INTERNATIONAL CONFERENCE ONENGINEERING DESIGN
ICED 05 MELBOURNE, AUGUST 15 – 18, 2005

INNOVATION AND DESIGN: COMPUTATIONAL SIMULATIONS†

Ricardo Sosa and John S. Gero

Abstract

This paper explores the role that social evaluation has in the definition of design as a source of innovation. The question of interest is how differences between the evaluative processes of group members can potentially determine the success or the failure of novel ideas in design. Computational simulations are presented of a simple design task in a social system where the effects of different types of group evaluation are analysed in relation to the diffusion of ideas, their differentiation, and the quality ascribed to them by evaluator groups. These results serve to formulate a series of qualitative hypotheses about group evaluation in design as well as practical guidelines for group brainstorming. The current contribution of these in silico studies is to assist in reasoning about innovation and design in a way that conventional research does not support.

Keywords: Creativity, design, innovation, social simulation.

1. Designers and societies

Design has been identified as a key source of innovation [1]. Novel ideas that present advantages over existing solutions can be introduced by designers and can be subsequently materialised by manufacturers, and ultimately evaluated by societies. This can be described as a generative-evaluative process in which initially the designer generates an idea that is evaluated by the immediate group such as the design team or department. If accepted, novel ideas continue through a series of evaluative stages until their materialisation and the ultimate decision-making process of adopters who embrace or reject novel designs. When new ideas trigger group changes, designers can be characterised as initiators of change.

The role of designers in triggering innovation is a difficult subject of inquiry. It is hard to experiment with novel ideas since it is challenging to generate them in the first place and in addition they can only be evaluated by those with no previous exposure, making replication difficult. Traditional research in this area is also constrained by the pro-innovation bias [2]: the implication that innovations in general should be adopted by a majority. Due to this bias, empirical research has focused more on rapidly diffusing innovations, more on adoption that on rejection, and more on continued use than on discontinuance. Computational social simulation is an appropriate method of study here because cases can be repeatedly experimented with in order to reveal the key factors for the triggering of social changes.

In the multi-level evaluative processes of novel ideas in design there is a large number of possible factors that may determine their success or failure. One such factor which has been overlooked to some extent is how different types of group evaluation may affect the role of designers as change agents of their societies. This paper presents a series of computational simulations of a simple design task in a social system that help to understand how differences

† This research is supported by an International Postgraduate Research Scholarship. Computing resources provided by Key Centre of Design Computing and Cognition.

between the evaluative processes of group members can potentially determine the success or failure of novel ideas in design.

1.1. Types of group evaluation

Studies of social aspects of creativity have previously identified the importance of evaluation of novel ideas by groups. Consensual agreement has been identified as one of the best available ways to distinguish creative solutions from their non-creative counterparts [3]. In some cases the evaluation criteria are made explicit to or by evaluators, whilst in other cases an operational notion of creativity as a combination of originality and utility are given without well-defined evaluation criteria. What has remained largely outside these types of studies has been the likely effects that different types of group evaluation may have in determining what ideas are ultimately considered creative. For instance, an aspect of interest is how diversity of evaluation criteria amongst members of evaluating groups could affect the influence of an individual to trigger a group change.

Different aspects of group interaction may be of interest in consensual evaluation. Relevant patterns of social influence include the propagation of errors; hidden profiles; cascade effects; conformity; and group polarisation [4]. These type of phenomena suggest that well-established solutions in design could be explained at least to some degree by group dynamics rather than by their intrinsic characteristics alone. Due to these types of factors, it is feasible to hypothesise that varying levels of individual differences across evaluators may help determine the acceptance or rejection of novel ideas.

An experimental range can be defined between two extreme types of situations of collective evaluation. On the one hand evaluation of novel ideas may take place in groups where evaluators are required to follow well-defined and widely shared criteria. In these cases the role of evaluators is to conduct an explicit and objective assessment and perhaps even include a well-established ranking of priorities collectively defined. On the other hand, members of evaluative groups may be free to apply their individual judgement to a higher degree. A variety of personal values and priorities may have more weight in these situations. This may be the case of more subjective group conditions where limited quantitative measurement is possible and individual idiosyncrasies have a greater impact on evaluation.

Given the range of possible situations between these two extremes, it would be interesting to reveal some fundamental patterns between types of group evaluation and the resulting innovation cycles. In this way we could start to understand what type of design behaviour is required to trigger innovation cycles across fields or disciplines that have different types of evaluation. Likewise, it would be useful to learn how often and under what group conditions innovation cycles are likely to take place given a type of group evaluation. At present, very little is known or has even been asked in these terms. Traditional research tends to place all agency upon the individual designer irrespective of these types of field factors [5].

1.2. Evaluation in teams and in society

The study of patterns of evaluative behaviour could apply in similar ways to small groups, committees, and teams (such as those that evaluate design competitions or a firm's design strategies) as well as to a broader notion of social evaluation, i.e. the market. In the former, team members may play the roles of generators and evaluators at different times for example during brainstorming sessions. In the latter, a range of stakeholders can be identified including opinion leaders, regulators, and consumers. What is common in both cases is the link between individual generation (the designer, the firm) and group evaluation (the team, the market). For the purpose of this paper what matters is that groups can be composed by evaluators that may have different criteria, leading them to varying degrees of agreement and therefore, different types of group judgements.

This paper presents computational simulations of social processes in a simplistic design task that captures some of the aspects of the evaluation and diffusion of novel ideas. An artificial society of agents is defined to conduct experimentation with the aim of better understanding the role of design in innovation. The model studied in this paper includes simple designer agents competing for the adoption decisions of an evaluating group. It also includes a notion of product repository or domain where groups collect selected products that receive high scores. This tripartite model is based on the DIFI model of creativity, i.e., Domain-Individual-Field Interaction [6].

