Creative Social Systems

Ricardo Sosa

rdsosam@itesm.mx
Departamento de Diseño
Instituto Tecnológico de Estudios Superiores de Monterrey, Mexico

John S. Gero

john@johngero.com
Krasnow Institute for Advanced Study and Volgenau School of Information Technology, George Mason University, USA

Abstract

In this paper we present an approach based on the view that creativity is essentially determined by social evaluation and that any computational modelling of creativity needs to account for this. Directions for modelling creative social systems are suggested based on results of computational social simulations of creativity and innovation.

Creative Systems

A mainstream assumption in creativity research is that creativeness is an exceptional ability or process by which original, useful and unexpected ideas are generated. The human capacity to produce ideas that radically change current dominant practices or beliefs in a social group is still today mostly unknown. Despite its social and historical dimensions, creativity is still largely taken for granted as something that takes place inside the head of distinctive people when they create or ideate new solutions.

This non-trivial emphasis on individual causality has led an increasing number of researchers to adopt the aim to emulate creativity artificially by capturing the underlying generative processes in a computer program. Under this dominant approach, all emphasis has been given to the synthesis of creativity assuming thus that its evaluation is limited to passive acknowledgement or recognition by their peers or some other kind of evaluators. In other words, it is often assumed that to define something as creative is a direct and obvious consequence of a special generative process that one must appreciate if well informed.

As a result, the output of a typical computational model of creativity is assessed by a group of experts who usually ascribe its creativeness independently of its underlying generative model. It follows from this view that creativeness is approached in an absolute way, without articulating a particular context in which synthesis and evaluation may interact.

One may challenge this dominant view by drawing analogies to other phenomena where rich interactions between complementary individual roles produce emergent results. For instance, in the same way as learning is not what necessarily happens as a direct result of teaching, defining an idea as creative needs not be a direct and exclusive consequence of an assumed special characteristic of the person or the process that generates that idea. It may be that in the attribution of creativeness to a person or an idea, evaluation plays a more active and crucial role than is currently assumed. Thus, it could be of interest to rethink the way creative systems may be developed in the future.

In this paper we approach the modelling of creativity in computational systems by studying the interaction between individuals that generate and introduce new ideas, and societies that collectively evaluate and decide to adopt or reject those ideas. Thus, rather than aiming to replicate or emulate the creative synthesis of solutions in painting, writing, music, or other creative tasks, we use computational modelling as an experimental testbed to implement and understand the complementary roles of generation and evaluation of ideas in social systems.

Computational Social Systems

Computational social simulation refers to the study of social agency through the design, implementation, and execution of computer models usually built under rather simple assumptions that, nonetheless, help to understand complex emergent results. The experimenter is able to define a series of hypotheses, implement and experiment with them to explore their consequences and veridicality. Early computational social simulations emerged from game theoretical approaches. Recently, a variety of modelling tools and frameworks have been developed by a growing research community, including useful multi-agent system toolkits (Conte et al. 2001).

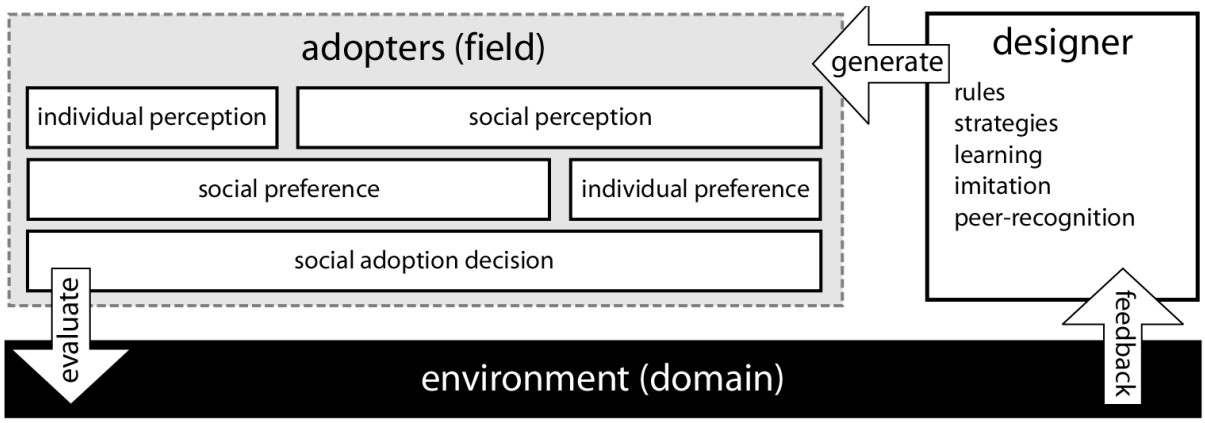

DIFI Multi-Agent Framework

In recent years we have developed a computational framework to study creativity at the relationship between individual agency and social groups (implemented either as multi-agent systems, MAS, or cellular automata, CA). Agents in this framework generate and evaluate a range of ideas and practices (represented by numeric values or geometrical shapes) by executing a set of simple rules (value exchange between agents in *n*-dimensional landscapes). This modelling is framed within a systems view of creativity and innovation: the Domain-Individual-Field-Interaction or DIFI framework (Feldman, Csikszentmihalyi and Gardner 1994).

