tag{3}$$

$$Y_{ij}(t) = \begin{cases} 1 & \text{if } U_{ij}(t) > \Theta_{ij}(t) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

$$\Theta_{ij} = e^{-\gamma_\Theta \Delta t} \Theta_{ij}(t-1) + V_\Theta Y_{ij}(t) \tag{5}$$

where, F_{ij}, L_{ij}, U_{ij}, Y_{ij}, Θ_{ij} define the *Feeding Field, Linking Field, Internal Activity, Output Field and the Dynamic Threshold*, respectively.

Using this approach an optimal path from v_o to v_d can be found. Firstly, a wave travels away from v_o until the wave reaches v_d. Then we can use time-dependant variables to traceback the optimal route [9]. We demonstrate this in a

$17 * 17$ network of neurons. We set $S_{2,11} = 1$ and then iterate the network. Behaviour produces autowave behaviour, where in this case network activity (number of pulsing neurons at time t) follows the simple linear sequence described above (see Figure 2). We call this sequence the *time series signature* of autowave search.

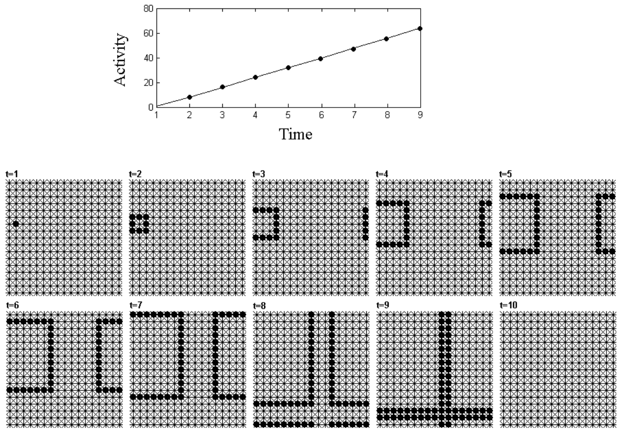

Fig. 2. Simple search sequences.

Search constraints

As mentioned above search is used to find the best move to make in the face of other pedestrians who occupy the lattice. In order to represent the presence, to pedestrian P_j, of ‘other’ pedestrians P_i, we fix the output of a given neuron ($Y_i = 0$), corresponding to the vertex location of P_i. This ensures that the expanding autowave travels *around* these pedestrians and the path returned by the trace back cannot contain v_i in the $v_o \mapsto v_d$ link [9], [11]. In this way the pedestrian moves towards a destination vertex by avoiding other pedestrians on the way.

We visualise this search at various densities by plotting the time series signature (see Fig 3). We can see how at higher densities ($\rho = 0.7$) little search is done, and at low densities ($\rho = 0.1,\ \rho = 0.3$) search produces more or less similar shapes to that of the simple sequence in Fig 2. At mid-range densities ($\rho = 0.5$) search becomes more complex and the expansive autowave travels for longer periods, lasting for 43 time steps.

We can appreciate the issue of robustness if we consider example $v_o \mapsto v_d$ linkage in two different regimes (see Fig 4). We choose $\rho = 0.4$ and $\rho = 0.59$, which is in a critical regime. We observe below the critical point that search is robust—many paths are returned and the trace-back is very *fat*. In this regime, if a few other pedestrians were to occupy a vertex in the link, then we would be confident that the link would not be broken. However, in the critical regime we would expect the opposite—almost certain breakage in the link $v_o \mapsto v_d$ because the trace back paths are very *thin*.

