Artificial Intelligence for Engineering Design, Analysis and Manufacturing (2004), *18*, 153–161. Printed in the USA.
Copyright © 2004 Cambridge University Press 0890-0604/04 $16.00
DOI: 10.1017/S0890060404040119

Curious agents and situated design evaluations

ROB SAUNDERS AND JOHN S. GERO

Key Centre of Design Computing and Cognition, School of Architecture, Design Science and Planning, University of Sydney, Sydney, NSW 2006, Australia

(RECEIVED May 27, 2003; ACCEPTED October 6, 2004)

Abstract

This paper presents a possible future direction for agent-based simulation using complex agents that can learn from experience and report their individual evaluations. Adding learning to the agent model permits the simulation of potentially important agent behavior such as curiosity. The agents can then report evaluations of a design that are situated in their individual experience. The paper describes the architecture of curious agents used in the situated evaluation of designs. It then describes an example of the application of such curious agents in the evaluation of the curating of an exhibition in an art gallery.

Keywords: Curious Agents; Emergent Design; Simulations; Situated Design

1. INTRODUCTION

Designers have a long tradition of using simulations to evaluate and analyze their work. Visual representations, such as renderings, scale models, and artistic impressions have been used for centuries to simulate the appearance of designs prior to manufacture. The advent of powerful computer graphics software has provided new technologies for producing visual representations of designs; polygonal modeling, raytracing, radiosity, and virtual reality have allowed designers to generate increasingly realistic and immersive simulations of their work.

The use of simulation in design is not limited to visual representations, however; mathematical models also play an important role in the daily lives of designers wishing to simulate the behavior of their designs prior to construction. Sets of mathematical equations are used by designers and engineers of physical structures to simulate the forces that affect their designs when in use. Physical models and simulators, such as wind tunnels, are also commonly used to analyze the behavior of designs when the mathematics required are formidable. A good example is the use of wind tunnels to evaluate the aerodynamic properties of vehicles under various conditions.

Reprint requests to: John S. Gero, Key Centre of Design Computing and Cognition (G04), University of Sydney, Sydney, NSW 2006, Australia. E-mail: john@arch.usyd.edu.au

Recently, the power of modern computers has allowed designers to use complex computational models instead of expensive physical simulators to analyze the behavior of their designs. Advanced mathematical techniques such as computational fluid dynamics and finite element analysis allow designers and engineers to analyze the performance of designs without the investment in time and money required to build scale models for physical simulation.

The development of artificial intelligence has provided new opportunities for designers to simulate the performance of their designs. In particular, agent-based simulations are proving useful for architects and town planners wishing to analyze their designs. Agent-based simulations allow designers to evaluate the behavior of individuals and groups inhabiting a space. For example, simulations of crowds of people have been used to analyze the performance of designs under emergency conditions to assess the ability of a space to support the rapid evacuation of a building on fire.

Typically, agent-based simulations have been modeled using reactive agent models that follow a small number of simple rules. The models are then analyzed by observing the simulation to identify significant behavior, for example, overcrowding during an emergency evacuation. Reactive agents are ideal for modeling the behavior of individuals in emergency situations where there is little time for decision making and people will tend to follow the crowd in a herd-like manner. Reactive agents are also useful in simulations

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

of overcrowding at large public gatherings like football stadiums and train stations when an individual's ability to take independent action is diminished by the lack of available space.

Unfortunately, purely reactive agents do not permit the simulation of individual behavior in many, more common, situations that would be desirable when analyzing the design of buildings for public use such as train stations, museums, and galleries where the support of problem solving, learning, and exploration are key functions of the building. The remainder of this paper presents a possible future direction for agent-based simulation using more complex agents that can learn from experience and report their individual evaluations. Adding learning to the agent model permits the simulation of potentially important agent behavior, for example, curiosity. The agents can then report evaluations of a design that are situated in their individual experience, such as their level of interest as they explore.

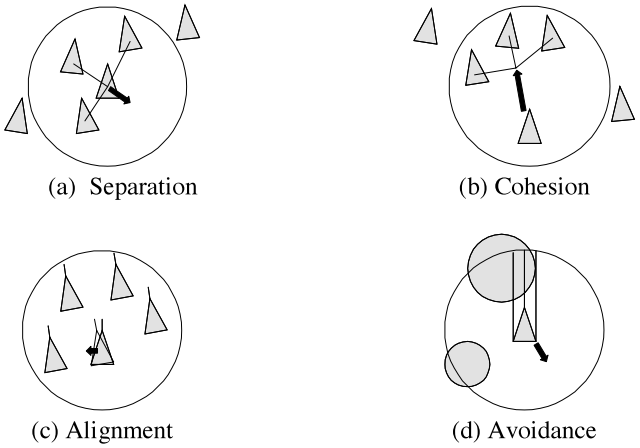


Fig. 1. The steering behaviors used in Reynolds' model of flocking.

2. SIMULATING CROWDS

Reynolds (1987) demonstrated that realistic simulations of groups of animals could be produced using simple reactive agents executing a small number of carefully chosen rules.