2. Artificial teams and artificial societies

Computational studies of innovation and creativity have focused on three main goals to date: a) to model the processes and products attributed to creative behaviour in humans [7]; b) to develop systems to support individual and team creativity of computer users [8]; and c) to experiment with types of system behaviours that have been associated to creativity and innovation. Computational explorations of the third type consist of designing the components of a target system and their interactions, and observing response patterns to controlled changes. From these explorations the modeller extracts generalisations of qualitative structure about a system which can then be investigated in relation to real situations.

One way to pursue the modelling of creativity and innovation with the aid of computers is to build exploratory computational models that contribute to our understanding by analogy. Computational social simulation is a tool to build hypotheses and discern patterns of behaviour in the social sciences. Simulation of multi-agent systems (MAS) supports a type of in silico experimentation where social behaviour is implemented computationally with the aim of discovering the interaction of local mechanisms and the resulting collective structures of interest [9].

Computational simulation of artificial societies using MAS is appropriate for this type of study because the target system (i.e., innovation) is complex, there are important non-linearities between variables, and there is interest in the dynamics of the system over time rather than in its components in isolation [10]. In recent years the need to model micro-macro causality in more detail and to account for circular and lateral emergence in MAS has been stressed [11]. In this paper we formulate a MAS which accounts for explicit group structures assumed to feed back into individual behaviour in circular or second-order emergence [12].

2.1. DIFI framework

A MAS to investigate design and innovation is proposed based on the DIFI framework (Domain-Individual-Field Interaction) [6, 13]. This framework focuses on the interaction of three system components: domain, field and individual. The individual is represented here by design agents which solve simple design tasks. The domain consists of the set of shared knowledge and selected artefacts shared by members of a group. The field includes groups of individuals who share a common domain. These may be consumers, critics, and other designers. The key implication of the DIFI model is that creativity is a system property and not something that takes place within the head of the designer. Situated in a dynamic environment and in balance between external and personal factors, creative individuals are said to generate "the right product at the right place and at the right time" where *rightness* is largely defined by the society.

Figure 1 depicts the DIFI model where creativity occurs in the relationship between the components of a system [6]. It describes the individual receiving information from the domain and adapting to the constantly changing conditions of a particular field, who evaluates novel ideas and decides to reject or incorporate them into their domain. In general terms, the domain is assumed to transmit information to the person, the person to produce a variation, and the field to include the selected variation to the domain [14].

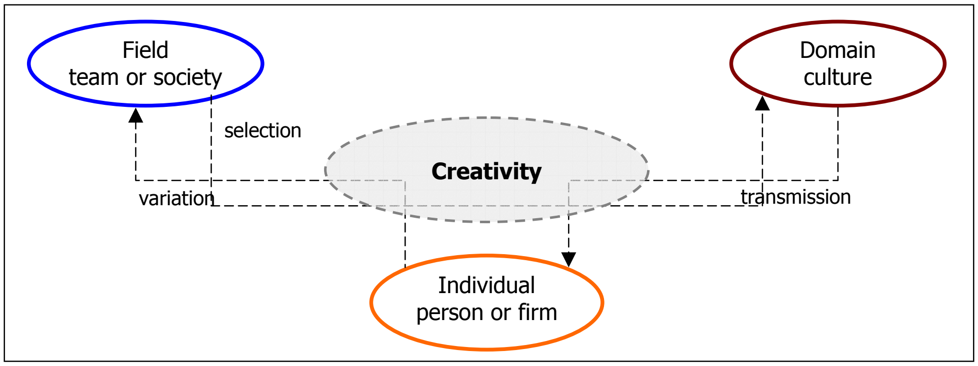

Figure 1 View of creativity as a property of a system of individuals, fields and domains .

This view holds that the only available way to evaluate the creativeness of a generative process produced by a person is social validation. However, the DIFI framework may resolve the issue of *where* to locate creativity, but it still fails to offer any details of what types of processes and how interaction takes place between person, field and domain.

The casual chain in the DIFI map is not likely to be a linear progression from individual variation to social selection to cultural retention and transmission. In a general-systems view the focus is on the linkages of system components where strong complementarities can be expected. Any real understanding of creativity and innovation in design is thus likely to emerge from the investigation of the dynamic interaction of all three levels: the psychological, the social, and the epistemological.

2.2. The system

The system integrates an architecture of social agency with a simple design task based on two dimensional shapes and geometric operations: a simple way to visualise interactions in the system. Here competing designers continuously generate artefacts and evaluator groups evaluate them. Artefact evaluation is based on a multi-objective adoption function that maximises geometric features and novelty, i.e., adopted artefacts tend to be more different from that of the competition.

Stake-holders represented in the system include a number of competing designers, a field or group of evaluators that includes potential adopters and opinion leaders, and a domain or repository of solutions selected by influential field members. Further details about the system that are not central for the scope of this paper can be found elsewhere [5]. The main variables of interest here are defined in Table 1.

