Elementary Social Interactions and Their Effects on Creativity: A Computational Simulation

Andrés Gómez de Silva Garza[1,1] and John Gero[2]
[1] Instituto Tecnológico Autónomo de México (ITAM), Río Hondo #1, Colonia Progreso Tizapán, 01080—México, D.F., México, agomez@itam.mx
[2] Krasnow Institute for Advanced Study, 4400 University Drive, Mail Stop 2A1, George Mason University, Fairfax, VA 22030, U.S.A., john@johngero.com

Abstract. This paper presents a multi-agent computational simulation of the effects on creativity of designers' simple social interactions with both other designers and consumers. This model is based on ideas from situated cognition and uses indirect observation to produce potential changes to the knowledge that each designer and consumer uses to perform their activities. These changes result in global, social behaviors emerging from the indirect interaction of the players with relatively simple individual behaviors. The paper provides results to illustrate these emergent behaviors and how the social interactions affect creativity.

1 Introduction

Computational models of creativity typically simulate the reasoning process of one creative agent that produces designs (whether they are of residences [1], stories [2], software [3], artwork [4], or other artifacts, whether abstract or physically realizable). In this view, the simulation ends as soon as the design of the artifact is generated by the simulated designer, and this design generation concludes when the simulated design process converges to an acceptable solution. The simulated designer is programmed with a particular body of knowledge, which may or may not change over time, that embodies its expertise and that includes evaluation knowledge that allows the agent to determine the acceptability of the designs it proposes and therefore halt the simulated design process.

This simulated design process might have some parameters that can be adjusted, but usually employs the same overall method in order to generate designs (whether it be evolutionary algorithms [5], analogy [6], constraint satisfaction [7], shape grammars [8], or other computational strategies for creating designs). If the design process is run again on the same problem, then the same solution, or at least the same type of solution, will be obtained as output. This is the "design as search or optimization" view, and does not account for the fact that most designers are able to continue producing creative output throughout their lives. Designers do not just produce one design and stop, and the designs that they have produced in the past

[1] Contact author.

influence the ones they produce in the future, instead of each episode of producing a design being done in complete isolation from all other episodes.

As Boden has pointed out, creative products must be both novel and useful/valuable [9]. In order for a designer's output to be considered creative it must be sufficiently distinct from the designer's previous body of work, and in order for this to occur in a computational model, the simulated designer must be dynamic: at minimum, some aspect of the way it analyzes and/or produces designs must change over time. In addition, in order for a designer's output to be considered creative it must be valued by others, and in order to be successful in this a designer must be aware (as much as possible) of what others look for or value. It thus appears that traditional computational models of design are limited in their veridicality because they do not take into account a designer's social context in modeling design activity and the factors that drive it.

This paper describes a computational model that embodies a broader view of design than single designer simulations. In this view the design decisions that a designer makes (*i.e.*, the evaluation criteria on which those decisions are based) are influenced by multiple factors, some of which are external to the agent. In particular, the knowledge that a designer uses both to produce and evaluate designs changes over time as a result of the interactions of the designer with other members of the world around it. The other members that can influence a designer's design decisions can be classified into those that are competitors (other designers in the same industry or domain, producing the same kinds of designs) and those that are consumers (of the type of artifact produced by the designer). The influence is not a result of direct communication between them, but rather results from each member being able to analyze the behaviors of the others around it, in particular their responses to the different designs being produced, and adjusting its own knowledge over time as a result. In broad terms, we can consider the computational model simulates that consumers' purchasing behavior in the world is independent of, but indirectly affects, the evaluation criteria used in producing new products.

This view of designing as including a social phenomenon is influenced by research in the branch of cognitive science known as situated cognition [10, 11]. One of the observations of situated cognition is that reasoning occurs within a world and is influenced by a designer's current worldview, called a “situation” [12]. The same designer confronted by the same requirements at a different time, or different designers confronted by the same requirements at the same time, might make different decisions while reasoning and therefore come up with different solutions to the requirements. Basing this computational simulation on ideas from situated cognition allows for the explanation of, and experimentation with, many of the phenomena involving social influences that are related to design activity.