2.1. Flocks, herds, and schools

Reynolds (1987) proposed a set of four simple rules, that, when executed together, within simulated agents, also known as boids, resulted in realistic group behavior similar to a flock of birds, a herd of cattle, or a school of fish. The rules executed by each agent are the following:

1. separation: steer to avoid local flock-mates.
2. alignment: steer toward the average heading of local flock-mates.
3. cohesion: steer to move toward the average position of local flock-mates.
4. avoidance: steer to avoid running into local obstacles or nonflock-mates.

Separation prevents agents from overcrowding under normal conditions. Alignment aligns each agent with its immediate neighbors so that they move forward as a group. Cohesion maintains a "natural-looking" closeness to a neighborhood of agents. Finally, avoidance allows an agent to go around obstacles and avoid potential predators.

The four rules described above are used to implement steering behaviors using a very simple model of locomotion that applies a force to the body of the agent that is calculated to achieve the desired consequence of applying a rule. Examples of the kinds of forces applied during a flocking simulation are illustrated in Figure 1. During a simulation the forces produced for each rule are combined into a single force applied to the body of the agent; often this is achieved simply by summing the forces.

The flocking algorithm has been extended to simulate the motion of crowds of people in simulations and games (Woodcock, 1999). Flocking is used in these instances as a way for crowds to follow paths determined using by path-finding routines. Because the original flocking model does not contain any notion of moving toward a goal, the applications of flocking in game environments often require the addition of a rule to move agents toward waypoints along a path to a goal location. In this way the extended flocking algorithm maintains a group's formation and local obstacle avoidance, leading to the natural-looking movement of agents between goals.

2.2. The social force model

The social force model is a microscopic model of pedestrian behavior that has been used to model self-organizing phenomena observed in crowds of people (Helbing & Molnár, 1995). Helbing and Molnár developed the social force model to simulate crowd behavior to gain a better understanding of empirical results. The social forces in the model do not represent physical forces exerted upon a pedestrian; rather, they are an approximation of the internal motivations of the individuals to move in certain directions. The social forces modeled by each agent follow.

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.
4. Pedestrians may be attracted to other pedestrians or objects (e.g., posters).

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

Obviously, the forces implemented in the social force model are very similar to the rules devised by Reynolds for flocking; the social forces 2, 3, and 4 are very similar to separation, avoidance, and cohesion. The social force model does not include a force to maintain alignment among pedestrians as with flocking, but it does add a force to model the movement between locations as used in gaming environments. Detailed mathematical descriptions of these forces can be found in Helbing and Molnár (1995).

Despite its simplicity, computer simulations have shown that the social force model is capable of realistically describing several interesting aspects of observed crowd behaviors. In one instance, predictions based on simulations of crowd behavior at junctions prompted new empirical research into human crowd behavior that confirmed the emergence of transient round-about motions (Helbing & Molnár, 1997).

2.2.1. Agent-centric evaluations

In their experiments with emergent crowd behavior around doors, Helbing and Molnár used some simple agent-centric measures to evaluate the efficiency and discomfort for each pedestrian (Helbing & Molnár, 1997). Efficiency is measured for a pedestrian as the average difference between the speed it is walking toward its goal and its desired walking speed. Discomfort is calculated as a function of the number of direction changes during a simulation that a pedestrian must perform to negotiate the built environment and other pedestrians.

Using agent-centric evaluations allowed Helbing and Molnár to evaluate the performance of simulated spaces using nonhomogenous crowds of pedestrian agents, for example, the agents used in crowd simulations varied in their desired walking speed to simulate younger and older pedestrians within the same crowd. This conveys an improvement in the nature of the evaluation: a simpler measure of efficiency, for example, number of pedestrians to pass a given point per minute, would not adapt to crowds consisting of pedestrians with differing preferences.

3. CURIOUS AGENTS FOR DESIGN EVALUATION

The agent model presented in this paper adds a model of curiosity based on learning to the social force model to support the evaluation of environments that are designed to stimulate exploration. This curious social force model extends Helbing and Molnár's model with the addition of a single rule, "Pedestrians are motivated to move toward potentially interesting areas."

3.1. Curious agent architecture

The architecture of the curious agents used to evaluate designs is illustrated in Figure 2. The curious agent comprises six primary functions: sense, learn, detect novelty, calculate interest, plan, and act. In addition, each agent requires a long-term memory to store category prototypes.

