# Designing for Interest and Novelty

*Motivating Design Agents*

Rob Saunders and John S Gero
*The University of Sydney*

Key words: design agents, novelty, curiosity, learning

Abstract: This paper is concerned with the motivation of design agents to promote the exploration of design spaces. A general form of motivation common to designers is a curiosity to discover interesting designs. This paper presents computational models of interest and curiosity based on the detection of novelty. We illustrate the behaviour of our model of interest by developing a design agent that is motivated to explore the effects of emergent crowd behaviours on the performance of doorways.

## 1. INTRODUCTION

The search for interesting designs is a primary motivation for designers. Interesting designs provide information about the design task and allow the designer to learn in advance of a need to apply the knowledge. This type of curious self-directed learning plays an important role in the weaving together of problem finding and problem solving, within and between design sessions.

Studies of preference judgements in designers and non-designers show that the subjective determination of interestingness depends upon the previous experiences of the individual (Whitfield and Wiltshire, 1982; Purcell and Gero, 1992; Martindale, 1990). A design is most likely to be considered interesting if it is similar-yet-different to previously experienced designs. In other words, a design is likely to be interesting if it is novel but not entirely unfamiliar. Consequently, a motivation to seek out novelty can be a useful general-purpose heuristic in design. Martindale (1990) proposed

that the search for novelty is the only constant motivation in the development of artistic and architectural styles as cultural and social conditions change over time.

## 1.1 Emergence in Design

One source of novelty familiar to designers is emergence. A property of a design that is not represented explicitly at the time of creation is said to be an emergent property if it can be made explicit (Gero, 1994b; Mitchell, 1993). *Design emergence* is the process of recognition and explicit representation of emergent properties (Gero, 1994a). A familiar example of design emergence is shape emergence.

Emergent shapes in a drawing or sketch are unintended consequences of the drawing actions that produced them (Schön and Wiggins, 1992). Protocol studies of designers while sketching have shown that unexpected discoveries of emergent shapes can have a significant impact on the course of further design activity (Schön and Wiggins, 1992; Suwa, Gero et al. 1999). *Shape emergence* is the process of recognition and explicit representation of emergent shapes.

### 1.1.1 Computationally Modelling Shape Emergence

Typically, computational models of shape emergence have created an unstructured intermediate representation of a sketch and then identified emergent shapes by combining elements of the intermediate representation in new ways. Computational systems using infinite maximal lines (Gero and Yan, 1993) have proved successful in identifying emergent shapes (Damski and Gero, 1996) and emergent shape semantics (Gero and Jun, 1995). Alternative computational models of shape emergence have used bitmap images as intermediate representations and image processing techniques to find emergent shapes. Liu (1993) used neural networks to identify previously learned emergent sub-shapes, Tomlinson and Gero (1997) used a neural model of early visual processing, and Edmonds and Soufi (1992) used Gestalt operators to construct emergent groupings of shapes.

### 1.1.2 Computationally Modelling Design Emergence

Shape emergence is not the only form of design emergence that can be computationally modelled. Also, to exploit emergence in future design tasks, design agents must learn about the initially unintended consequences of their actions. Most of the computational models of shape emergence have lacked

the ability to learn. As a consequence all of the emergent shapes discovered had to be considered potentially "interesting" and presented to a user for further evaluation. In contrast, the computational model of design emergence presented here builds on previous work that developed a model of shape emergence capable of learning to expect emergent shapes (Gero and Saunders 2000).

The computational model of interest described below has been developed in recognition of the fact that design emergence is more than shape emergence: it models interest in the emergence of unexpected group behaviour in crowds of simulated pedestrians. The model of interest is used to develop a computational model of curiosity that uses the evaluation of interestingness to motivate the actions of a design agent. The task of the curious design agent, in this example, is to explore a space of possible doorway designs that allow crowds of simulated pedestrians to pass in opposite directions.

While the problem of designing a doorway is conceptually simple, the complex interactions between the pedestrians mean that emergent group behaviours play a critical role in determining the performance of different designs. Therefore the initial statement of the design problem is necessarily ill defined: it *cannot* include a description of every relevant detail of emergent group behaviour in advance. This provides a similar problem to those faced by human designers: our design agent's task includes both problem finding and problem solving.