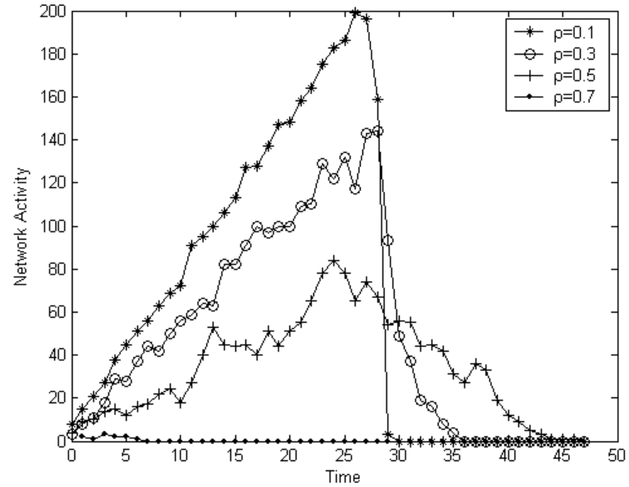

Fig. 3. Various search signatures according to ρ

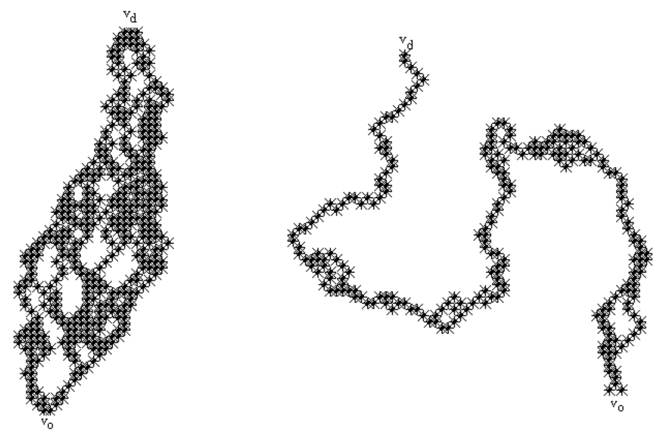

Fig. 4. Path robustness: left $\rho = 0.4$, right $\rho = 0.59$

Search heuristics

Direction dependant relaxation—Previous work has used the idea of a bias network input to allow the production of higher-level lane formations in abstract periodic spaces—in a bi-directional scenario where a population P^z is split in half ($P^n,\ P^s$) the constraints in autowave search for each pedestrian are constrained by the rule that each pedestrian P_i^z accounts for other pedestrians P_j^z in Γ_i^1, but where $\Gamma_i^{>1}$ inputs are applied to pedestrians traveling in the opposite direction $P_j^{n \leftrightarrow s}$, only [11]. This approach is adopted in the model used below, but is applied in a more complicated context where the population will have an exit choice based on exit sign information. This approach effectively relaxes the constraints and allows for higher-level pattern formation across much wider density ranges [8].

Time series heuristic—We have seen in previous sections that the network activity or time series information indicates what kind of regime the environment of a given pedestrian is like. What we are looking for is an *indicator* that the autowave front will die before reaching v_d. It is not our intention here to analyse the time series in any depth, but only to identify obvious features that may be used as fairly useful predictors. An obvious question to ask about the autowave search is *if the wave front shrinks to a value of 1, what is the probability of quick failure thereafter?* The only way to answer this is statistically, firstly by identifying whether or not the autowave from v_o to v_d contains a value of 1. If so, we refer to this as a *'one-type'* signature.

We present the fraction of failed searches that contain a value of 1 in the search signature (see Fig 5). We can see that the fraction of one-type signatures quickly increases at the critical ranges identified above. However, an important feature of this information is that above the critical range ($\rho \approx 0.55 - 0.59$), the fall off of one-type signatures in the population is relatively *smooth*—the majority of pedestrians still produce one-type signature search in a fairly large range ($\rho \approx 0.6 - 0.88$).

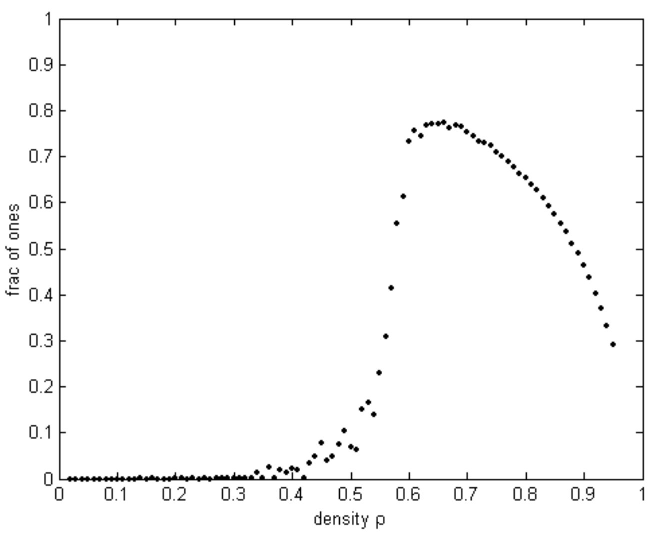

Fig. 5. One type signatures.

