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 in emergent properties of a design and develop a curious design agent that is motivated to discover interesting situations whilst designing.

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. 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.

Shape emergence is the recognition of emergent shapes in a drawing or sketch that 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).

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). *Figure 1* illustrates a good example of the emergence of multiple shape representations from a single building floor plan using infinite maximal lines.

Alternative computational models of shape emergence have used bitmap images as intermediate representations. Image processing techniques are used to find emergent shapes by recognising structure in the bitmap representation. Liu (1993) used neural networks to identify previously learned emergent sub-shapes, Edmonds and Soufi (1992) used Gestalt operators to construct emergent groupings of similar shapes, and Tomlinson and Gero (1997) used a model of early visual processing to emerge familiar optical illusions.

To exploit emergence in future design tasks, designers 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 "interesting" and presented to a user for further evaluation. In contrast, a computational model of shape emergence capable of learning to expect emergent shapes was presented in Gero and Saunders (2000).

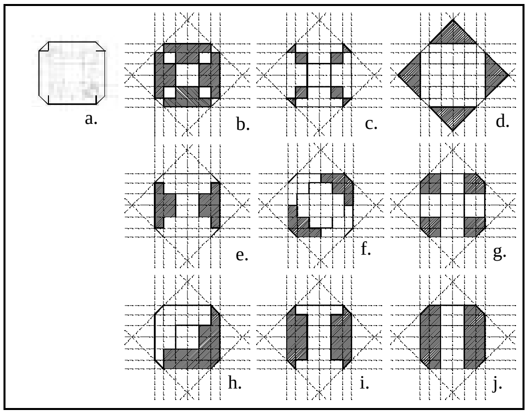

Figure 1. An example of the emergence of multiple representations for a floor plan design using infinite maximal lines (from Reffat and Gero, 1998)

1.1.2 Computationally Modelling Design Emergence

Shape emergence is not the only form of design emergence that can be computationally modelled. The computational model of interest described below has been developed in recognition of this fact: it models interest in the emergence of unexpected group behaviour found in crowds of simulated pedestrians. The curious design agent developed using this model of interest uses the evaluation of interestingness to guide its design activity.

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.

Emergent group behaviour can affect the evaluation of designs in two important ways. Firstly, designs that appear similar may perform very differently with the same numbers of pedestrians attempting to pass through them. Secondly, a single design may perform very differently with similar numbers of pedestrians as the group behaviour changes in character.

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 before examining issues of determining interest and how curiosity can be implemented.

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.

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 learn about new problems because the design recognises interesting situations 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 novel experiences (Schmidhuber, 1991). Curious design agents must be able to recognise problems and solutions as interesting. Fortunately, as we shall see 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.*

Empirical research suggests a positive connection between the novelty and aesthetic preference in various creative fields including art, literature, music and architecture (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 model of a curious design agent described here uses reinforcement learning to construct a world model. The agent uses a set of neural networks that learn to predict the consequences of taking design actions. The predictions of the neural networks are rarely perfect for several reasons. The most obvious sources of prediction failures are small differences in the environment that have a significant impact of the performance of a design but are not detected by the design agent. Another major source of prediction failure is the fallible nature of the learning process.

Machine learning algorithms often have to strike a balance between being plastic enough to learn about new experiences while being stable enough to retain memories of previous experiences. Plasticity can cause useful memories to be forgotten, while stability can cause new experiences to be misclassified. Both of these can have a significant affect on the accuracy of predictions.

Prediction errors are used by a process called novelty detection to determine a measure of the novelty of a situation that is then used to produce a positive reinforcement signal for the controller.

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-Organizing Map (Kohonen, 1993) with an additional neuron connected to every neuron of the map by habituating synapses (Marsland et al., 2000).

The first HSOM estimates the novelty of a doorway design by categorizing a representation of the design solution. The second HSOM estimates the novelty of the performance of the current doorway by first 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 the doorway design. The novelty detected by the second HSOM is used to estimate the familiarity of the doorway design performance. The novelty of a design situation is calculated as a product of the confidence assigned by the first network and the novelty detected by the second network. Consequently, significant novelty is only detected when a familiar design has an unfamiliar performance. A reinforcement signal proportional to the novelty of the design situation is produced to reward the controller for finding novel situations, indicating the degree of interest.