The remainder of the paper is organized as follows. Section 2 briefly presents the mechanics of the computational simulation of a social environment in which creative agents are present, using ideas from multi-agent systems [13]. Section 3 presents some details about the makeup of the agents used in this simulation. Section 4 describes and presents the results of some experiments performed with this simulation. The paper concludes by discussing, in Section 5, some of the important outcomes of this research.

2 Multi-agent Simulation

This simulation was implemented in MASON (Multi-Agent Simulation Of Networks), a multi-agent simulation platform, developed at George Mason University [14]. In this simulated world there are 1,000 agents, of which 2.5% are designers (which are also called *producers*) and 97.5% are observers or consumers (which are also called *receivers*) of the designs produced by the designers. These proportions are based on statistics gathered by the U.S. Census Bureau [15] that show that approximately 2.5% of the U.S. population is involved in some sort of creative activity or industry.

Each designer and consumer is modeled as a single agent in MASON resulting in 25 designers agents and 975 consumer agents. Each of these agents has its own value system, modeling its situation at any time: a set of interests and preferences, or biases, that are used to evaluate designs. In addition, each of the designer agents has its own set of skills: generative knowledge that it uses to produce new designs. The sets of preferences and skills are different in each agent.

The "lives" of the agents are divided into time-steps, and a simulation is run for each agent for 1,000 time-steps. Within each time-step each designer agent produces a new design based on its set of generative skills and its evaluation criteria for deciding what makes a good design. The consumers then observe the produced designs and use their own evaluation criteria to assign a value to the quality of the designs.

Once all the consumers have had a chance to evaluate the designs produced by all the designers the results are gathered together to obtain mean values of the population of designs produced in that time-step. The mean values are used to rank the designs and the designers according to their success (the relative quality of the designs they produced, as judged by the consumers) and the consumers according to their enthusiasm (for the overall set of designs produced by the designers). The results of this procedure are used by the agents as a catalyst for potentially making adjustments to the knowledge that they use in their activities in the next time-step (evaluating designs and, in the case of designers, also producing designs).

In order to simulate the adoption of technologies and methods that have been used by others and have been proven to be successful, in a previous time-step, the least successful designers change their situation by adopting some of the knowledge (both generative and evaluative) that the most successful designer used in the time-step that has just ended (and thus try to improve their own success in the future). In the real world this adopted knowledge could have been obtained through licensing, patents, reverse engineering, industrial espionage, or other means. In order to simulate the membership behavior of consumers, where consumers are influenced to adopt products based on which products have been adopted by large groups, the least enthusiastic consumers adopt some of the evaluative knowledge that the most enthusiastic consumer used in the time-step that has just terminated in order to try to improve their enthusiasm for the overall set of designs in existence.

The above procedure is then repeated for each subsequent time-step in the simulation. Fig. 1 schematically shows the simulation framework just described. The agents in the simulation undergo gradual changes in their way of viewing the world around them (and of producing designs, in the case of the designer agents) as the simulation proceeds. These gradual changes occur as a result of each agent observing the behavior (skills, evaluation criteria, and opinions) of others, rather than as a result of direct communication between the agents. As a result of these gradual changes, our hypothesis is that interesting global (social) behaviors that were not programmed directly into the simulation emerge on the basis of the elementary social individual agent behaviors.

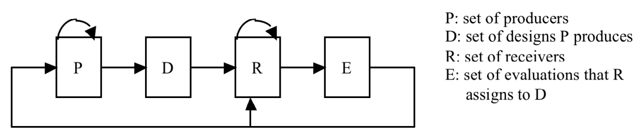

Fig. 1. Framework for the simulation of a society of producer and receiver agents.