Sensing samples the world to produce a stimulus pattern that characterizes its environment according to the abilities of the agent. Learning updates prototypes stored in long-term memory to better reflect the agent's experiences as new types of stimulus pattern are produced. The differences between a new stimulus pattern and the closest matching category prototype are used to calculate a measure of the

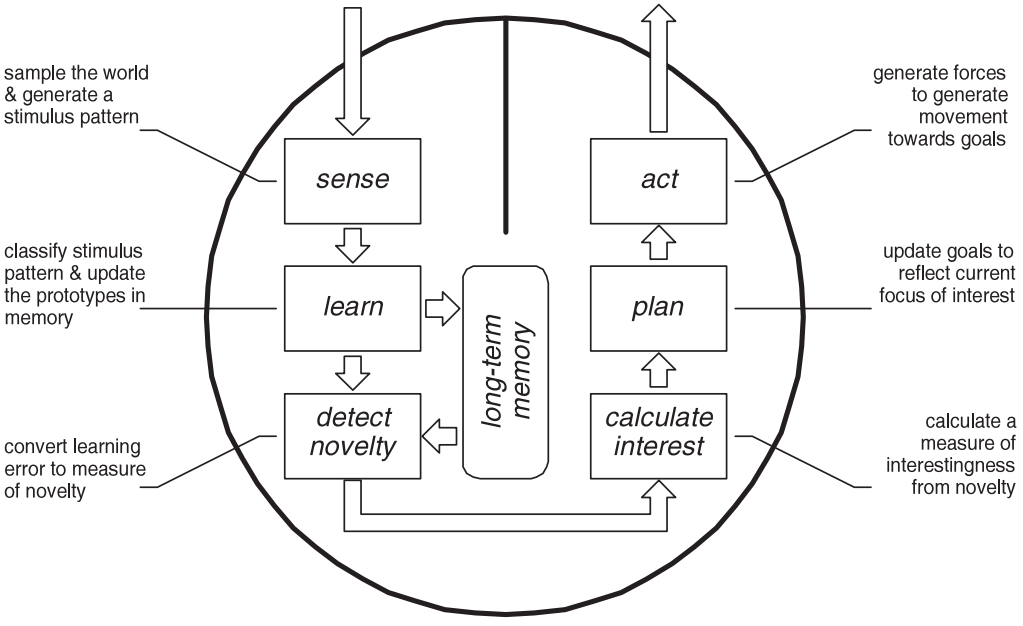


Fig. 2. The architecture of a curious agent.

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

novelty for the experience. A measure of the interestingness of the current situation is then calculated based on a psychobiological model of preference to arousing stimuli. The goals of the agent are then updated to reflect the agent's current interest in its surrounding, and forces are generated to propel the agent in an appropriate direction.

3.2. Interestingness, novelty, and curiosity

The determination of interestingness requires the agent to learn from its experiences. Interest in a situation depends upon an agent's goals, its previous experiences, and its ability to predict future situations. Interestingness is often based on the novelty of a situation when an agent's goals include learning. Curiosity is the term given to the exploratory behavior displayed by agents intended to reduce the uncertainty produced by novelty.

3.2.1. Detecting novelty

Detecting novelty in a situation requires a comparison between expectations and observations and this requires that the agent learn and predict aspects of a situation from previous experiences. The degree to which a stimulus pattern is novel will be inversely proportional to

1. how often similar patterns have been experienced,
2. how similar these patterns have been, and
3. how recently these patterns have been experienced.

To understand how the degree of novelty might be measured, it is useful to think in terms of conceptual spaces, as defined by Gärdenfors (2000). A conceptual space represents the knowledge held by an agent as a space where similar concepts are located closer together than dissimilar ones. Gärdenfors (2000) argued that such spaces can be defined in terms of prototypes, and that the recognition of a stimulus is in terms of its distance from nearby prototypes. Empirical evidence appears to show that humans and other sophisticated mammals categorize stimuli in a way that is consistent with this mechanism.

Figure 3 illustrates the different notions of novelty described above using a conceptual space containing three categories, C_1, C_2, and C_3. Each category is represented by a prototype, marked by a cross. The group of points around each prototype indicates the experiences mapped to each category, and the circle surrounding each prototype represents the standard deviation of those experiences from the prototype. The lines dividing the space are the boundaries of the Voronoi cells around each category; any point falling within a particular cell is attributed the category in the same cell.

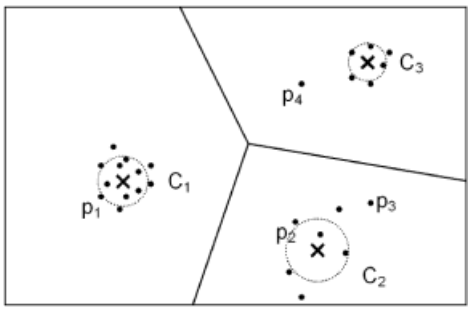


Fig. 3. A conceptual space occupied by three prototypes, C_1, C_2, and C_3.

The four points, p_1, p_2, p_3, and p_4, represent four experiences, mapped onto the conceptual space with different degrees of novelty. According to the first measure of novelty given above, the experience represented by p_2 has greater novelty than that represented by p_1 because fewer previous experiences have been mapped to the category C_2 than C_1. According to the second measure of novelty, the experience represented by p_3 is more novel than the experience mapped to p_2 because it is a greater distance from the prototype of C_2 than p_2. By the same rule, the experience mapped to p_4 is even more novel than p_3 because it is further from the other experiences that have also been mapped to C_3. In other words, it is further from the standard deviation marked by the circular region around the category prototype.