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-organizing 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). Microscopic models of pedestrian motion take the individual pedestrian as the elemental unit to simulate rather than the collective behaviour as in macroscopic models (Helbing, 1991).

3.1.1 The Social Force Model

The social force model of pedestrian behaviour was developed to model self-organising phenomena in the motion of pedestrian crowds (Helbing and Molnár, 1995). The "social forces" do not represent forces exerted on a pedestrian; rather they are an approximation of the internal motivations of the individuals to move in certain directions. Despite its simplicity, previous computer simulations have shown that the social force model is capable of describing several interesting aspects of collective pedestrian behaviours in a realistic manner (Helbing and Molnár, 1997). The social forces modelled in the simulations of pedestrian crowds 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 comfort for each simulated pedestrian. 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 doorway design space and the pedestrian problem space.

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 performances of the doorways were evaluated using the pedestrian simulator. The three doorway designs are a narrow door, a wide door, and a combination of two narrow doors, as shown in *Figure 2*.

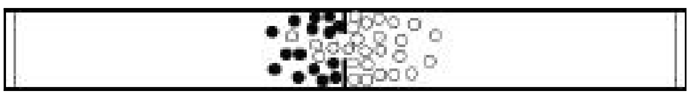

(a) Narrow door (5 units wide)

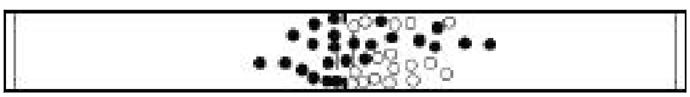

(b) Wide door (10 units wide)

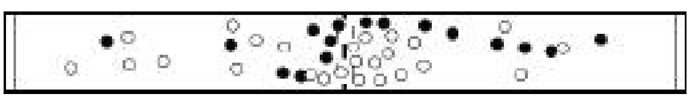

(c) Double door (2x5 unites wide)

Figure 2. 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 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 presented to a novelty detector as the evaluations of the doorways.

The evaluations of the narrow doorway design were presented in order of ascending pedestrian numbers to the novelty detector. The evaluations of the wide doorway were presented second and the evaluations of the double doorway were presented last. The resulting novelty measures are presented in *Figure 3*.

Figure 3. The results of using novelty detection on three different door designs (narrow, wide and double) for different crowd sizes (1–51 pedestrians).

Very little novelty was detected for the narrow doorway design. This was partly due to the lack of experiences against which the novelty detector could compare performances, and also because simulations involving more than one pedestrian passing through the door at a time resulted in similarly bad performance, this can be seen in the narrow doorway evaluations in *Figure 4*.

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. For 41 and 51 pedestrians the performance of the wide doorway drops as it becomes too crowded to work effectively. The evaluations for these numbers of pedestrians are more similar to the evaluations of the narrow doorway design. The fall in the novelty of the wide doorway design is partly in response to this and partly because some of the characteristics of the wide doorway independent of the narrow doorway design have been learned.

The assessments of the double doorway design shows very high novelty measures for simulations using 21 pedestrians, highlighting the resistance of the double doorway to the fall in performance suffered by the wide doorway design. *Figure 4* shows how the performance of the double doorway design under increasingly crowded conditions degrades more slowly than the wide doorway design.

The novelty of the performance of the double door design for simulations involving 21 pedestrians is particularly pronounced because it is the first occasion when the double doorway outperforms the wide doorway. 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 become more pronounced.

The reason for the double doorway design's superior performance in crowded conditions is the formation of an emergent organisation whereby 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 2*, pedestrians travelling from left to right pass through the top door while pedestrians travelling right to left pass through the bottom door. In contrast, the narrow and wide doorways display oscillatory behaviour where one group of pedestrians gains control of the whole door at a time, switching back-and-forth in direction as the numbers of pedestrians on either side of the doorway change.