In this application of the DIFI framework, the domain represents the set of values shared by a field, the field is defined by the aggregate characteristics of the group of agents and their interaction over time, and the individual by the set of behaviour rules for different agent roles. By manipulating independent experimental variables at these three different levels, we are able to explore in our computational models the formation of patterns over simulated time of social influence, diffusion, and emergence of new values.

The role of designers in these systems has been modelled as instances of change agents that work towards providing novel solutions to a set of problems shared by large social groups. Typically, in these social simulations a small subset of the population are designer agents, i.e. they compete over a time period by iteratively interpreting the problem and proposing a solution which is conversely evaluated by the rest of the social group.

The designer agents learn from the feedback provided by the social group including their adoption decisions and a measure of satisfaction with their adopted solutions. Designer agents also have a learning mechanism that influences their future behaviour based on the overt actions of their competitors and the social adoption of their solutions.

Although the evaluation process carried by adopter agents follows a set of rules that define individual differences such as perception and preferences (following a normal distribution), social interaction is included as the potential of adopter agents to influence each other's decisions to either adopt or reject solutions generated by the designers.

Figure 1 illustrates the architecture upon which our systems are built including the type of behaviour rules of individual designers, the layers of interaction in fields or social groups of evaluators, and the resulting domain or set of aggregate solutions.

In groups which comprise a few thousand adopters, patterns of interest arise such as the emergence of opinion leaders and cycles of convergent-divergent adoption. During a simulation, the system tracks the behaviour of every agent as well as the global patterns of group behaviour. Despite their apparent simplicity, these multi-agent models (MAS) generate non-linear effects that emerge from the interaction of their components over time. In this way, researchers are equipped with in silico laboratories where they can 'grow up' different states from a set of initial conditions, gaining insights into the role of designers as change agents of complex social systems.

[Figure 1 here]

Figure 1 System architecture based on Domain-Individual-Field Interaction (Feldman, Csikszentmihalyi and Gardner 1994).

Typically, computational social simulations of this sort have focused on the aggregate effects of individual behaviours interacting in large groups and over large time spans. For instance, the canonical social agent simulation consists of a population of agents represented by cells in a two-dimensional grid which are individually assigned initial numeric values at random and the simulation run shows how a simple instruction to exchange values between adjacent individuals, tends towards group convergence in one or more dominant values (Axelrod 1997).

Such systems have been extensively replicated showing that the final outcome (group convergence) is unavoidable although variations do occur depending on variables such as the range of values assigned and the rules of interaction between neighbours or adjacent cells. When group characteristics are studied in this type of models, the main focus has been on the manipulation of group size showing that the size of a society may only determine the time length to group convergence but not the final result.

A converged society represents one in which a dominant value is shared by all or a majority of individuals. This could represent a generalised practice or belief throughout that society such as a cultural feature of a group. In order to change a dominant practice or value, single agents or minorities may intervene by generating alternative values and presenting them to their social group for evaluation. Previous models had shown that this phenomenon of triggering a social change by a minority in these systems is unlikely to succeed under most circumstances due to the high probability of the new value to be 'eaten' by the continuous exchange of the dominant or wide-spread value amongst most neighbours (Axelrod 1997).

However, modified systems can be implemented to track the diffusion or spread of new alternative values, showing as a result a number of variables that may have a direct influence on how difficult it is for an individual or a minority to diffuse a new value in their society replacing dominant values.

Experimental Settings

In a DIFI-inspired framework to study creativity and innovation, variables can be defined and implemented at three levels: roughly, differences between individual attributes of agents, differences of field characteristics such as social structure, and differences between domains such as the nature and complexity of the task at hand. At implementation, key decisions are made regarding the selection of programming languages, multi-agent development kits, agents' behaviour, tasks or problems to be solved and evaluated by agents, and data capture and analysis. As a result, validation of this type of systems is a key issue that has been discussed extensively in recent years (Conte et al. 2001). To develop our systems we have iteratively tested and developed a number of alternatives discussed at length elsewhere (Sosa 2005).

Experimental variables in this framework can be defined and manipulated by the experimenter to build and test hypotheses. These variables define the characteristics of individual designers, or properties of their interaction with adopter groups, or domain rules. The aim is to vary conditions and run a number of simulations (i.e., cases) a number of times in order to observe patterns that link initial conditions with observed behaviour. Keeping every other condition constant, the experimenter can attribute causal associations. Within a society, a number of competing designer agents are initialised. The number of designers can be varied as well as the frequency and other parameters of their behaviour.