We also present interval histograms, which represent a critical range. Each histogram shows the conditional probability, given a value of 1 in the signature, that the time interval Δt elapses before failure occurs (see Fig 6). Results suggest that the presence of a one-type autowave is a very good indicator that wave failure is about to occur. We can use this information to determine at what range for i constraints are chosen in Γ^i relating to other pedestrians in the lattice. As the presence of a one in the signature is a good indication of wave failure and a break in $v_o \mapsto v_d$, we record t when activity is 1 and where $\Gamma^{>t}$, we relax the inputs to the network so that the wave can reach v_d and the link remain intact. Each pedestrian P_i given some destination v_{d_i} can then traverse the space by moving inside the dynamically changing $v_o \mapsto v_d$ link [11], which adapts to environmental change defined by other pedestrian motion.

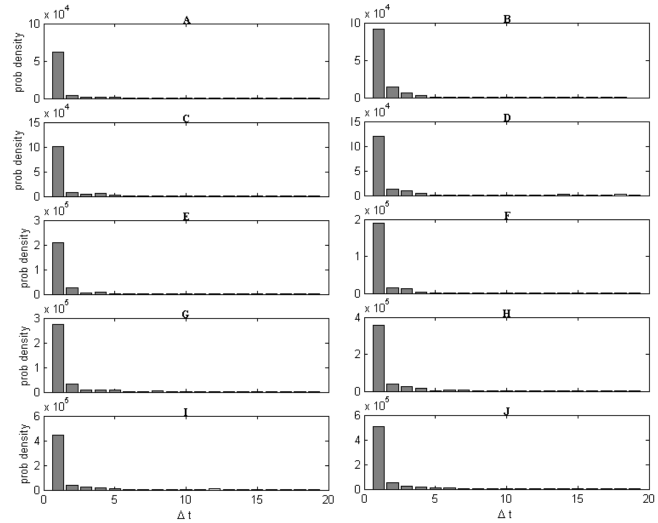

Fig. 6. One type signature failure.

IV. Evolutionary Approach

Genotype-phenotype

Portland Square (University of Plymouth) is a typical office environment. We intend to evolve exit sign locations and we constrain possible locations to corridor sections extracted from the original floor-plan (see Fig 7). The subset is mapped to a 1-D array of length $L = 209$, because there are 209 vertices in the extracted subspace. We constrain the number of exit signs to 30 and 1 byte for each and therefore length $l = 30 \times 8 = 240$—30 bytes decoded in the range $0 - 208$. Each byte corresponds to an exit sign location, free to be placed anywhere on the corridor subspace mapped back from the $1d$ array. An extra bit after each byte codes for either 'Exit 1' or 'Exit 2'. The phenotype is therefore an arrangement of exit-signs chosen freely within the constraints imposed by the genotype.

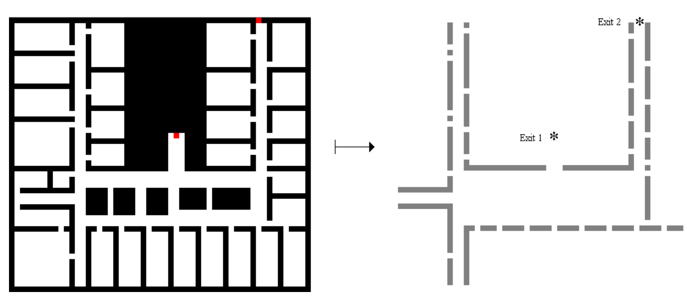

Fig. 7. Configuration and extracted exit-sign subspace.

Phenotype-fitness

We place a number of pedestrians P_i at v_i constrained so as 2 pedestrians are placed to each office or 'room' subspace. Each origin vertex will have a characteristic path length from either exit vertex v_1^e or v_2^e. We denote this path length by φ_i. Each pedestrian begins movement in a random walk and when an exit sign is discovered in Γ^1, then $v_d = v_1^e$ or $v_d = v_2^e$ depending on the corresponding bit in the genotype. On exit sigh discovery, pedestrian P_i then traverses the space between it's current location and the destination v_d, using the methods outlined above. We record the time taken from the onset of the simulation to the time P_i reaches v_d and call this μ_i. Phenotype fitness is then calculated by:

$$f(x) = \frac{\sum_{i=1}^{N} \frac{\varphi_i}{\mu_i}}{N} \tag{6}$$

where the minimum possible time, φ_i, for a given pedestrian to escape is divided by the actual time taken μ_i and summed over the size of the population and normalised so $f(x) = [0, 1]$. For example, in the special case where each pedestrian moves to v_{d_i} in minimal time, then $f(x) = 1$ and were each pedestrian to take twice the minimal time $f(x) = 0.5$.