Unsupervised neural networks are ideal for detecting novelty because they often rely on classification error of new stimuli to guide training. Errors are often measured using distance metrics in the vector space defined by the output of the network and are similar to the measurement of distances in the conceptual space illustrated above. We have previously used a special type of neural network called a novelty detector in our agent-based research to model curiosity (Saunders & Gero, 2001*a*) as have other researchers (e.g., Schmidhuber, 1991).

3.2.2. Calculating interestingness

Novelty is not the only determinant of interestingness: interest in a situation is also related to how well an agent can learn the information gained from novel experiences. Consequently, the most interesting experiences are often those that are similar yet different than previously encountered experiences because these experiences provide the most opportunity for rapid learning (Schmidhuber, 1997).

Berlyne (1971) proposed that a nonlinear function called the Wundt Curve could be used to model the typical response that organisms have to many types of stimuli, including novelty. The Wundt Curve is illustrated in Figure 4. The Wundt Curve is calculated as the sum of two nonlinear functions that are used to model the reward and punishment generated internally by an agent as a consequence of experiencing a stimulus. Importantly, the Wundt Curve peaks at a maximum value for a moderate degree of stimulation, meaning that the most interesting forms of novelty are those that are similar yet different than previously encountered experiences.

For a given stimulus, the novelty detector described above will generate a value for the novelty of the stimulus pattern

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

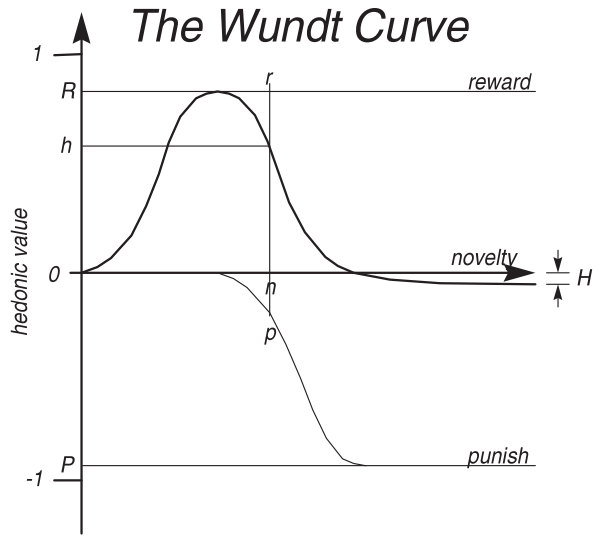


Fig. 4. The Wundt Curve.

(n) using the Wundt Curve this is transformed to an hedonic value, $h = r + p$, where r is the reward generated for the discovery of a novel stimulus and p is the punishment generated for the discovery of a highly novel stimulus. The hedonic value of a stimulus represents how interesting the novelty of the experience is. For moderate values of n the hedonic value will be positive, but as n gets larger the hedonic value can fall below zero, indicating that the novel stimulus is repellent to some degree. For very large values of n the hedonic value of the experience will approach, $H = R + P$, where R and P are the most positive reward and the most negative punishment, respectively. This is another important characteristic of the Wundt Curve as a model of interest based on novelty, as it is a familiar observation that too much novelty in a design can result in the repulsion of its audience.

Agents can be instantiated with different preferences for novelty by adjusting the response curves used to calculate reward and punishment for discovering novel situations. For example, agents can be created with a preference for less novelty by adjusting the reward and punishment response curves to move the part of the Wundt Curve with the greatest hedonic value closer to the origin. Agents with different preferences for novelty have been used in previous simulations to investigate the emergence of social notions of creativity in groups of curious design agents exploring a space of artworks (Saunders & Gero, 2001*b*).

3.3. The curious social force model

To model curiosity the interestingness of an experience must be translated into a behavior with the goal of learning more about the stimulus producing the interest and thereby reducing the agent's uncertainty. Using the same simple model of agent locomotion used in flocking and the social force model provides one way to model curious behavior by generating a force that will tend to move an agent toward stimuli with high hedonic value. Figure 5 illustrates one way that curious social force can be calculated.

In the example given in Figure 5 an agent has sensed three potentially interesting objects in the environment and has assigned hedonic values to each one. Three forces are then generated in the direction of the objects with magnitudes equal to the maximum force applicable scaled by the amount of interest the agent has in each one, that is, the assigned hedonic value between 0.0 and 1.0. The individual forces for each object are used to determine the curious social force for the agent by averaging the directions and magnitude of each force.

Other methods for calculating the curious social force are possible, for example, a winner-takes-all method using the greatest force calculated for a single object or by taking into consideration other factors such as the distance to the object of interest, but the above method is sufficient for demonstration purposes.

3.4. Situated design evaluations

Situatedness is an important concept for designers because evaluations of a design are generally located within a context so that decisions that are taken are a function of both the situation and how the situation is constructed or interpreted. Good designs should support situated decision-making processes by providing an appropriate context for constructing useful situations. Agent-centric evaluations of the interestingness as they explore a designed space are good examples of "situated design evaluation" because, as described above, the evaluations are situated in the history of experiences remembered by the agents.