Section 2 introduces our approach to developing curious design agents. Section 3 describes some experiments with an implementation of a curious design agent applied to the design of a doorway. We conclude with a discussion of the potential benefits of using curious design agents to assist human designers.

## 2. DEVELOPING CURIOUS DESIGN AGENTS

In this section we describe our approach to developing curious design agents. We begin by examining the role that curiosity and interest can play in computational models of designing. We then describe the components of a curious design agent.

### 2.1 Curiosity

In humans and animals the drive that we call curiosity rewards self-directed learning through inquisitive exploration in advance of a need to apply the knowledge gained. Berlyne (1971) describes curiosity as follows:

> Uncertainty can generate a kind of motivational condition that we call "curiosity". […] It will impel action to obtain further information from, or relating to, the object of curiosity so that information capable of relieving the uncertainty can be absorbed.

Curiosity motivates a designer to explore interesting designs to relieve the uncertainty that accompanies an incomplete understanding of the design space. A designer can be motivated by curiosity to investigate a new approach to solving a problem simply because it is interesting rather than because it is successful. Alternatively, curiosity can motivate a designer to explore new problems because the designer can recognise interesting problems where familiar designs do not perform as expected, whether for the better or for the worse.

Computationally speaking, *curiosity is a process that internally generates reinforcement signals sent to an agent's controller that rewards the discovery of interesting concepts.* The main difference between curious agents and other types of reinforcement learning agents is that some of the reinforcement signal is generated internally to reward the discovery of novelty (Schmidhuber, 1991). Curious design agents must be able to recognise both problems and solutions as interesting, fortunately, the same mechanisms can be used for both types of recognition.

## 2.2 What's Interesting?

In general, determining interestingness depends upon the knowledge of the agent and their computational abilities; things are boring if either too much or too little is known about them (Schmidhuber, 1997). Hence situations that are similar-yet-different to previously experienced situations are the most interesting and this is what we mean when we say that something is novel. *A novel situation is one that is similar enough to previous experiences to be recognised as a member of a class but different enough from the other members of that class to require significant learning.*

It is a relatively straightforward to develop a computational model of interest based on this definition of novelty. A very simple model of interest used in the following experiments maintains an average of the novelty detected over a fixed window of the ten most-recent situations. A boredom threshold is used to determine when the interest in the current area of a design space is low enough to warrant a change in the design process, e.g. a switch from problem solving to problem finding.

Empirical research suggests a strong connection between novelty and aesthetic preference in various creative fields including literature, art, architecture and music (Martindale, 1990; Gaver and Mandler, 1987). These

reports lend weight to the argument that novelty is an important determinant of interest in many creative fields including architecture.

## 2.3 A Curious Design Agent

The implementation of a curious design agent described here uses a combination of neural networks and reinforcement learning. Neural networks are used to construct a world model, i.e. a mapping from designs to evaluations. The world model constructed by the neural networks is rarely perfect and predictions of evaluations from design descriptions often contain errors. Some errors stem from a lack of adequate training and are of little interest but others are potentially more important.

Two sources of potentially interesting errors found in designing are consequences of emergent properties of the design task and the nature of learning and recall processes. Emergent properties of a design task can be sensitive to small differences in design parameters that can make a big difference to performance. World models that do not take these small differences into account can contain significant errors. Machine learning algorithms trade off being plastic enough to learn about new experiences with being stable enough to recall memories of previous experiences. The balance struck between stability and plasticity can have a significant affect on the accuracy of predictions.

A process called novelty detection is used to determine a measure of novelty that is proportional to the amount of error in the predictions of the neural networks. The level of interest in the current area of design space is calculated from the novelty that is used to produce a reinforcement signal for the agent's controller. The goal of a curious design agent is to maximise the reinforcement signal by seeking out novel situations.

### 2.3.1 Novelty Detection

The purpose of novelty detection is to identify unexpected or abnormal situations from examples of normal behaviour. Novelty detection has been used in domains as varied as medical diagnosis (Tarrasenko, 1995), industrial plant monitoring (Worden, 1997), robot navigation (Marsland et. al., 2000) and text retrieval (Yang, 1998).