Table 1. Main variables of the system

Variable	Description
ω_ε	Perceived artefact feature
υ_ε	Perception threshold

E	Artefact feature representation
ΩE	Artefact line graph representation
$\psi(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots})$	Evaluation function of artefacts
$\rho(\psi)$	Evaluation preference or bias
α_A	Adoption decision
υ_s	Adopters' satisfaction
SDI	Strategic Differentiation Index
μ_{Comp}, μ_{Diff}	Design strategies (competition and differentiation)
T, E_T	Domain and domain entry
T_Ψ	Domain score

2.2.1 Artefacts

In this system, design artefacts are the product of design behaviour and are the subject of evaluation by adopters. Artefacts are kept simple enough to support reasoning for adoption decisions. They are implemented as two-dimensional line representations constrained by 12 boundary points as shown in Figure 2(a). This is a simple way of visually representing features of design artefacts with nomological constraints [15]. This simple design task is a type of layout problem with negotiable constraints, delayed feedback, cumulative solutions, and with a range of acceptable solutions that depend on different viewpoints. Multiple representations and ambiguity are possible because these types of artefacts are perceived and interpreted by evaluators according to a set of randomly distributed perception biases. Figure 2(b) shows sample perceived features of an artefact. The assumption is that people perceive and base their evaluations on different features of design products.

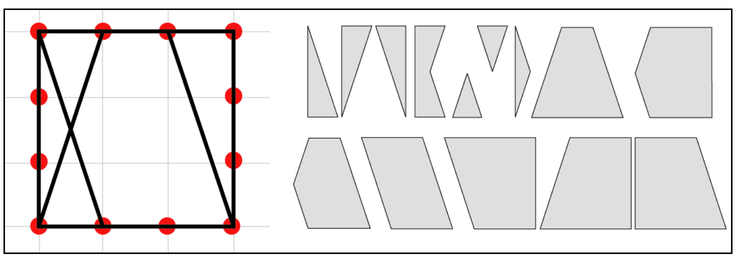

(a) *(b)*

Figure 2 (a) A simple two-dimensional representation, Ω_E, and (b) possible interpretations of an artefact E built as Hamiltonian closed shapes by adopters with individual perception biases.

A perceived artefact in this system is defined as a set E of perceived features ω_ε represented by closed shapes built as Hamilton cycles [16] from a linear representation Ω_E in a two-dimensional space. A perception threshold υ_ε is individually assigned to adopters from a Gaussian distribution with mean and standard deviation as independent variables. With a tolerance of ±1, υ_ε indicates the number of boundary points that a perceived feature ω_ε can include. Therefore, with $\upsilon_\varepsilon = n$ adopters perceive a set $\{\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots}\}$ that includes *n-1* to *n+1* boundary points. Given such set of perceived features, an artefact E is defined by:

$$E = \left\{ \frac{\upsilon_\varepsilon \pm 1}{\Omega_E} \right\}, \{\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots}\} \ . \qquad (1)$$

where an artefact E is perceived as a set of features ω_ε (closed shapes here) given a branch limit of a Hamilton cycle, or perception threshold υ_ε over a line graph Ω_E. Since artefact features ω_ε are represented here by polygons, it is possible to evaluate artefact geometry $\psi(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots})$. Geometric relations $\psi_{\ldots}$ are evaluated between perceived artefact features and are expressed as $\psi_i(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots})$.

Figure 3 shows five artefacts with features that exhibit (a) uniform scale, (b) vertical alignment, (c) rotation and intersection, (d) flip, and (e) uniform scale and rotate. Geometric relations are built based on the perception of the corresponding features, i.e., the scale property of Figure 3(a) is only built if both triangles in Figure 3(f) are perceived. In such case the relation is written as $\psi_{scale}(\omega_{t1}, \omega_{t2})$.

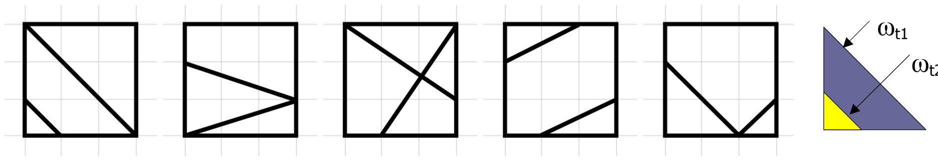

(a) *(b)* *(c)* *(d)* *(e)* *(f)*

Figure 3 Type of geometric properties of perceived artefact features

The evaluation of an artefact $\Psi(E)$ is given by the set of geometric relationships between the perceived features ω_ε used for evaluation.

$$\Psi(E) = \{ \psi_0(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots}), \psi_1(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots}), \psi_{\ldots}(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots}) \} . \qquad (2)$$

An artefact is, in effect, a compromise between conflicting evaluation criteria. Namely, a single artefact cannot contain all geometric characteristics (i.e., some aspects of scale and alignment are mutually exclusive). The design task therefore consists in adapting artefacts to evaluators' criteria. At initial time, artefacts are configured and assigned to each designer. Typically, artefacts are all set to a common configuration from which designers and evaluators start their tasks.

2.2.2 Field

The field here is composed primarily by evaluators or potential adopters. The decisions of evaluators are determined by individual preferences and by group influence. In modelling different design fields more specific assumptions can be incorporated into the evaluation functions. For explanatory purposes, in this system potential adopters evaluate artefacts based on six geometric criteria, which include horizontal and vertical alignment, intersection, rotation, similar bounds, and number of sides.

Individual preferences of evaluators are implemented in the form of biases for every geometric criterion. These biases are assigned at initial time from a random Gaussian distribution with control mean and standard deviation. This individual bias is a source of evaluation disagreements. This is a way to implement the assumption that even if all adopters can agree on performance measures, feature priority is individualised. For instance, we all may agree that a sports car is fast, but speed is an important adoption factor to select a vehicle only for some.