The fitness of a given phenotype then depends on the interaction between a pedestrian movement component and the exit sign locations, which are placed under genetic control and evolved with the use of a genetic algorithm (GA), details of which are published elsewhere [11].

Evolution

The idea behind the use of the GA is therefore to improve the exit-sign locations in order to increase the speed pedestrians can escape from random walks and traverse the space. We can visualise a phenotype by placing a line from a given exit sign to the intended exit specified genetically. We present two example phenotypes initialised with random genotypes (fig 8).

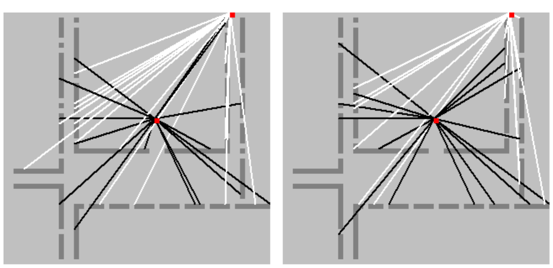

Fig. 8. Random phenotypes.

We run the GA for 8 trials and 500 generations (see fig 9). We can see how the fitness increases, and from two sample runs, how noisy the fitness mapping is, due to the random walk nature of the pedestrian model. However, evolution is still observed, although the power of selection appears quite weak, perhaps because the problem is too easy for the GA—we discuss this issue below when considering future work.

From the presentation of fittest phenotypes taken from the final generation we observe that the GA discovers a solution, which splits the exit signs into two spatial domains. We present two

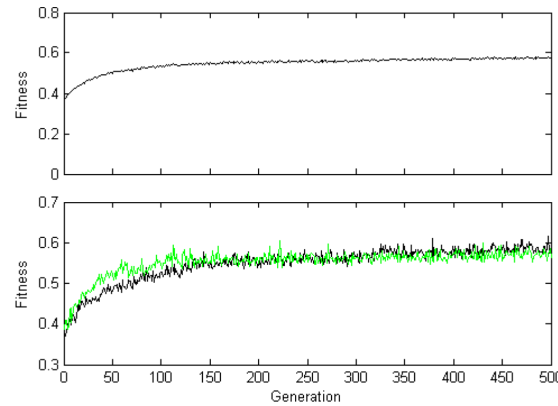

Fig. 9. Evolution: top = average, bottom = noisy evolution.

example phenotypes after evolution (see fig 10). This split ensures that conflict between pedestrian movement is reduced and flow improved. We also observe that the GA has optimised by distributing exit signs more widely than in the random cases.

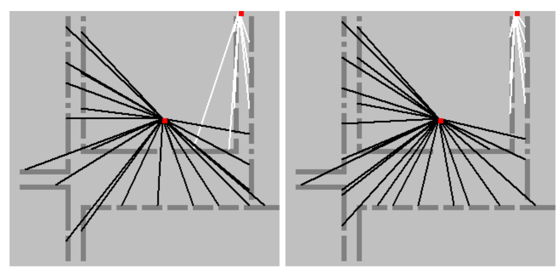

Fig. 10. Fit phenotypes.

V. Summary and Discussion

In this paper we presented general methods for discrete models of pedestrian dynamics, based on an inclusive search and a relaxation of inherent constraints of the 2-lattice. The methods allow a local to global link—between the current location and a longer-range destination, which is used to allow pedestrian agents to traverse a complex labyrinthine space. We then demonstrated the use of these techniques within an applied evolutionary approach to the design of spatial configuration. This demonstrates a coupling of MAS with evolution to explore spatial forms with a built-in safety function.

Phenotype-fitness

Concepts of space—The pedestrian model has been shown elsewhere to produce realistic pattern formations [8]. However, the efficacy of the methods and validity of their general use can only be gauged properly with more work. We believe the use of the 2-lattice is a severe constraint on *any* discrete MAS model of pedestrian flow. It was mentioned above that continuous models have shown quite promising pedestrian-pedestrian and pedestrian-obstacle interactions, but only in very simple configurations.