Our implementation of novelty detection uses two Habituated Self-Organizing Maps (HSOMs) to estimate the novelty of a situation. An HSOM consists of a standard self-organising map (SOM) with an additional neuron that outputs the novelty of the current input (Marsland et al., 2000).

A self-organising map consists of a lattice of neurons that are used to represent different categories of inputs (Kohonen, 1993). Each neuron has an

associated vector of weights of the same dimension as the inputs. When a new input is presented to the SOM each neuron compares the similarity of its weight vector to the inputs. The neuron with the best matching weights is declared the winner. Learning is accomplished by updating the winner to reduce the difference between its weights and the inputs. In addition a neighbourhood of neurons around the winner are updated to reduce the difference between their weights and the inputs. This process results in a topographic map of the input space, with similar categories being represented by nearby neurons.

In an HSOM every neuron in the SOM is connected to the output neuron by habituating synapses that become less effective at transferring activation between neurons with use. The more frequently a map neuron fires the lower the efficacy of the synapse and hence the lower the output of the novelty detector.

The first HSOM estimates the novelty of a design by categorizing a representation of the design solution. The second HSOM estimates the novelty of the performance of the design by categorizing a profile of the design situation that includes representations of the design solution, the design problem and an evaluation of the design's performance.

The inverse of the novelty detected by the first HSOM is used to estimate the familiarity of a design. The output of the second HSOM is used to estimate the novelty of the design performance. The novelty of a design situation is calculated as a product of the familiarity assigned by the first network and the novelty assigned by the second. Consequently, significant novelty is only detected when a familiar design has an unfamiliar performance.

## 3. DESIGNING VIRTUAL ENVIRONMENTS FOR SIMULATED PEDESTRIANS

A simple crowd management problem is used to illustrate the behaviour of our curious design agent. The problem is to design a doorway to facilitate the efficient and comfortable movement of crowds of pedestrians travelling in opposite directions. A pedestrian simulator was developed to evaluate doorway designs. Pedestrian movement is simulated using a microscopic model of crowd behaviour developed to account for empirically observed self-organising phenomena.

## 3.1 Simulating Pedestrians

Computer models of pedestrian movement have been used to provide valuable tools for designers when planning or modifying pedestrian areas in large buildings like railway stations or shopping malls (Major et al., 1998). The "social force model" is a microscopic model of pedestrian behaviour that simulates the behaviour of individual pedestrians to model self-organising phenomena in crowds (Helbing, 1991).

### 3.1.1 The Social Force Model

Helbing and Molnár (1995) developed the social force model of pedestrian behaviour to simulate the pedestrian crowd movements to gain a better understanding of empirical results. The "social forces" in the model do not represent forces exerted upon a pedestrian; rather they are an approximation of the internal motivations of the individuals to move in certain directions. Despite its simplicity, computer simulations have shown that the social force model is capable of realistically describing several interesting aspects of collective pedestrian behaviours observed in empirical studies (Helbing and Molnár, 1997). The social forces modelled in these experiments are listed in *Table 1*. Detailed mathematical descriptions of these forces can be found in Helbing and Molnár (1995).

*Table 1.* The social forces modelled in the simulations of pedestrian crowds.

| | Description of social force |
|---|---|
| 1. | Pedestrians are motivated to move as efficiently as possible to a destination. |
| 2. | Pedestrians wish to maintain a comfortable distance from other pedestrians. |
| 3. | Pedestrians wish to maintain a comfortable distance from obstacles like walls. |
| 4. | Pedestrians may be attracted to other pedestrians (e.g. family) or objects (e.g. posters). |

### 3.1.2 Evaluating Virtual Environments

Designs are evaluated using measures of the efficiency and discomfort for each simulated pedestrian (Helbing and Molnár, 1997). Efficiency is measured for a pedestrian as the average difference between actual walking speed during a simulation and desired walking speed. Discomfort is calculated as a function of the number of direction changes during a simulation that a pedestrian must perform in order to negotiate the built environment and other pedestrians.

Like an architect, the primary concern of our design agent is the "subjective experience" of the simulated pedestrians visiting its environment. However, it should be stressed that our curious design agent does not attempt to optimise its designs in the computational sense. Instead

the design agent is motivated to explore the space of possible designs. It is equally motivated to investigate good and bad designs, e.g. inefficient designs can be interesting if their inefficiency is unexpected.