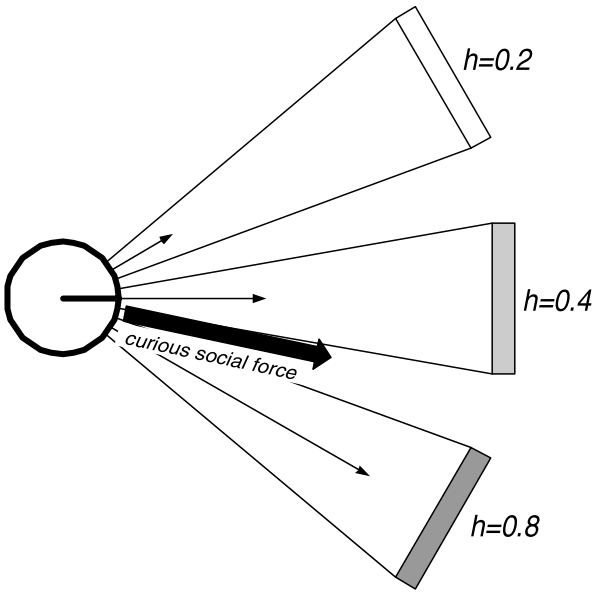


Fig. 5. An example of calculating a curious social force.

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

The agent-centric evaluations used by Helbing and Molnár (1997) are extended in this work by using situated design evaluations such as interestingness measures. The use of an agent-centric approach to evaluating the interestingness of a design allows the model to support crowds of agents with different preferences for novelty. Agents with a preference for low amounts of novelty can be modeled alongside agents with a preference for high levels of novelty, and the interest levels of both types of agent can be used to evaluate the design.

Alternative situated design evaluations to interestingness include the related concept of boredom, measures of learning, and the ability of agents to achieve their goals. For example, measures of an agent's learning during the course of an interaction with a design could be very useful in the simulation of museums and galleries, and could be implemented using relatively simple tests conducted on agents before and after their visit to the simulated building to determine how much the agent has learned during its visit. Similarly, other situated design evaluations can be determined through a combination of continuous monitoring and specific testing of knowledge.

4. AN EXAMPLE DESIGN PROBLEM: CURATING A GALLERY

Modeling curiosity in crowds of pedestrian agents would permit the simulation of design problems where maintaining an interest in a space is as important as efficient movement or comfort. To illustrate this idea, consider the problem of designing an art gallery exhibition. Problems for a curator of a gallery include keeping an exhibition interesting throughout an individual's visit. To keep the experience interesting, a gallery's design would have to take into account the preference of curious pedestrians to encounter similar yet different experiences in their travels.

4.1. Implementation

The gallery environment is implemented in a stylized form, with a number of rooms connected by doorways, allowing the contents of each room to be partially visible from its neighbors. The artworks in the gallery are modeled as areas of flat color with different hues that allow the agents to use computationally inexpensive vision and learning processes. The implementation of the curious agent architecture used in the gallery simulations is described in the remainder of this section.

4.1.1. Sensing

The curious agents used in the gallery simulation make use of a simple raycasting model of vision. To sense the environment an agent sends out a number of rays into the environment within its "field of vision." Where a ray collides with an object in the environment the agent senses the object. Figure 6 illustrates the raycasting model of vision.

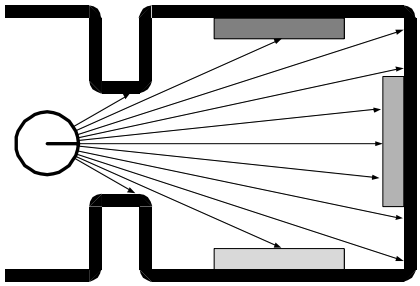


Fig. 6. A simple vision model implemented using raycasting.

This approach to modeling vision is simple to implement as an extension of the detection system needed to implement obstacle avoidance behaviors. It is particularly well suited to the task of simulating an agent's visual sense in the gallery example because the gallery's simple geometry allows it to be adequately sampled using a 2-dimensional raycasting engine and only a few rays per agent. In the implemented system, three rays are used to implement short-range detection of possible collisions with walls, and so forth, and three additional rays are used to implement long-range vision.

The information returned by the raycasting engine contains a number of color samples from the objects that have been detected in the environment. The agent's sensory system has been designed to transform the color samples into hues and to discard any samples that return black or white. In this way the agent is programmed to pay attention to the "paintings" in the gallery and to ignore the walls for the purposes of novelty detection. The obstacle avoidance behavior still makes use of the distance recorded by each ray to the walls and other obstacles to steer away from potential collisions.

4.1.2. Long-term memory, learning, and novelty detection

The advantage of filtering the colors returned by the raycasting vision system is that the resulting hue can be represented as a single angle. A standard representation of hue is the color wheel illustrated in Figure 7 that shows how hues can be translated into degrees around the wheel.