Individuals, Field and Domain. At the person level, individual differences are implemented through probabilistic distributions in the multi-agent system (MAS) representing preferences, abilities, and learning mechanisms that define each agent's behaviour. The roles of designer and adopter agents are assigned by the experimenter as a ratio based on what is usually defined as the "creative industries" in occupational census.

A decision-making process of interest in design is the adoption of solutions. Design artefacts (implemented as geometric compositions with features such as rotation, scale, angles, etc) are assumed to be evaluated by a population in two complementary ways. Individually, adopters form their own feature perceptions and develop their own geometrical preferences. In a group, adopters exchange, compare, and influence each other on their perceptions, preferences, and choices. These shared elements constitute what is common to members of a community or a group, and is subject to trends of influence in a social space.

A social space or social network is defined in this framework as an arrangement of neighbouring agents, and can be implemented in matrices or arrays where every agent has a specific location. However, social interaction can be assumed to take place in more than one social environment simultaneously, so they are organised in a number of social spaces where they have different connections with other agents (neighbours) and different ways of interacting with them, i.e., different positions within kinship, work, acquaintances, and other types of social environment or space.

Agent interaction on every social space can be modelled with different parameters, i.e., different extroversion thresholds and rules of interaction. In our studies we typically implement three social spaces based on the non-convergent property of >2-dimensional random walks.

Social spaces are also charecterised by ties, i.e., linkages between nodes in a social network. These links determine what adopters (nodes) have contact with each other. The strength of social ties refers to the likelihood that nodes in the social network are maintained over time. Strong ties are characteristic of resilient social relationships such as kinship or friendship, whilst weak ties characterise temporary social networks such as school peers or travel acquaintances. In networks with strong social ties, adopter agents maintain contact with each other over longer time periods, whilst in networks with weak ties adopter agents constantly change contact with different neighbours.

The role of designers is assigned to agents that aim to solve tasks that consist of generating artefacts that adopter groups evaluate based on a multi-objective adoption function that maximises key domain features and promotes novelty.

Designer agents generate new artefacts by learning and applying rules that increase their adoption. If appropriate, they may also imitate more successful competitors. The performance of designers in this type of tasks is estimated along the following dimensions: a) the size of adopter bases, b) the number and type of design rules generated, c) the influence on other designers' work, and d) scores given by experts. These are illustrative indicators of creative performance that provide a simple account of phenomena such as popularity, quality, novelty, peer-recognition, and expert juries –factors usually associated to creative benchmarking in the literature.

A multi-objective adoption function guarantees that in this type of tasks there is no single best solution. Constantly evolving preferences of adopters over time causes designers to continually update or replace the artefacts that they introduce to their social groups for evaluation. Whilst there are clear evaluation criteria based on geometric features, the appropriateness of artefacts in this framework is given by the existing conditions of the population as well as the actions of other designers at a particular time. Designers develop artefacts by modelling the evaluation decision of target groups. Feedback is only obtained after the artefact is made available to the society. There are fixed constrains given by the geometric representation.

The domain in this framework consists of geometric representations. It is also complemented as a cumulative collection or repository of selected designs. These entries to the repository are selected as an aggregate effect of field interactions, i.e., gatekeepers (field agents that collect high influence levels) emerge from the interaction of adopter groups. A repository thus characterises a society over time: it can contain a varying quantity and quality of entries as a result of the interaction between designers, artefacts and adopters. The mechanism by which gatekeepers add new artefacts to the repository can be described as a selection of 'better or different' entries. A threshold of entry is set (initialised to zero), which is continuously raised by previously selected entries. Namely, 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.

Parameters for experimentation in this framework include individual differences between designers on the generation of rules and their use. Temporal parameters include the rates at which designers modify their artefacts and the frequency at which gatekeepers evaluate artefacts. A type of normative parameter is the type of access that designers have to existing design rules: in public mode designers can use the rules generated by all others, whereas in private mode they can only use the rules they generate. Another available parameter for exploration is the weight of individual adoption preferences in the adoption decision process. Lastly, rules of social interaction and the size of adopter populations are available for experimentation.

Whilst these parameters represent only a small part of the large number of issues around the theme of creative design, they support experimentation with the following types of questions: How well do individual differences of designers predict creative performance? What are the creative effects of having a higher rate of design activity in a population? Or more frequent gatekeeping? How do different types of access to knowledge affect the generation and adoption of creative contributions? Under what circumstances are adopters likely to be more satisfied? When will a creative design concentrate large adopter groups? Or receive more peer recognition? When will domains be larger or contain entries with higher scores of creativeness?