## 3.2 Experimental Results

This section describes two experiments using the models of interest and curiosity described above. The first experiment investigated the detection of novelty as emergent group behaviours affect the performance of three doorway designs. The second experiment investigated the behaviour of a curious design agent autonomously exploring the problem and solution spaces of doorway design.

### 3.2.1 Experiment 1: Assessing the Novelty of a Two Door Design

To illustrate the performance of the novelty detector, three designs for a doorway were created. The three doorway designs were for a narrow door, a wide door, and a combination of two narrow doors, as shown in *Figure 1*.

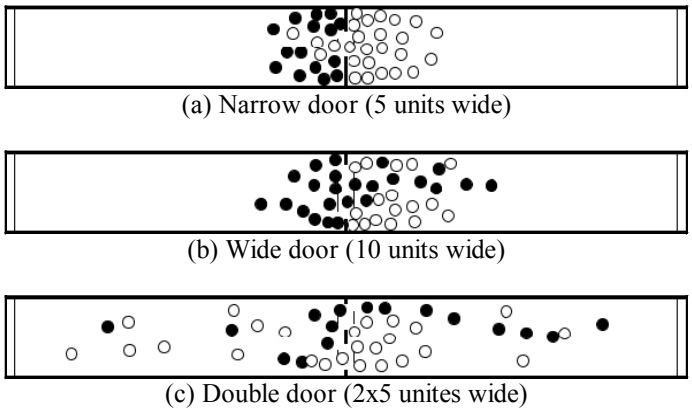


*Figure 1.* Screenshots of the simulations of pedestrian flow through (a) a narrow, (b) a wide, and (c) a double doorway design with a crowd of 40 pedestrians. The black circles indicate pedestrians travelling from left-to-right across the doorway and the white circles indicate pedestrians moving from right-to-left.

The doorway designs were tested using different numbers of pedestrians simultaneously trying to get through the doorway, crowds ranged in size from 1 to 51 pedestrians in increments of 10. The efficiency and discomfort measures from the simulations were combined into a single evaluation

measure for each simulation. The best evaluations of three trials conducted at each crowd size are shown in *Figure 2*.

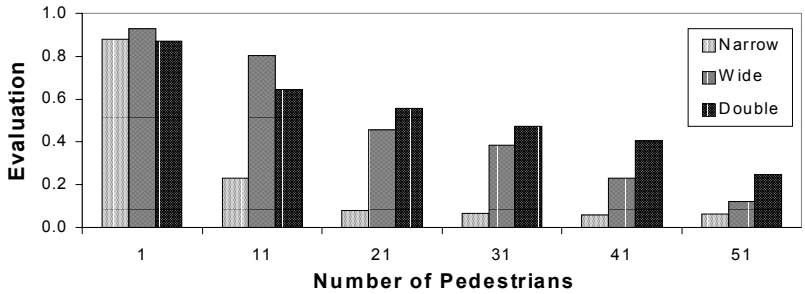


*Figure 2.* The best combined efficiency and discomfort evaluations for narrow, wide and double doorway designs for different crowd sizes (1–51 pedestrians).

All doorways performed equally well with only one pedestrian passing through it at a time. As the number of pedestrians increases the crowds display an oscillatory behaviour around doorways where one group of pedestrians gains control of the whole door at a time. The control of the doorway switches back-and-forth in direction as the numbers of pedestrians on either side of the doorway change.

The performance of the narrow doorway design quickly deteriorates to give consistently bad evaluations as the number of pedestrians increase. The wide doorway design maintains a very high performance for 11 pedestrians but its performance reduces dramatically, by almost 30%, as the number of pedestrians increases to 21. The performance of the wide door degrades more slowly over as the crowd sizes continue to increase from 31–51 pedestrians.

The performance of the double doorway design degrades even more slowly than the wide doorway design. For small crowds with less than 11 pedestrians the wide doorway design performs better but as the numbers of pedestrians increase the double doors outperform the wide door.