An appropriate learning system for a simple conceptual space such as the one presented by the color wheel is a 1-dimensional self-organizing map (SOM). SOMs are a type of unsupervised neural network commonly used for classification of data that have been applied to a wide range of applications (Kohonen, 1995). In the implemented system, the classification error reported by the SOM learning algorithm is used as a measure of the novelty inherent in the sense data produced by an experience.

Figure 8(a) illustrates how an SOM can learn a good representation of the conceptual space of a color wheel. The color wheel, represented in the figure by a shaded ring, has been mapped by an SOM with just six evenly spaced

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

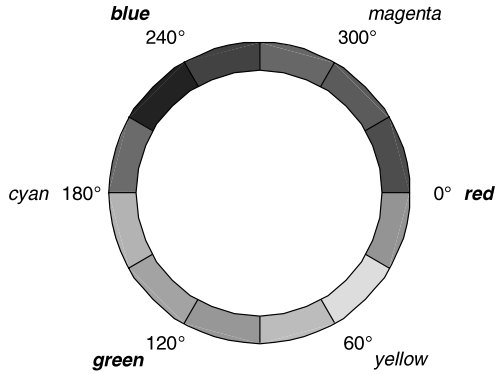


Fig. 7. A color wheel.

neurons arranged in a hexagonal pattern around the center of the color wheel. A 2-dimensional space is used, even though hue can be represented by the 1-dimensional space of angles around the color wheel because the space is continuous in two dimensions, making the recognition of the similarity of colors like magenta and orange natural even though they have quite different hue values.

The very small size of the SOM used to map the color wheel space may not be appropriate for many applications, but for the modeling of curiosity in our implementation it is ideal because the limited number of neurons mean that agents may have to "forget" certain colors to map new colors as they are encountered, thus distorting the SOM. Figure 8(b) illustrates how the learning process can distort an SOM's representation of the color wheel due to overexposure to colors in the red–orange part of the color wheel. Although this type of distortion is undesirable in many applications of SOMs, it is key to simulating the changing interests of agents as they explore a conceptual space.

4.2. Emergent design problems

Figure 9 illustrates the kinds of problems that can emerge during simulations. Inappropriate positioning of artworks around the gallery causes the production of curious social forces within agents that impede or otherwise impoverish their visit. The curious social forces are indicated in Figure 9 by large light gray arrows.

The two problems illustrated in Figure 9 are the formation of a crowd blocking the passage at the gallery entrance and the streaming of pedestrians past paintings at the end of the gallery tour. In the first case, the painting in the second room visible from the first is so different from those in the first room that the pedestrians prefer to remain in the first room causing a blockage. The curious social force generated in response to the painting in the second room causes the agents to behave as if there is a physical obstacle blocking the path to the second room. Eventually, agents will move from the first to the second room as the "pressure" produced by agents crowding into the small space behind them overcomes the curious social force pushing them back. This is similar to the "pressure" reported by Helbing and Molnár as an important factor in their simulations of group behavior around overcrowded doorways (Helbing & Molnár, 1997).

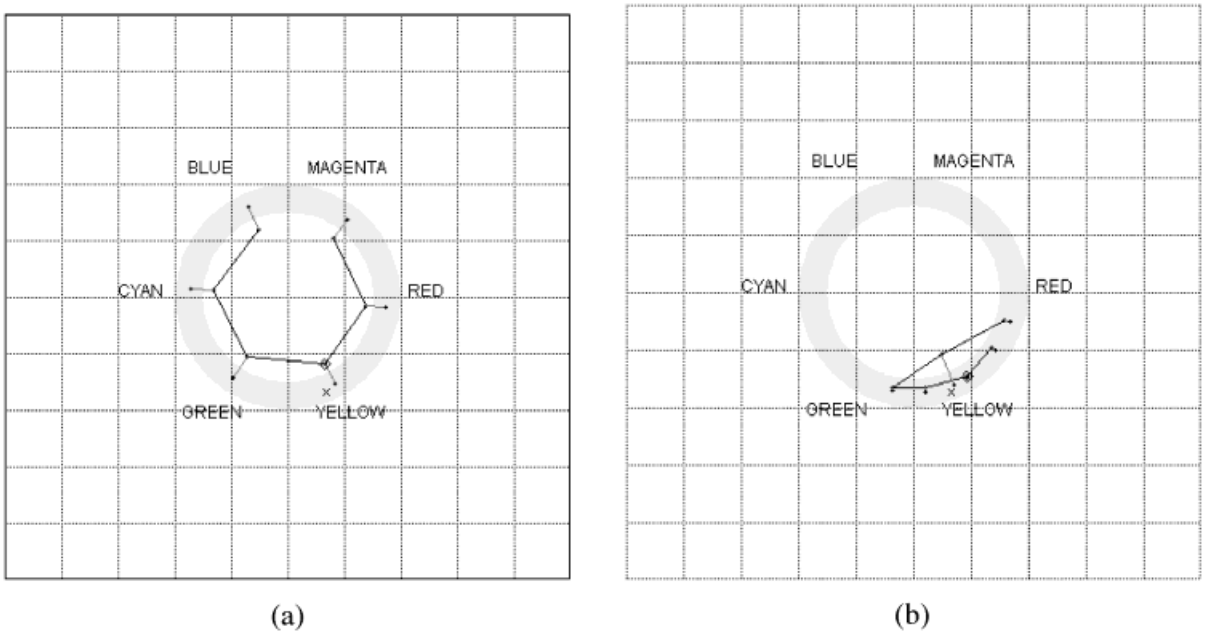


Fig. 8. Two examples of the categorization of inputs with SOMs trained on uniform and nonuniform sampling of colors. (a) Balanced and (b) unbalanced sampling of color space.