Evaluators assign a preference bias to every geometric criterion such as alignment, rotation, etc. A preference $\rho(\psi)$ is implemented as an evaluation bias between 0.0 and 1.0 for each geometric relation of a perceived artefact. An individual preference $\rho(\psi)$ is implemented as a weight to the evaluation ψ of and artefact.

$$\psi_i + \Delta\rho(\psi_i) \quad (3)$$

where individual preference $\rho(\psi)$ biases the evaluation of a geometric relation ψ by adding a weight assigned at initial time as a random value from a distribution with independent mean and standard deviation.

The evaluation function is translated into the adoption decision as a function of the perceived novelty of the artefact. Perceived novelty is given by the difference per criterion between perceived artefacts. This mechanism promotes the adoption of artefacts with the most distinctive attributes. Given that geometric functions are mutually exclusive, perceived artefacts are expected to receive a high evaluation in some geometric functions and low in others. As shown above, evaluators compare artefacts on every geometric criterion. Once each evaluator adds its own individual preference or bias, the artefact with the criterion with the highest difference from the mean performance in that criterion is chosen. This choice is called the adoption decision, α_A.

$$\alpha_A = \max(\psi_i) \quad (4)$$

where the adoption decision α_A is based on the most distinctive geometric criterion between artefacts.

Therefore, an artefact with a high index ψ in a criterion where all other artefacts have a low evaluation is more likely to be adopted than an artefact with perhaps a higher mean evaluation Ψ if other artefacts also rank high on the same criteria. Since evaluations are based on different perceived features and different individual biases are applied, the population need not converge to the adoption of a single artefact. Likewise, when adoption decisions converge, they need not be based on the same criteria.

Individual preferences of evaluators are dynamic. One way in which evaluation preferences change is by a mechanism of habituation. Evaluators update their preferences $\rho(\psi_i)$ at every adoption decision by marginally increasing their preference for the criterion with the highest score if other artefacts do not perform well in that criterion. This mechanism generates a gradual trend by which adopters increase their preferences for good artefacts as long as these maintain a degree of novelty.

In this framework adopter satisfaction υ_s is used as a simple post-adoption measure of quality. υ_s is a discrete value given by the comparison of adoption decision α_A and individual preference $\rho(\psi)$. It is represented as $\upsilon_s = \{-1, 0, 1\}$. These values correspond to three levels of satisfaction: not satisfied, satisfied, and very satisfied, respectively.

Satisfaction here does not determine adoption, it is a measure of how satisfied is an adopter after it formulates its adoption decision α_A, based on perceived performance, individual biases, and novelty. An adopter may choose an artefact that performs very well in unique criteria and still be dissatisfied.

$$\left\{\begin{array}{l} \rho(\psi) = \max(\Psi(E)) \Rightarrow \upsilon_s = 1 \\ \rho(\psi) > \mathrm{mean}(\Psi(E)) \Rightarrow \upsilon_s = 0 \\ \text{otherwise} \Rightarrow \upsilon_s = -1 \end{array}\right\} \quad (5)$$

If the geometric criterion with highest preference $\rho(\psi)$ receives the highest evaluation score $\Psi(E)$, then adopter satisfaction $\upsilon_s = 1$, i.e., the artefact's features match the preferred criterion. If $\rho(\psi)$ is not the highest of $\Psi(E)$ but is above the mean, $\upsilon_s = 0$, or the adopter is just satisfied. Lastly, if $\rho(\psi)$ is below the mean evaluation component, $\upsilon_s = -1$, or the adopter is not satisfied.

An adopter may be unsatisfied even when the artefact performs well if the artefact is differentiated from others in criteria other than the adopter's highest preference. If the artefact does not perform well or is not different from other competing artefacts, then the adopter abstains to adopt and its satisfaction level $\upsilon_s = -1$ by default.

Social interaction complements the evaluation process and consists of contact with neighbouring evaluators in the field. The aim of this interaction is to influence the adoption decisions of one's neighbours. Evaluators in the field are assigned random positions in different social networks where they are represented by nodes and their adjacency by links or social ties.

The various mechanisms of social influence are given elsewhere [5]. Influential adopters gain the role of opinion leaders in this system. These agents are representative of their groups inasmuch as they have spread components of their adoption decisions. Since there are various social spaces, there are different opinion leaders on every space. This means that there are expert adopters on different topics: some are influential on perceived features, others on preferences, others on adoption choices. Opinion leaders are adopters who concentrate an influence that is one and a half standard deviations above the groups' average.

2.2.3 Designers

A simple design task is assigned to a group of designer agents. The size of this group is determined as a ratio of field size at initial time. During a task, designer agents may develop knowledge in the form of *if-then* rules where the condition is an artefact change and the action is the feedback provided by artefact adoption in the field. As a result of this task designers generate new artefacts which are presented to adopter groups in the field for evaluation. Designers can also evaluate the work of their peers and learn new knowledge from artefacts that receive positive feedback from the field. As designers evaluate each other they modify their peer-recognition levels.

At initial time all design variables are set to zero. The outcomes of the design task are partly determined by the decisions made by adopter group and partly by the ability of competing designers to generate novel solutions. The goal of designers here is to generate artefacts that are chosen by adopters, selected by critics, and imitated by their peers.

Designers may engage in different types of behaviour depending on a number of internal and external factors. Contingent design strategies are defined in this system as the product of the confluence of these conditions. As determined by their strategies, design behaviour may seek to increment adopters' satisfaction levels, to extend the adopter base by capitalising on relative superiority (competition), and/or to maximise differences to other artefacts (differentiation).

Designer agents seek a type of contingent strategy where they learn a *design rule*, i.e. an instance of domain knowledge tied to the artefact representation. Rules are generated based on the designer's model of the population's adoption process construed by retrieving preferences and choices of opinion leaders. This is a way to implement positive feedback since otherwise a designer would not have access to target criteria and target perception, i.e. an opinion leader may be an adopter of a competing artefact or may be abstaining from adopting. A designer can emulate the collective decision process by generating hypotheses of possible alternative artefacts.