One fruitful avenue of research would be an investigation into combinations of discrete approaches to spatial representation with continuous models of pedestrian movement superimposed on the 2-lattice in order that catastrophic updates, which are a result of discrete moves be avoided. However, as part of this research, problems may result from the inherent multiple optima produced from the Moore-constructed lattice [11]. Continuous models rely on the concept of desired velocity [4], which is represented as a *single* vector quantity. In a situation where many optima are available the dynamics already produced by these models would be severely affected because many desired velocity vectors can be available at any one time. It may be that a complete reevaluation of underlying spatial concepts and algorithms in pedestrian and evacuation modelling is needed. Just because many models within the community use lattice spaces, this is no reason not to abandon the method completely.

The design of an algorithm capable of approximating continuous optimal paths in labyrinthine spaces would be beneficial in this resect, although to the best of our knowledge no appropriate algorithms are, as yet, available.

Fitness measures—More specific details such as the calculation of fitness are also important to the evolutionary outcome. In the above mapping we have calculated fitness based on the simple assumption of average escape time. Other important measures may include the amount of blockages inherent to certain spaces, and although this is implicit in average escape time, we could imagine more specific definitions of fitness in this respect, perhaps by recourse to measures of pedestrian clustering through the course of a simulation.

Generative architecture

We mentioned above how the current problem is perhaps too easy for a genetic algorithm. However, the longer-term goal is to evolve architectural configurations form scratch. Generative architectural algorithms have been developed in nature-inspired approaches based on the concept of social insects and stigmergy [12]. We see no reason why, given reliable and general pedestrian flow behaviours, generative algorithms may not be used to develop more complex architectures for human occupation, to be evaluated and explored with evolution. There may be many possible ways to develop architectures and an important emphasis in this respect would be the evolvability of various developmental mappings. Also, an architectural configuration is built from many inter-related components and current applications of the GA to combinatorial problems may be relevant.

VI. Conclusion

Our research suggests that paradigms within the Pedestrian Crowd Dynamics community are limited in several ways. In an attempt to overcome these limitations we have developed complementary methods, which produce more robust behaviour. We believe however, that before the realisation of evolved architecture is possible much more work in this area is needed. The current methods do produce some selection pressure, but many issues need to be addressed while designing genotype $\mapsto$ fitness mappings in this sense. We have briefly discussed some important issues in order to motivate future research.

Acknowledgements

The author gratefully acknowledges the University of Plymouth, School of Computing, which provided financial support during this research.

References

[1] A. Frazer, *Evolutionary Architecture*, London: Architectural Association, 1996.

[2] U.M. O'Reiley and P. Testa, "Representation in architectural design tools", 2000: *ACDM Adaptive Computing in Design and Manufacture Proceedings*, London: Klewer, 2000.

[3] R. Sole and B. Goodwin E. *Signs of Life: How Complexity Pervades Biology* New York, Basic Books, 1999.

[4] T. Broughton, A. Tan and P.S Coates, "The Use of Genetic Programming in Exploring 3D Design Worlds", 2000: *CAAD Futures Proceedings*, London: Klewer, 1997.

[5] D. Helbing, I. Farkas, and T. Vicsek, "Simulating Dynamical Features of Escape Panic", *Nature* ***407***, *487-490*, 2000.

[6] S. Gwynne, E. Galea, and M. Owen, "A Comparison of Predictions from the BuildingExodus Evacuation Model with Experimental Data",*1998 Human Behaviour in Fire Proceedings*, 711-720, 1998.

[7] V. Blue, J. Adler, "Cellular Automata Model of Emergent Collective Bi-Directional Pedestrian Dynamics", *2000 Artificail Life VII Proceedings*, 437-445, 2000.

[8] R. Holden and A. Cangelosi, "Limits to Locality in Cellular Automata Models of Pedestrian Dynamics", *2004 Interflam 10 Proceedings*, 761-766, 2000.

[9] R. Holden and A. Cangelosi, "A General Computationally Intelligent Model for Egress Simulation", *2004 Interflam 10 Proceedings*, 387-398, 2000.

[10] J. Caulfield and J. Kinser, "Finding the Shortest Route in the Shortest Time using Pcnn's", *IEEE Transactions on Neural Networks*, ***10***, 604-607, 2000.

[11] R. Holden, *Multi-Agent Fitness Functions for Evolutionary Architecture* PhD Thesis, University of Plymouth, UK 2005.

[12] E. Bonabeau, M. Dorigo and G. Theraulaz *Swarm Intelligence: from Natural to Artificial Systems* Oxford, England: Oxford University Press, 1999.