Results

The most significant results obtained from experimentation with this framework are related to the varying social impacts that certain individuals may have depending on field and domain factors independently of the actual attributes of the individuals or the actions which trigger the change. Such results are of interest because they shift the emphasis from the generation to the evaluation stages of creativity.

For instance, individuals or minorities with diverging ideas from their social groups may exhibit a short time span due to group influence during social interaction. However, in other similar cases, minorities may stay isolated when the degree of dissent is high enough to impede interaction –and in only a few cases, dissenters may exert influence over the entire group triggering a social change. The key implication is that the same individual behaviour may have entirely different global consequences not due to the specific actions, but due to its coupling in space and time with the rest of the system.

This type of results illustrate the relative importance of individual differences in creativity. Although it is possible that talent may indeed play a key role in determining individual actions, it has no direct relation to the social consequences of those actions. Persistence may thus be of special importance inasmuch as the same individual action may have distinct social impacts at different times.

A number of field characteristics further illustrate situational conditions that can determine the creativeness of individuals or their new ideas. These include the strength of ties in a society that evaluates a new idea (Sosa and Gero 2008) or the population size of the social group (Sosa 2005). Domain factors that may become determinant include the compatibility of new ideas with dominant values and practices (Sosa and Gero 2007) and the complexity and ambiguity of new ideas (Sosa 2005).

Conclusions

The Fundamental Attribution Error (FAE) is well documented in social psychology research to indicate that people tend to overestimate the role of individual traits and underestimate the importance of the situation and context when interpreting other people's behaviour. At present, we suggest that the field of computational modelling of creativity shows a pervasive Fundamental Attribution Error by entirely placing causality on the generative algorithms.

On the grounds of putative talent, Simon (2001) insists on placing extraordinary achievements toward the upper end of a continuum of common human mental activities, without claiming that they are qualitatively different from more commonplace behaviour. Similarly, creative systems need not be models of extraordinary talents, rather the dominant individualistic paradigm may be complemented

by a deeper understanding of the social effects of common individual processes.

Creative Social Systems (CSS) can be distinguished in general from the mainstream paradigm of Creative Individual Systems (CIS) by a shift of focus regarding the causality of creativity. Whilst the latter aims to model artificial creativity exclusively as an autonomous and individual generative process, the former's aim is to embed agents in complementary roles of generative-evaluative interaction from which micro (individual) and macro (structural) phenomena contribute to shape situations that determine how an idea or its proponent are considered as creative by their peers or evaluators. Table 1 presents the main differences between CSS and CIS approaches.

Table 1. Creative Social Systems vs. Creative Individual Systems

Creative Individual Systems	Creative Social Systems
Intent is to constantly seek autonomous creative output	Intent is to interact with other systems and as a result produce lots of failures and some hits
Output is assessed by an external system (humans) applying assessment criteria used to define creativity in humans	The evaluation of its generative processes is carried within the system and becomes input to future generative cycles
Creativeness is directly attributed to the system and compared to human achievements	Creativeness is product of interactions within the system and it is only relative to the domain embedded into the system
Emphasis is exclusive on generative phases. Seeks to discover creative processes of synthesis	Causality is balanced between generative and evaluative phases. Seeks to understand how creativity emerges within a particular context.
Individuals have special attributes or talents and are sufficient to explain social impact	Interaction between individuals and their situations create special conditions that yield social change and individual recognition
Individual actions are assumed to constantly seek originality, utility and unexpectedness	The previous experience of an individual and its context determine actions. These may include behaviour that triggers social change
Systems are "solipsistic", i.e., closed and independent of external influence	Systems are socially aware, actions are interdependent upon the behaviour of others
System generates output without modelling an expectation of the reaction by the evaluator	Systems may exhibit high degrees of persistence and be prepared to repeat their output expecting to get different responses within varying situations

References

Axelrod, R. 1997. The Dissemination of Culture: A Model with Local Convergence and Global Polarization, *The Journal of Conflict Resolution* 41(2): 203-226.

Conte, R., Edmonds, B., Moss, S., and Sawyer, RK. 2001. Sociology and Social Theory in Agent Based Social Simulation, *Computational and Mathematical Organization Theory* 7: 183-205.

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

Sosa, R. 2005, Computational Explorations of Creativity and Innovation in Design, PhD dissertation, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia.

Sosa, R., and Gero, JS. 2007. Computational Explorations of Compatibility and Innovation, in Leon-Rovira, N ed. *Trends in Computer Aided Innovation*, Springer, 13-23.

Simon, HA. 2001. Creativity in the Arts and the Sciences, *The Kenyon Review* 23(2): 203-220.

Sosa, R., and Gero, JS. 2008. Social Structures that Promote Change in a Complex World: Complementary Roles of Strangers and Acquaintances in Innovation, *Futures* 40(5). Forthcoming.