3 Individual Agent Models

In this simulation the designer agents produce simple shapes consisting of sets of colored unit squares through an evolutionary algorithm. Any generative approach can be substituted for the evolutionary algorithm. Each agent uses several criteria in parallel to evaluate designs. In the case of the designer agents, these criteria are used to evaluate the designs that they themselves generate, and guide their generation towards convergence in each time-step. In the case of the consumer agents, the criteria are used to evaluate the designs that the designer agents produced during that time-step.

The sets of evaluation criteria, which are used to model the notion of "situation," available to designer and consumer agents overlap but are distinct. The initial state of the agents is randomly set (choosing for each agent a fixed number of criteria from the set of possible criteria that corresponds to it) before commencing the simulation. In this example the evaluation criteria relate to geometric properties of the designed shapes (such as their tallness, flatness, area-to-perimeter ratio, bumpiness, degree of convexity, and symmetry) as well as criteria that relate to color properties of the shapes (such as degree of color saturation, contiguousness of the colors, and the existence of different color patterns within the unit squares that make up the shapes).

Each of the designer agents uses a set of genes in order to create genotypes that describe moves that can be made to describe a shape (design). The set of genes that each designer agent uses is initialized at random at the beginning of the simulation, and is chosen from a set of 32 possible genes.

Each gene represents making a unit move from a given start position in one of eight possible directions (during the creation of a shape) and placing a unit square (of a particular color) in the position resulting from that move. A genotype is a sequence of such moves and placements of colored unit squares, read from left to right, that together creates an entire shape. The start position for each gene in the sequence (genotype) is the end position for the previous gene. Fig. 2 shows a subset of the set of genes available to designers (the subset shows the eight possible genes that can exist for a given color of unit square).

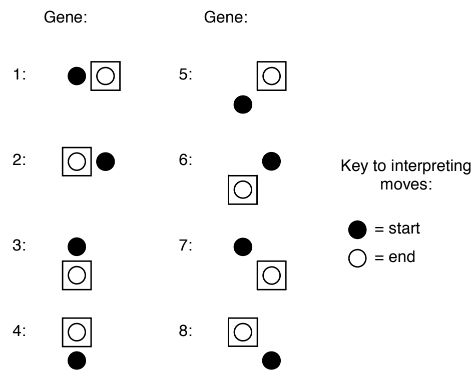

Fig. 2. Subset (for a given color) of the genes available to designer agents.

4 Experimental Results

Fig. 3 shows snapshots of the state of the simulated world after each of the first four time-steps in a typical simulation. In the snapshots, designers are shown as hollow squares distributed in five rows of five columns each, and consumers are shown as small circles that are rendered in the vicinity of the designer whose design they liked the most upon terminating the corresponding time-step of the simulation. From one time-step to the next the designers remain immobile, but the consumers travel within the window from their location in the previous time-step to the vicinity of the designer whose designs they evaluated most highly in the current time-step. The density of the cloud of consumers depicted in the vicinity of each designer is a measure of how popular/successful that designer's design was in the current time-step.

A wide range of responses can be observed in the sequence of snapshots shown in Fig. 3. If the designers are numbered from left to right and from top to bottom, Designer 8 (second row, third column) maintains an above-average level of popularity throughout the four time-steps. Designer 25 (last row, last column) has an above-average number of "followers" only in the third of the four time-steps shown in the sequence of snapshots. Designer 5 (first row, last column) is not successful at all at the beginning of the simulation, then has an average number of followers during the next two time-steps, and then has very few by the fourth time-step. Designer 21 (fifth row, first column) oscillates between being relatively unpopular and being relatively popular in each of the four time-steps.