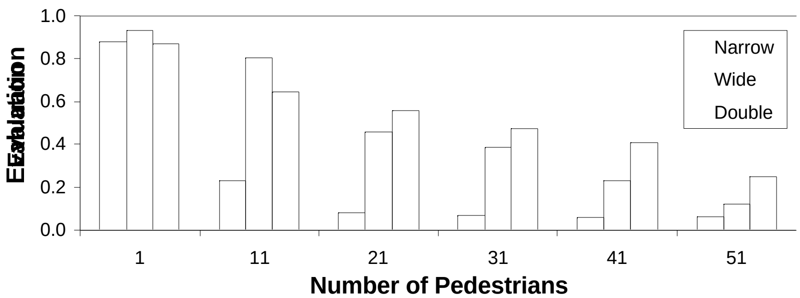

Figure 4. The best combined efficiency and discomfort evaluations for narrow, wide and double door designs for different crowd sizes.

The results of this experiment show that novelty detection can identify the most interesting aspects of the design problem by comparing the relative performance of similar-yet-different designs under similar conditions. The same novelty detector was used in the next experiment to implement curiosity in 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 total width of the doorway and the number of doors making up the doorway. The problem space was defined by a single variable: the number of pedestrians in the crowds trying to get through the doorway. All other variables of the simulation remain constant throughout.

Figure 5 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 he agent learns to predict the doorways performance. The design agent's interest, calculated as a function of the average novelty over the ten most recent simulations quickly falls below the agent's boredom threshold and it begins to explore the problem and solution spaces for more interesting situations.

Figure 5 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.

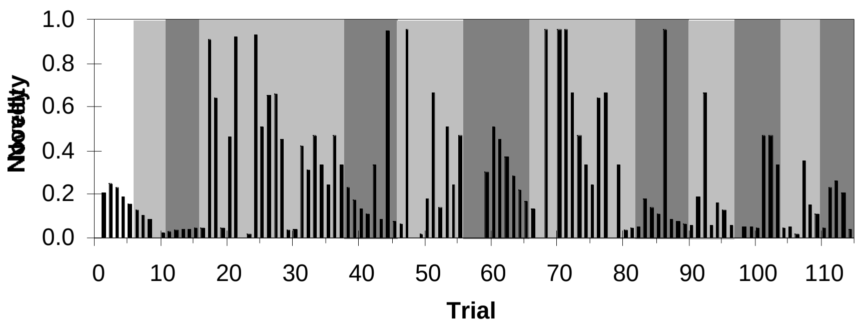

Figure 5. The results of using a curious design agent to explore the problem and solution spaces for doorway design. The chart shows the novelty detected for each simulation run and the switching between problem finding and problem solving.

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.

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 as a consequence of the design agent learning its performance characteristics.

The interest in the wide doorway designs is promoted by similarity of the representations of the wide and narrow doorways. The design agent is discovering an array of interesting situations where the wide door does not perform in the same way as the double doorway design.

At lower numbers of pedestrians the wide door 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 the 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 emergence plays 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 applications of the system presented here will include more complex tasks, including the design of large public spaces like train stations. The design of large public spaces must take into account the interactions between pedestrians that occur frequently. The appropriate placement of facilities in a public space with this in mind can have a significant impact.

The design of public spaces and the analysis of crowd movements within them would be an appropriate application of curious agents in larger CAAD systems because simulations like these can take a long time to run. An autonomous design agent can run them independently and then report only the interesting designs. This would allow a designer to concentrate on a smaller subset of potentially interesting design cases to quickly learn about the design problem.

Designers using computational tools that can generate and analyse designs face the problem of "information overload" as they try to understand the significance of the many designs produced. Similar problems are being faced in other industries as information becomes ever more plentiful. The growing popularity of data-mining techniques and collaborative filtering technologies are testament to this.

Technologies similar to curious design agents may play an important role in future CAAD systems as effective methods of reducing the number of designs produced by generative design systems that need a designer's attention. By providing design agents with motivations that reward the discovery of interesting rather than just successful design, future CAAD systems may be able to provide a more natural and rewarding basis for collaborative partnerships between designer and machine.

5. 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*, Key Centre of Design Computing, The University of Sydney, Sydney, p.?-?.

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, p.?-?.

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, p.?-?.

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.

This is a copy of the paper: Saunders, R. and Gero, J. S. (2001) Designing for interest and novelty: Motivating design agents, *CAADFutures* (to appear)