Designer agents evaluate their artefacts to decide what features to change in order to increase adoption. This decision determines the strategy to follow. We assume that designers are able to sample their adopter groups. Firstly, designer agents use the mean perception threshold υ_ε of their adopters to model the set of perceived features ω_ε of their artefacts. Secondly, designer agents use the most frequent adoption preference of their adopter group as a target preference. If the artefact of a designer has no adopters, then random perception threshold and criteria preference are set.

Based on the mean perception threshold υ_ε, designer agents simulate the perception and the evaluation of adopters to obtain the estimated artefact performance $\Psi(E)$. The result is an approximate perceived performance of the artefact by criterion $\psi_i(\omega_{\varepsilon 0}, \omega_{\varepsilon 1}, \omega_{\varepsilon \ldots})$. This is rated against the highest adoption preference $\rho'(\psi)$ of the adopter group. Two types of strategies are defined based on this rating: competition and differentiation [17]. In this way the designer agent modifies an artefact based on an estimate of its performance against the preferences of its adopter group.

$$\rho'(\psi) = \mathrm{mode}(\mathrm{P}(\Psi)); \mathrm{P}(\Psi) = \{\rho(\psi)_0, \rho(\psi)_1, \rho(\psi)_{\ldots}\} \quad (6)$$

where the set of adopter group preferences $\mathrm{P}(\Psi)$ is a set of the individual preferences of adopters. $\rho'(\psi)$ is the statistical mode, i.e. the most frequent preference of the group.

A competition strategy is defined by a designer agent when the performance of the main evaluation preference is above the mean of all geometric criteria. In competition, designer agents aim to improve the features that provide good performance on the preferred geometric criterion. For instance, if the main preference of adopters is rotation, designers aim to modify their artefacts to increase rotation relationships. Competition indicates that the designer is likely to increase adoption by gradually increasing the performance of the preferred criterion, i.e., it has a chance to improve.

$$\rho'(\psi) > \mathrm{mean}(\Psi(\mathrm{E})) \Rightarrow \mu_{\mathrm{Comp}} \quad (7)$$

A differentiation strategy is chosen when artefact features perform below average on the evaluators' highest preferences. In differentiation, designers aim to improve features that perform best on different geometric relationships. This is because an implicit criterion of novelty controls the evaluation function, so potential adopters may prefer an artefact that performs well in a different criterion. For instance, if the main evaluation preference is rotation and an artefact performs poorly on rotation but performs very well in alignment, the designer may choose to improve alignment. Differentiation indicates that a designer is not likely to improve an artefact's performance on the preferred criterion but is likely to improve performance on other criteria.

$$\rho'(\psi) < \mathrm{mean}(\Psi(\mathrm{E})) \Rightarrow \mu_{\mathrm{Diff}} \quad (8)$$

At initial time, when all artefacts are set with the same configuration, all designers engage in differentiation strategies until they start finding out what features perform well in the field. Then, designers start to modify their strategies. The field does not converge into adopting one solution: when artefacts get too similar, designers seek to improve other features. Strategies therefore indicate what features of their artefacts designers change (and why). The actual change is based on a type of learning process in which designers generate and apply rules.

Designer agents here are not equipped with creative abilities per se. The aim is not to introduce special traits to assess the effects of agents' creativeness as defined by the experimenter. Instead, all designers are given equivalent sets of mechanisms. No extraordinary process within the individual is hardwired but in time agent interaction renders a social self-organised construct of how a designer may exhibit behaviour considered as creative within its society.

A measure of difference between artefacts is implemented in this system as the Strategic Differentiation Index (SDI): an index estimated collectively by evaluators that reflects the perceived differentiation across available artefacts. With artefacts initialised in a converged state, SDI = 0.0. As designers seek to generate artefacts that differ from each other, SDI becomes greater than 0.

These system mechanisms encapsulate in a simple way some of the characteristics of design problems including ill-structuredness and interpretation; incremental solutions; hypothesis generation; nomological constraints; no right or wrong answers; and delayed feedback [15].

2.2.4 Domain

The last component of the system is the repository or domain. A domain is defined as the collection of design artefacts of a field. Entries to the repository are selected by opinion leaders which emerge from the organisation of adopter groups. A domain characterises the field at a point in time: it can contain a varying quantity and quality of entries as a result of the interaction between designers and evaluators.

The mechanism by which opinion leaders add new artefacts to the domain can be described as a selection of 'better or different' entries. A threshold of entry is set (initialised to zero), which is raised by selected entries. If future artefacts are selected for the same geometric relationships, then they have to receive a higher score. Otherwise, new artefacts can be selected as entries if they receive high values on other geometric relationships.

A domain is defined as a collection of artefacts selected by opinion leaders. At initial time the domain is empty.

$$T = \{E_{T0}, E_{T1}, E_{T\ldots}\} \quad (9)$$

Opinion leaders select entries E_T when their evaluations $\Psi(D)$ of an artefact produce a domain score above an entry threshold Ψ_T. The entry threshold of a repository is initialised at $T_\Psi = 0$. With every entry, Ψ_T is increased to match the value $\Psi(D)$ of the last entry.