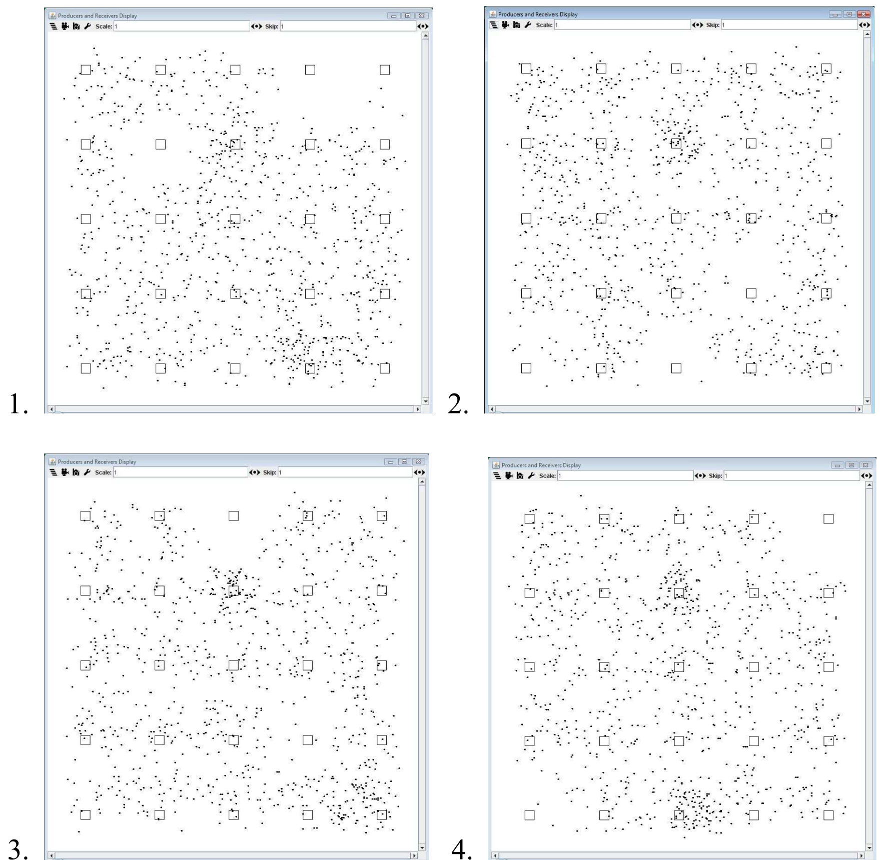

Fig. 3. Snapshots of the first four time-steps in a typical simulation. Designers are shown as hollow squares distributed in five rows of five columns each, and consumers are shown as small circles that are rendered in the vicinity of the designer whose design they liked.

This first experiment shows that many different types of social responses to creative agents can emerge in this computational simulation. This is despite the simplicity and indirectness of the knowledge transfer mechanism employed by the individual agents in each time-step in the simulation (which is what originates this range of behaviors) and despite the fact that only four time-steps were observed in detail in order to analyze the agents' individual behaviors.

The second set of experiments is designed to determine global emergent trends based on these simplified concepts of situated cognition. A Monte Carlo simulation [16] (with 1,000 time-steps, 20 runs) was run. The mean and standard deviation of the distributions of the evaluation knowledge used by the sets of agents and the genes used by the designer agents at each time-step in each run were measured. The distributions of knowledge allow us to observe whether some evaluation criteria in both types of agent and some production knowledge in the designers tend to dominate in time (their initial distribution is set at random, and is therefore statistically uniform). The standard deviation of these distributions was used as a measure of the variability within the population of agents. The mean and the standard deviation of these standard deviations were measured to obtain a global measure of the variability within the population across all runs (i.e., the variability of the variabilities).

Without the concepts of situated cognition, the simulation should tend over time to produce agents that will use the same knowledge that only a few of them used at the beginning of the simulation (the ones initially that turned out to be the most successful or enthusiastic). Situated cognition encompasses changes in the worldview of the participants. In the design world we are simulating this could be brought about when new knowledge (methods, technologies, ideas, etc.) appears. This new knowledge may augment or supplant some of the knowledge that was being used earlier. To account for this change in worldviews new values are introduced in an "onionskin" model of an open world.