The double doorway design's superior performance in crowded conditions is a consequence of an emergent organisation. The two doors become specialised in the transfer of pedestrians moving in a single direction for relatively long periods of time. This can be seen in the double doorway simulation shown in *Figure 1*, pedestrians travelling from left to right pass through the top door while pedestrians travelling right to left pass through the bottom door.

The evaluations of each doorway design were presented to the novelty detector in ascending order of pedestrian numbers. The evaluations of the

narrow doorway were presented first, the wide doorway evaluations second and the evaluations of the double doorway were presented last. The best novelty measures of three trials are presented in *Figure 3*.

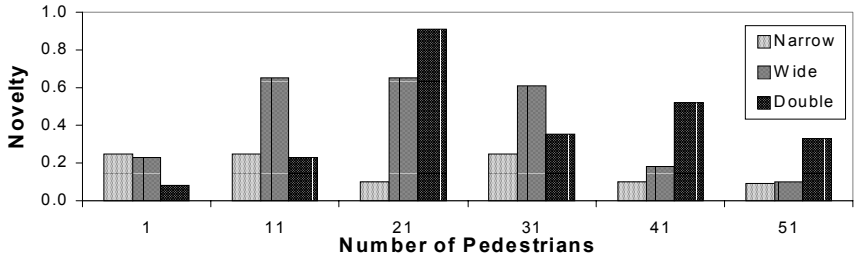


*Figure 3.* The greatest novelty detection for narrow, wide and double doorway designs for different crowd sizes (1–51 pedestrians).

Very little novelty was detected for the narrow doorway design at any crowd size. This is due to the lack of experiences against which the novelty detector could compare performances and the fact that the narrow doorway had consistently bad performance with more than one pedestrian.

The relatively high (~0.6) novelty measure for the wide doorway simulations with 11, 21 and 31 pedestrians indicate the improved performance of the wide doorway over the narrow doorway. The novelty of the wide doorway design drops at larger crowd sizes as the characteristics of the wide doorway are learned.

The novelty assessments of the double doorway design show very high novelty measures for simulations using 21 pedestrians, highlighting the resistance of the double doorway design to the fall in performance suffered by the wide door. The subsequent levels of novelty for simulations involving 31, 41 and 51 pedestrians reflect the relative differences in evaluations as the advantages of the double door design are maintained and the characteristics of the new design are learned.

The results of this experiment show that novelty detection can identify the most interesting designs without extensive reasoning by comparing the relative performance of different designs under similar conditions. The same novelty detector was used in the next experiment to implement models of interest and curiosity for an autonomous design agent.

### 3.2.2 Experiment 2: Curious Problem Finding and Problem Solving

In this experiment a curious design agent was given two conceptual spaces to explore: a problem space and a solution space. The solution space was defined by two variables: the number of doors making up the doorway and the combined width of doors. The problem space was defined by a single variable: the total number of pedestrians in the two crowds trying to get through the doorway. All other variables of the simulation remained constant.

*Figure 4* shows the novelty detected over the course of a design session. The design agent was initially given a narrow doorway as a solution to the problem of moving a single pedestrian. The novelty of exploring this design soon decreases as the agent learns to accurately predict the doorway's performance, the agent's interest level quickly falls below its boredom threshold and it begins to explore the problem and solution spaces for more interesting situations.

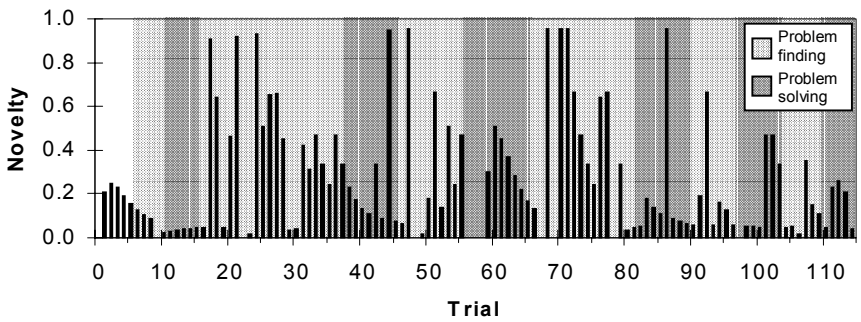


*Figure 4.* The results of using a curious design agent to explore the problem and solution spaces for doorway design. The chart shows the novelty detected for each simulation trial. The light shaded regions indicate that the design agent is problem finding and dark shaded regions indicate that the design agent is problem solving.