A decay mechanism of the entry threshold is implemented so that when no entries are selected by opinion leaders, the entry bar is gradually lowered. The ratio of decay is another parameter for experimentation. In different fields, the rules of selection are assumed to vary.

$$(E_T = \varnothing) \Rightarrow \Delta^{-}\Psi_T \quad (10)$$

With these mechanisms in place, the system can be run over a predetermined number of iterations or until certain conditions are met. The type of design task, the number of agents and other social and domain experimental factors can all be varied. In this paper all conditions are kept constant except for the value of evaluation or adoption biases (w) that adopters apply in their decision-making process.

3. Conformity and dissent in group evaluation

The objective of this experiment is to explore how group evaluation and design may vary when potential adopters have different levels of individual biases in a shared set of evaluation functions. At one extreme, individual evaluation preferences are assumed to play a strong part in the adoption decisions. At the other, individual preferences are weaker so the decisions of the group become more uniform. In the former case adoption is more 'dissocialised' whilst in the latter social influence is stronger.

Evaluation bias (w) is defined in this paper as the strength of evaluation preferences $\rho(\psi)$ among members of a group. The range of experimental values is $(0 \leq w \leq 1)$ so that $w \approx 0$ indicates that individual preferences play a marginal role in evaluation, whilst $w \approx 1$ causes individual evaluation preferences to significantly affect the evaluation process and thus the adoption decisions in a group.

The evaluation bias (w) is analogous to varying the degree of subjectivity across evaluation groups. With certain types of design products evaluation can be expected to be more objective leaving only marginal room for individual biases. This is the case when products are likely to be evaluated by measurable functions such as speed or performance. In such cases the evaluation bias is defined in our experiments as $w \approx 0$. In contrast, other types of products leave more room for interpretation and subjective independent judgements. In such cases the evaluation bias in our experiments is given by $w \approx 1$.

3.1. Experimental setup

With all other field, domain, and design variables in the system kept constant, the experimental variable w is manipulated across 30 cases. Monte Carlo simulations are conducted to explore the range $w = 0$ to $w = 1$ in ten increments. Experimental pseudo-random generator seeds are used to replicate each case. The objective is to assess whether consistent patterns of change are observed in the outcome of the system's components by changing this variable alone.

The behaviour of each designer is analysed through a simulation by recording the number and types of strategies developed, the number of design rules or knowledge that they generate and use to modify their artefacts, and the amount of peer-recognition exchanged between competing designers.

The behaviour of the field is captured by the total adoption of its members, the distribution of these adoption choices (who adopts what artefact), and by the history of satisfaction levels in the field. Lastly, dependent variables of interest in the domain include the number of artefacts selected as domain entries and the scores assigned to them by opinion leaders.

Results for this experiment are obtained by comparing the mean values of thirty cases for every increment of evaluation bias w. Simulations presented in this paper are run for a period of 2,500 iterations with three designer agents and a field size of one hundred individuals.

3.2. Results

This experiment shows a variety of effects that differences in the evaluators of solutions can have in systems of this type. The main effects are observed in cumulative adoption and adoption distribution; adoption satisfaction; perceived differentiation; and scores assigned to domain entries. Results show that increasing the strength of individual biases has a number of effects. As the evaluation bias is manipulated, evaluators vary their selectivity of adoption choices and vary their frequencies of adoption. They also change their levels of satisfaction as the evaluation bias is changed. The way designer agents choose their strategies is also consistently affected, resulting in artefacts that receive different levels of differentiation by evaluators. Finally, the selection of domain entries by opinion leaders varies as a result of manipulating the evaluation bias, in particular the scores assigned to these entries. In the following section results are presented for the two extreme cases where w = 0 and w = 1 for clarity.

4. Effects of group evaluation

4.1. Adoption and abstention

The first result of varying evaluation biases across groups is related to adoption satisfaction and to total or cumulative adoption. Figure 4 shows the effects of the evaluation bias w = 0 and w = 1 against adoption satisfaction levels and aggregate adoption. The first result in Figure 4(a) is rather straightforward: adoption satisfaction increases significantly with evaluation bias. This can be expected to occur because the adoption decision is directly biased by individual preferences, so group members tend to adopt artefacts with features that they prefer. As individual preferences have more weight in the adoption decision, evaluators tend to adopt what they prefer more often. In contrast, when their preferences are not strong enough they may adopt artefacts with high fitness in features for which they have low preference. However, the predictability of satisfaction levels also changes as a function of evaluation biases: satisfaction variance is directly proportional to evaluation bias. Namely, with more independent adoption decisions, satisfaction is likely to be higher but it is also less stable and less predictable.

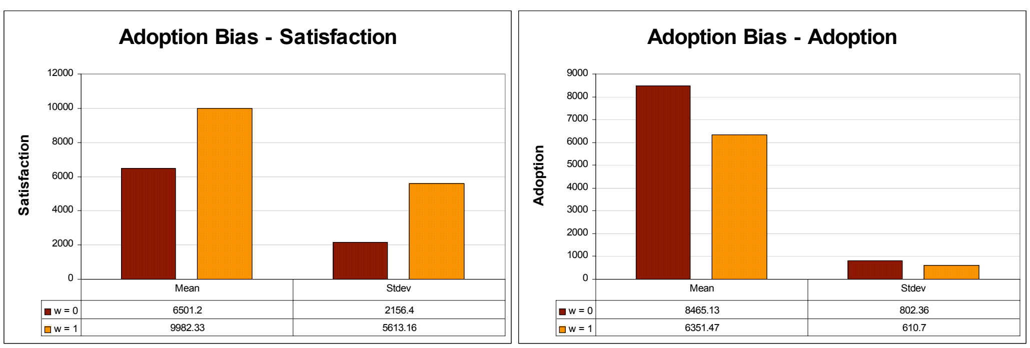

(a) (b)

Figure 4 Effects of evaluation bias in satisfaction levels and in cumulative adoption. When evaluators make more independent choices, or w = 1: (a) their choices yield more unpredictable but higher satisfaction levels and (b) they adopt less artefacts or abstain more.