In the onionskin model the open world is modeled as a sequence of closed worlds, one embedded in the other. Each "skin" completely envelopes the previous world, thus the previous world becomes an open world embedded within the next world as the constants that define that world are turned into variables by the next closed world. In this work we treat both the criteria and the generation knowledge as fixed in each time-step. This makes each time-step operate within a closed world defined by those criteria and the generation knowledge. In the next skin the criteria that were previously fixed become part of a larger set as does the generation knowledge. In this way the current time-step becomes an open world for the previous time-step. Here a set of new values is regularly introduced to account for changes that emerge from the current state of the world. These new values are added to or substituted for existing values. This is repeated at regular intervals. At every 200 time-steps new evaluation criteria are introduced for both designers and consumers and new genes are added to the pool from which designs are produced by designers. Fig. 4 shows the graph of the resulting variability (for the distribution of both designer and consumer evaluation criteria). The eleven values shown in the horizontal (time) axis of Fig. 4 and Fig. 5 correspond to eleven key time-steps in the simulation: the initial (time-step 0) and final (time-step 1,000) state of the simulation, and just before and just after the introduction of new knowledge in time-steps 199, 200, 399, 400, ...

The effect of this introduction of new knowledge can be seen in Fig. 4, which shows that the variability of the evaluation knowledge is maintained and does not converge. Having the agents react to these changes in the world by altering the way they do things crudely models the way they construct situations (interpretations of the world around them) for themselves, and thus change, as they interact with other agents in that world in the course of performing their activities [10, 11].

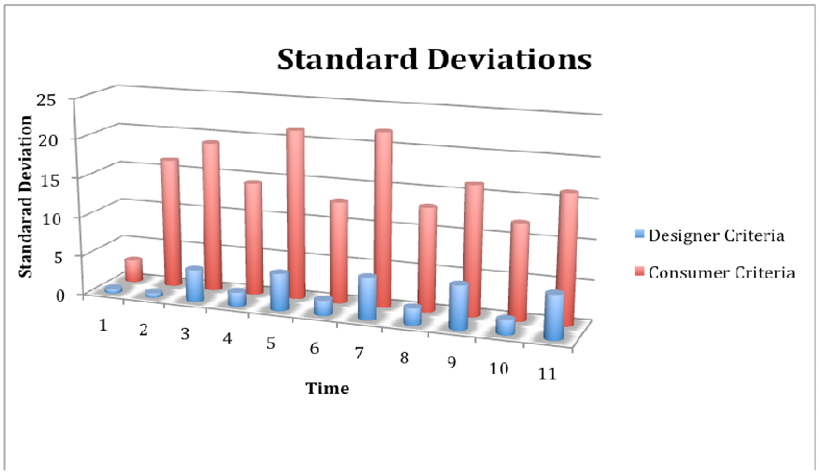

Fig. 4. Graph of the variability in terms of the standard deviations of the standard deviations of the evaluation knowledge, expressed as criteria, used by the designers and consumers.

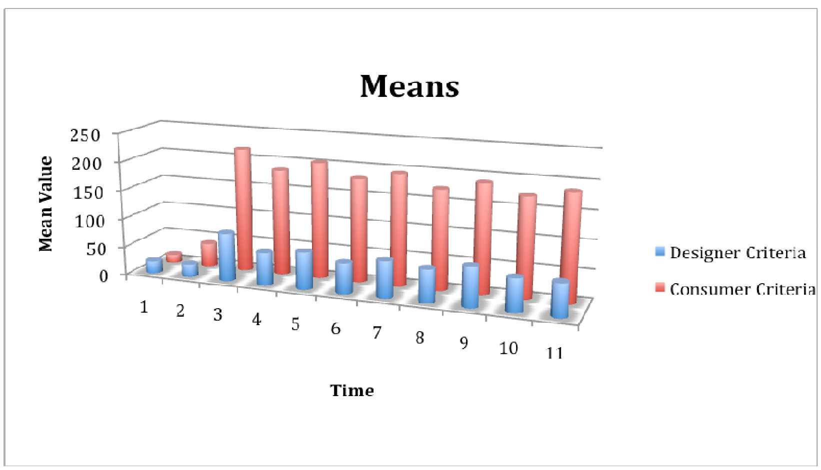

Fig. 5. Graph of the variability in terms of the means of the standard deviations of the designer and consumer criteria. The means of the designer criteria have been multiplied by 10 to make them viewable at the same scale as the consumer criteria.