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

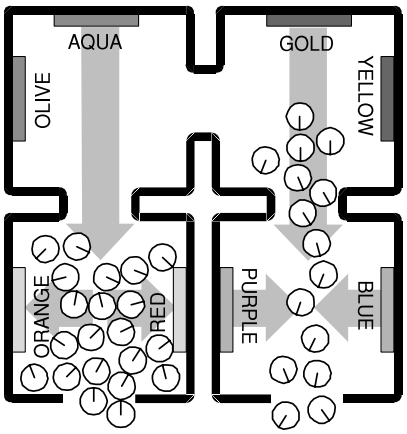


Fig. 9. An illustration of two emergent design problems found when curious agents explore a gallery with an inappropriate arrangement of artworks.

The second problem illustrated in Figure 9 demonstrates the problem of using very different artworks in the final room that cannot be seen from the previous room. The paintings in the last room are discovered upon entrance to be too different from what was expected by the agents. Consequently, the agents entering the final room produce a curious social force that hastens their exit from the exhibition space and the paintings gain little attention.

4.3. Situated design evaluations

The problems illustrated in Figure 9 affect the evaluations reported by the curious agents in a number of ways. First, the interestingness measures reported by the agents will show to be low during their visits due to the long period of time spent in the first room and the short period of time spent in the last room. Second, the support that the gallery provides to satisfy educational goals can also be tested using situated design evaluations. In this example, the educational goal is to expose the agents to a wide range colors during their visit, and the gallery's success in accomplishing this can be evaluated by testing the learned representation of the color space of each agent before and after their visit to the gallery. Well-designed galleries will result in agents having learned a good representation of the color space, similar to that illustrated in Figure 8(a), while badly designed galleries will result in representations of the color space biased toward the color of the artworks found near areas of overcrowding and away from the color of artworks in areas passed through quickly.

Third, in addition to the situated design evaluations the overcrowding in the first room can be assessed using the agent-centric evaluations devised by Helbing and Molnár (1995). A badly designed gallery negatively affects both the efficiency and discomfort measures by reducing the efficient movement of individuals through the space and increasing the discomfort of agents as they try to negotiate crowded areas. The fast transit of agents toward the exit of the gallery can also negatively affect the agent-centric evaluations of the space, as agents with slower desired speeds may be pushed along by agents with a higher desired speed, leading to reduced efficiency and increased discomfort for both agents.

Figure 10 illustrates one possible solution to the emergent deign problems illustrated in Figure 9. In this example, the paintings in the second and third room have been exchanged, resulting in a more gradual progression between the rooms that better suits the curious agents. The consequences of a more appropriate arrangement of paintings is that the agents explore effectively and comfortably, they report a higher level of interest in the artworks throughout their exploration, and as a result, they learn a better representation of the conceptual space of the colors used in the artworks.

5. CONCLUSIONS

This paper proposes that reactive models used for agent-based simulations can be extended with the ability to learn from experience, and that this extension allows agents to provide useful situated design evaluations. In addition, using a model of curiosity, agents can provide useful information about the interestingness of designs that can be used to evaluate designs intended to entertain or educate its users.

The implementation of the simulation of curious individuals exploring a gallery is at an early stage, but already we can see the emergence of interesting crowd behavior. Figure 11 provides some screenshots of different runs of the gallery simulation displaying different crowd behaviors that depend upon the layout of the gallery space and the number of times that the agents are allowed to visit the gallery. The restriction that agents are only allowed to visit the gallery once is achieved by replacing agents that have visited the

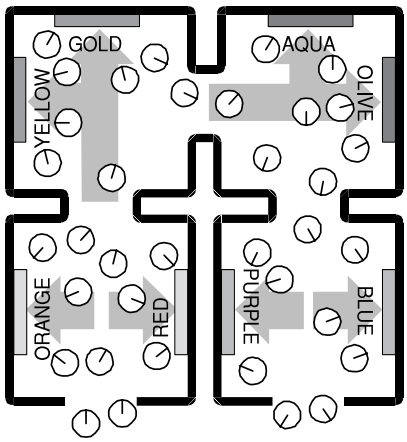


Fig. 10. An illustration of curious agents exploring a gallery with an arrangement of artworks that promotes exploration and learning.