*Figure 4* shows the design agent switching between searching the problem and solution spaces as interest in a particular problem or solution wanes. The chart shows the "tailing-off" of novelty values as the characteristics of situations are learned. It also shows how the detection of novelty extends the period that an agent spends searching a particular space, especially the exploration of the problem space for trials 17–37 and 67–82.

The highest peaks in detected novelty (~0.9) in the first half of the experiment (up to trial 68) all correspond to simulations using double doorway designs as these have significantly different characteristics to single doorway designs initially explored.

The high peaks in the second half of the chart correspond to simulations using wide doorway designs. This change in fixation occurs when the interest in double doorway designs subsides. In the second half of the design session the design agent is discovering an array of interesting situations where a wide door does not perform in the same way as a double door. At lower numbers of pedestrians the wide doorway does better than the double doorway, while at higher numbers of pedestrians it performs worse. Either way, the design agent finds situations involving wide doorway designs novel and maintains a higher level of interest in exploring this area of the design space than would otherwise be expected.

The change in fixation of the design agent from double to wide doorways illustrates a difference in exploration between a more conventional optimisation approach and one based on curiosity. The curious design agent did not explore the situations using wide door designs because they performed better than the double door designs. Instead, it explored the space of wide door designs because they did not perform as expected from previous experiences of the similar-yet-different double door designs.

## 4. DISCUSSION

The experimental work described has investigated models of interest and curiosity using processes that detect the novelty of similar-yet-different design situations. Experiment 1 showed that novelty detection could be used to identify interesting situations where unexpected emergent properties play an important role in the evaluation of designs. Experiment 2 showed that using this model of interest a curious design agent can autonomously explore problem and solution spaces to identify interesting design situations.

Future experiments will include investigations of curious design agents applied to more complex design tasks. For example, a natural progression is to apply curious design agents to the design of large public spaces like train stations where frequent interactions between pedestrians and the resulting emergent group behaviours have a significant impact on the performance of the space.

The ultimate goal of this work is to develop design agents that can assist an architect explore the issues involved in complex design tasks. Architects increasingly face the problem of “information overload” as they try to explore complex design spaces for innovative solutions. Although generative design tools relieve some of the burden of designing, they can make the problem of information overload worse as designers attempt to understand the significance of the designs produced. Technologies similar to curious design agents may play an important role in future CAAD systems by

reducing the number of designs presented to an architect to a subset of those that are judged to be potentially interesting.

Providing design agents with motivations that reward the discovery of interesting designs more closely matches the motivations behind human exploration of design spaces. The application of curious design agents may allow future CAAD systems to provide more natural and rewarding collaborative partnerships between designer and machine.

## 5. ACKNOWLEDGEMENTS

This research is supported by an Overseas Postgraduate Research Scholarship and by a University of Sydney Postgraduate Award. We wish to thank Hsien-Hui (Michael) Tang for his valuable "non-computational" insights.

## 6. REFERENCES

Berlyne, D. E., 1971, *Aesthetics and Psychobiology*, Appleton-Century-Crofts, New York.

Damski, J. C. and Gero J. S., 1996, "A logic-based framework for shape representation", *Computer-Aided Design*, 28(3): 169–181.

Edmonds, E. and Soufi, B., 1992, "The computational modelling of emergent shapes in design", in: Gero and Sudweeks, (eds.), Computational Models of Creative Design, The Key Centre of Design Computing, The University of Sydney, Sydney, p.173–189.

Gaver, W. W. and Mandler, G., 1987, "Play it again, Sam: On liking music", *Cognition and Emotion*, 1(3): 259–282.

Gero, J. S., 1994a, "Computational models of creative design processes", in: Dartnall (ed.) *Artificial Intelligence and Creativity*, Kluwer, Dordrecht, p. 269–281.

Gero, J. S., 1994b, "Towards a model of exploration in computer-aided design", in: Gero and Tyugu (eds.), *Formal Design Methods for Computer-Aided Design*, North-Holland, Amsterdam, p. 271–291.