Somewhat unexpectedly, as Figure 4(b) shows, cumulative or total adoption is negatively related to evaluation bias. When adopters have low adoption preference biases (w = 0) total adoption is higher than with high evaluation biases (w = 1). In other words, abstention increases as adopters fail to see differences between artefacts.

This seems paradoxical: adopters are more 'free' to make their choices and they are more satisfied with these choices, yet they adopt less than when they are constrained by artefact features and social influence. A key to this apparent contradiction is in adoption variance, defined as the distribution of adoption decisions between designer agents. When adoption variance is high, adopters tend to concentrate their choices in artefacts from a few designers, whereas a low adoption variance refers to a more competitive environment where adopters distribute their choices across all designers.

When $w = 0$, mean adoption variance is 0.43 whilst for $w = 1$, mean adoption variance falls to 0.33. This shows that when individual evaluation biases are stronger, adopter agents select artefacts from a smaller number of designers, presumably those that best satisfy their preferences. Therefore, when adopters have more independence on their adoption choices, they adopt less but are more satisfied with their choices.

4.2. Idea differentiation

Effects of evaluation bias (w) on two aspects of design behaviour are shown in Figure 5. The strategic differentiation index (SDI) measures the difference of artefacts as perceived by adopters. The mean SDI of a simulation run is recorded for every case. Figure 5(a) indicates mean and standard deviation of SDI for cases where $w = 0$ and $w = 1$. The number of strategies chosen during every simulation run is recorded for every designer. Figure 5(b) shows the aggregate for each strategy of all designer agents.

SDI is negatively correlated with evaluation bias showing that perceived differentiation decreases when evaluation bias is high, $w = 1$. This is not only a change in adopters' perception but an actual variation of strategies by designers as observed in Figure 5(b). When adopters are assigned high evaluation biases ($w = 1$), designers adapt their behaviour towards strategies of competition and differentiation. An explanation for this consistent effect may be that when adopters take more independent decisions ($w = 1$) designer agents are less likely to perceive that their artefacts address the population's preferences. This could be due to the reduced impact of the social influence mechanism, i.e., less group agreement emerges.

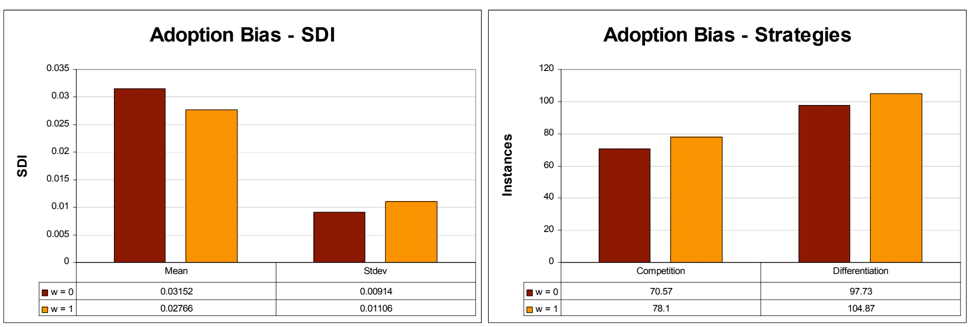

(a) (b)

Figure 5 Effects of evaluation bias in SDI and design strategies. In cases where evaluators take more independent decisions (w = 1), (a) design artefacts are perceived as more similar and (b) designer agents engage in more competition and more differentiation.

4.3. Quality ascribed to ideas

The last variable considered is the score assigned by opinion leaders to domain entries. These are entered during a simulation run with their assigned score and time of entry. Figure 6 shows mean score and standard deviation of all cases for $w = 0$ and $w = 1$. Evaluation bias is positively correlated with score, i.e., in cases where adopters take more individualised decisions or $w = 1$, domain scores are higher.

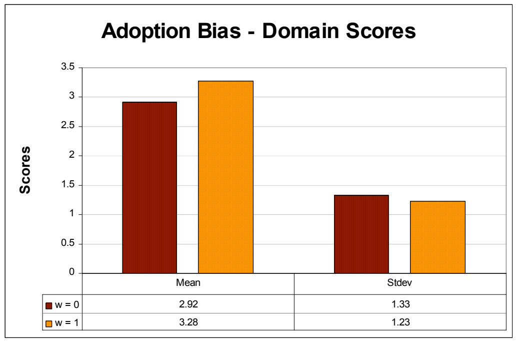

Figure 6 Effects of evaluation bias in domain score. In cases where adopters take more independent decisions (w = 1) gatekeepers tend to assign higher scores to selected artefacts.

No associated effects are seen for other domain variables such as the number of entries. This effect could be attributed to the individual skew produced by the impact of the evaluation bias in the selection of artefacts by opinion leaders, which are also evaluators and are therefore affected by changes in evaluation biases. In sum, when opinion leaders in a group have higher preference thresholds, the merit of selected artefacts can be expected to be perceived as higher.

5. Discussion

The main consequences of individual evaluation biases in the system under analysis are that novel ideas are likely to be better accepted or become more ‘popular’ in groups where evaluation biases are low than in those where individual biases are stronger. Popularity refers in this experiment to the cumulative or total number of adoptions.