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press

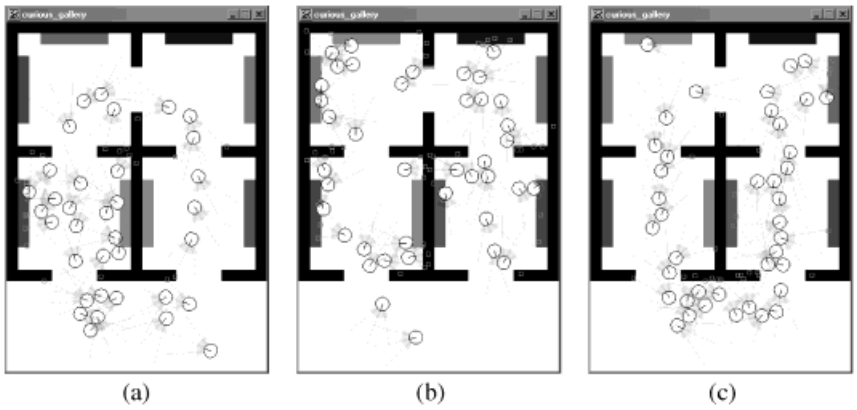


Fig. 11. Screenshots of the curious gallery system running different simulations of gallery arrangements and agent visit strategies. (a) A poorly arranged gallery populated by agents on single visits, (b) a well-arranged gallery populated by agents on single visits, and (c) a well-arranged gallery populated by agents on multiple visits.

gallery with new ones while they are in the open space below the gallery.

Figure 11(a) shows crowd behavior similar to that illustrated in Figure 9 above, where agents become congested around the early works in the gallery but stream past the latter works because the gallery has been poorly arranged. Figure 11(b) shows crowd behavior similar to that illustrated in Figure 10 above, where agents become evenly spread out in the gallery because they find each room equally interesting.

Curious agents have complex behavior that changes over time with exposure to new experiences. The example problem given in this paper of designing an interesting gallery is further compounded if one assumes that agents will visit the same gallery more than once. Figure 11(c) shows the behavior of agents when they are allowed to visit the gallery multiple times in succession and retain the memory of their previous visits. It shows how the gallery eventually becomes "boring" for all of the agents, indicated by the tight formation that the majority of agents have taken while streaming past all of the works because they are no longer of any interest. The question of how a curator maintains the interest of visitors that have already experienced many of the works in previous visits is one that may be explored in future research using curious agents and situated design evaluations.

The example problem of curating a gallery illustrates how the use of curious agents in a simulated environment allows a designer to experiment with different layouts to maximize the interestingness reported by visitors. The use of situated design evaluations opens up new possibilities for using optimization techniques, such as genetic algorithms, to explore the space of possible gallery layouts systematically. Given the complex nature of the group behavior displayed by groups of people, and modeled by curious agents, the use of intelligent design tools that can assist in the planning process would be of great benefit to designers.

REFERENCES

Berlyne, D.E. (1960). *Conflict, Arousal and Curiosity*. New York: McGraw–Hill.

Berlyne, D.E. (1971). *Aesthetics and Psychobiology.* New York: Appleton–Century–Crofts.

Gardenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought.* Cambridge, MA: MIT Press.

Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E 51*, 4282–4286.

Helbing, D., & Molnár, P. (1997). Self-organization phenomena in pedestrian crowds. In *Self-Organization of Complex Structures: From Individual to Collective Dynamics* (Schweitzer, F., Ed.), pp. 569–577. London: Gordon & Breach.

Kohonen, T. (1995). *Self-Organizing Maps*. Berlin: Springer.

Reynolds, C.W. (1987). Flocks, herds, and schools: A distributed behavioural model. *Computer Graphics 21(4)*, 25–34.

Saunders, R., & Gero, J.S. (2001*a*). Designing for interest and novelty: Motivating design agents. In *CAADFutures 2001* (de Vries, B., van Leeuwen, J., & Achten, H., Eds.), pp. 725–738. Dordrecht: Kluwer Academic.

Saunders, R., & Gero, J.S. (2001*b*). Artificial creativity: A synthetic approach to the study of creative behaviour. In *Computational and Cognitive Models of Creative Design* (Gero, J.S., & Maher, M.L., Eds.), Vol. V, pp. 113–139. Sydney, Australia: University of Sydney, Key Centre of Design Computing and Cognition.

Schmidhuber, J. (1991). Curious model-building control systems. *Proc. Int. Joint Conf. Neural Networks*, Vol. 2, pp. 1458–1463. Singapore: IEEE Press.

Schmidhuber, J. (1997). *What's interesting?* Technical Report No. TR-35-97. Lugano, Switzerland: IDSIA.

Woodcock, S. (1999). Flocking: A simple technique for simulating group behavior. In *Game Programming Gems* (DeLoura, M., Ed.), pp. 305–318. Hingham: Charles River Media.

Rob Saunders is a Computer Scientist who works in London. He received his PhD from the Key Centre of Design Computing and Cognition, University of Sydney, on curious agents in design. He has developed a number of systems that exhibit situatedness and curiosity.

John S. Gero is Professor of design science and Director of the Key Centre of Design Computing and Cognition at the University of Sydney. His research spans both design computing and design cognition.

https://doi.org/10.1017/S0890060404040119 Published online by Cambridge University Press