Gero, J. S. and Jun, H. J., 1995, "Visual semantics emergence to support creative designing: a computational view", in: Gero, Maher, and Sudweeks (eds.), *Preprints Computational Models of Creative Design*, Key Centre of Design Computing, The University of Sydney, Sydney, p. 87–116

Gero, J. S. and Saunders, R., 2000, "Constructed representations and their functions in computational models of designing", in: Tang, Tan and Wong (eds.), *CAADRIA 2000*, CASA, Singapore, p. 215-224.

Gero, J. S. and Yan, M., 1993, "Discovering emergent shapes using a data-driven symbolic model", in: Flemming and Van Wyk, (eds.), *CAAD Futures '93*, North Holland, Amsterdam, p. 3–17.

Helbing, D., 1991, "A mathematical model for the behavior of pedestrians", *Behavioral Science*, 36: 298–310.

Helbing, D. and Molnár, P., 1995, "Social force model for pedestrian dynamics", *Physical Review E*, 51: 4282–4286.

Helbing, D. and Molnár, P., 1997, "Self-organization phenomena in pedestrian crowds", in: Schweitzer (ed.) *Self-Organization of Complex Structures: From Individual to Collective Dynamics*, Gordon and Breach, London, p. 569–577.

Kohonen, T., 1993, *Self-Organization and Associative Memory*, 3rd ed., Springer, Berlin.

Liu, Y–T., 1993, "A connectionist approach to shape recognition and transformation", in: Flemming and Van Wyk, (eds), *CAAD Futures '93*, North Holland, Amsterdam, p. 19–36.

Major, M. D., Stonor, T. and Penn, A., 1998, "Passengers, pedestrians and shoppers: Space syntax in design", *Passenger Terminal World*, April.

Marsland, S., Nehmzow, U. and Shapiro, J., 2000, "A real–time novelty detector for a mobile robot", in *Proceedings of EUREL European Advanced Robotics Systems Conference*, Salford, England.

Martindale, C., 1990, *The Clockwork Muse*, Basic Books, New York.

Mitchell, W. J., 1993, "A computational view of design creativity", in: Gero and Maher (eds.) *Modelling Creativity and Knowledge-Based Creative Design*, Lawrence Erlbaum, p. 25–42.

Purcell, T. A. and Gero, J. S., 1992, "Effects of examples on the results of a design activity", *Knowledge-Based Systems*, 5(1): 82–91.

Schön, D. A. and Wiggins, G., 1992, "Kinds of seeing and their functions in designing", *Design Studies*, 13: 135–156.

Schmidhuber, J., 1991, "A possibility for implementing curiosity and boredom in model-building neural controllers", in: Meyer and Wilson (eds.) *Proceedings of the International Conference on Simulation of Adaptive Behaviour: From Animals to Animats*, MIT Press/Bradford Books, Cambridge, MA, p. 222–227.

Schmidhuber, J., 1997, "What's interesting?" *Technical Report TR-35-97*, IDSIA, Lugano, Switzerland.

Suwa, M., Gero, J. S. and Purcell, T., 1999, "Unexpected discoveries and S-invention of design requirements: A key to creative designs", in Gero and Maher (eds.), *Computational Models of Creative Design IV*, Key Centre of Design Computing and Cognition, The University of Sydney, Sydney, p. 297–320.

Tarassenko, L., Hayton, P., Cerneaz, N. and Brady, M., 1995, "Novelty detection for the identification of masses in mammograms", in: *Fourth International Conference on Artificial Neural Networks*, Cambridge, p. 442–447.

Tomlinson, P. and Gero, J. S., 1997, "Emergent shape generation via the boundary contour system", in: Junge (ed.), *CAAD Futures 1997*, Kluwer, Dordrecht, p. 865–874.

Whitfield, A. and Wiltshire, J., 1982, "Design training and aesthetic evaluation: An intergroup comparison", *Journal of Environmental Psychology*, 2: 109–117.

Worden. K., 1997, "Structural fault detection using a novelty measure", *Journal of Sound and Vibration*, 201(1): 85–101.

Yang, Y., Pierce, T. and Carbonell, J. G., 1998, "A study on retrospective and on-line event detection", *Proceedings of ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'98)*, p. 28–36.