This implies that novel ideas may be adopted by a larger majority in groups where adoption choices are less independent and social influence is stronger. This could seem at first paradoxical inasmuch as conformity and uniformity are intuitively regarded as inappropriate for creativity. However, under closer scrutiny it is easy to understand how higher levels of group agreement may in fact facilitate unanimous adoption decisions, including the acceptance of new ideas. In groups where evaluation criteria are shared, adopters of novel ideas may be less satisfied with their choices, and experts may also assign lower values to domain artefacts.

These effects represent important ways in which the individual characteristics of evaluators may determine the course of diffusion as well as who is regarded as creative in a group, and when. Although these results are constrained by the assumptions embedded in this model of artificial societies, they show consistency with the literature.

In a thorough study of seven creative figures, Gardner [18] concludes that uniformity in the evaluation of novel ideas facilitates the emergence of prominent figures. This is replicated in our experiments firstly by showing that groups with stronger shared biases exhibit an increase in cumulative adoption. Arguably stronger group biases means that the view of a few evaluators influence the adoption choices of the rest. Secondly, these experiments show that in such types of groups the perceived differentiation of competing products also increases. This reinforces the idea that when evaluators are unable to emphasise their individual biases, their choices become normalised and this supports the rise of a few prominent designers.

However, our experiments exhibit an opposite effect in regards to domain scores. According to Gardner [18], increased adoption should be accompanied with higher value ascribed to novel ideas in cases with stronger group biases. In our experiments we obtain a different pattern: when individual evaluation biases are weaker, domain entries are ascribed lower values. It will be the subject of future work to determine the type of selection mechanisms that are likely to affect this outcome.

In relation to group brainstorming, these experiments suggest some practical guidelines. During idea discussion, if members develop an agreement on explicit criteria for evaluation, novel solutions are likely to receive more support by evaluators. In these conditions, group members are also likely to perceive higher variety of ideas, and more ideas will tend to originate from a smaller number of participants.

The relation between computational modelling and target socio-cognitive systems remains an open question. In this particular type of modelling there is a characteristic tension between transparency and validation or veridicality [19]. Future work will address the comparability between the computational model and experimental and field evidence.

References

[1] V. Walsh, R. Roy, M. Bruce, and S. Potter, *Winning by Design: Technology, Product Design and International Competitiveness*. Cambridge: Blackwell, 1992.

[2] E. M. Rogers, *Diffusion of Innovations*, 4th ed. New York: Free Press, 1995.

[3] T. Amabile, *Creativity in Context*. Boulder: Westview Press, 1996.

[4] P. B. Paulus, *Psychology of Group Influence*. Hillsdale: L. Erlbaum Associates, 1980.

[5] R. Sosa and J. S. Gero, "A computational framework for the study of creativity and innovation in design: Effects of social ties," in *Design Computing and Cognition '04*, J. S. Gero, Ed. Dordrecht: Kluwer Academic Publishers, 2004, pp. 499-517.

[6] D. H. Feldman, M. Csikszentmihalyi, and H. Gardner, *Changing the World: A Framework for the Study of Creativity*. Westport: Praeger, 1994.

[7] P. Langley, H. A. Simon, G. L. Bradshaw, and J. M. Zytkow, *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge: MIT Press, 1987.

[8] B. Shneiderman, *Leonardo's Laptop: Human Needs and the New Computing Technologies*. Cambridge: MIT Press, 2002.

[9] J. M. Epstein and R. Axtell, *Growing Artificial Societies: Social Science from the bottom up*. Washington, D.C.: Brookings Institution Press, 1996.

[10] R. Conte, R. Hegselmann, and P. Terna, *Simulating Social Phenomena*. Berlin: Springer, 1997.

[11] R. Conte, B. Edmonds, S. Moss, and R. K. Sawyer, "Sociology and social theory in agent based social simulation: A symposium," *Computational and Mathematical Organization Theory*, vol. 7, pp. 183-205, 2001.

[12] G. N. Gilbert, "Varieties of emergence," in *Social Agents: Ecology, Exchange and Evolution*, C. Macal, M. and D. Sallach, L., Eds. University of Chicago: University of Chicago, 2002, pp. 41-50.

[13] M. Csikszentmihalyi, "Society, culture, and person: a systems view of creativity," in *The Nature of Creativity*, *Contemporary Psychological Perspectives*, R. Sternberg, Ed. Cambridge: Cambridge University Press, 1988, pp. 325-339.

[14] M. Csikszentmihalyi, "The domain of creativity," in *Theories of Creativity*, M. A. Runco and R. S. Albert, Eds. Newbury Park: Sage Publications, 1990, pp. 190-214.

[15] V. Goel, "A comparison of design and nondesign problem spaces," *Artificial Intelligence in Engineering*, vol. 9, pp. 53-72, 1994.

[16] F. Rubin, "A search procedure for Hamilton paths and circuits," *Journal of the ACM*, vol. 21, pp. 576-580, 1974.

[17] P. M. Nattermann, "Best practice does not equal best strategy," *The McKinsey Quarterly*, vol. 2000, pp. 38-45, 2000.

[18] H. Gardner, "The creators' patterns," in *Changing the World, a Framework for the Study of Creativity*, D. H. Feldman, M. Csikszentmihalyi, and H. Gardner, Eds. Westport: Praeger, 1994, pp. 69-85.

[19] K. M. Carley, "Simulating society: the tension between transparency and veridicality," in *Social Agents: Ecology, Exchange and Evolution*, C. Macal, M. and D. Sallach, L., Eds. University of Chicago: University of Chicago, 2002, pp. 103-113.

Corresponding author: Ricardo Sosa Medina, Key Centre of Design Computing and Cognition, University of Sydney NSW2006, Australia. Phone 61+290365493, r.sosa@usyd.edu.au, www.arch.usyd.edu.au/kcdc.