Another measure of the variability is the means of the standard deviations. If these drop that is an indication of a drop in variability. If, however, they stay high then variability is sustained. Fig. 5 shows the means of the standard deviation values of the designer and consumer criteria.

Both graphs show that the variability is maintained throughout the entire process.

5 Discussion

In this paper we have presented a computational simulation that uses ideas from situated cognition to model some of the social aspects of creative activity. In our simulation, creativity does not stop as soon as an agent finishes producing some design for a particular set of requirements, as in many traditional computational models. Instead, we view creativity as an ongoing process that is influenced by factors that are external to creative agents. Our model fits well within, and contributes a computational implementation of, the DIFI (Domain-Individual-Field-Interaction) framework proposed in [17], which views creativity as a property of the interaction between individuals in a society (field) that belong to a given culture (domain).

Another model that is conceptually similar to the one we present here is described in [18]. The model in [18] uses a direct interaction between the agents, unlike the one we describe in this paper, but shares our interest in observing the emergence of complex social behavior from the elementary interactions of individual behaviors. It does so by having agents' knowledge not be static, and their "lives" not end as soon as they produce satisfactory designs, but rather modify agents' knowledge based on their changing situation as they proceed with their activities and interact with other agents, and keep agents active throughout many problem-solving episodes. A preliminary version of the model described in this paper appeared in [19].

There are no causal models of the relationship between consumer preferences and the designers of the consumed designs. However, computational simulations like this permit the testing of hypotheses and the observation of the resulting systemic behaviors. The focus of this paper has been on the hypothesis that peer pressure and market pressure are drivers of change in the way designers design creatively, and that this occurs through the indirect observations that designers make of the opinions that consumers and other designers have on their previous designs, rather than direct communication between them. The paper described and showed the results of experiments in which social behaviors emerged from this kind of indirect interaction.

Computational social science, from which this work is derived, provides the techniques to experiment with behavior in silico, behavior that is too difficult to experiment with in vivo. Complex social behavior can result from simple individual behaviors. The results produced here demonstrate that the hypothesis that creativity is both an individual and social phenomenon can be tested. The results indicate that social interactions play a role in designers being continuously creative and that the concepts of situated cognition play a role in our understanding of creativity.

Further experiments will be carried out where different ideas about how designs and design criteria substitute for existing ones in order to model Schumpeter's [20] foundational concept of "creative destruction" will be tested.

Acknowledgements

This research supported by grants form the US National Science Foundation, grant numbers CNS- 0745390 and SBE-0915482.

References

1. Jupp, J., Gero, J.S.: Towards a Computational Analysis of Style in Architectural Design. Journal of the American Society for Information Science 57:11, 1537-1550 (2006)
2. Pérez y Pérez, R., Sharples, M.: MEXICA: A Computer Model of a Cognitive Account of Creative Writing. Experimental and Theoretical Artificial Intelligence 13:2, 119-139 (2001)
3. Zitzler, E., Teich, J., Bhattacharyya, S.S.: Evolutionary Algorithms for the Synthesis of Embedded Software. IEEE Transactions on Very Large Scale Integration (VLSI) Systems: 8:4, 452-456 (2000)
4. Gómez de Silva Garza, A., Zamora Lores, A.: An Evolutionary Process Model for Design Style Imitation. In: Gero, J.S., Goel, A.K. (eds.) Design Computing and Cognition '08. Springer, Heidelberg, Germany (2008)
5. Mitchell, M.: An Introduction to Genetic Algorithms. MIT Press, Cambridge, MA, USA (1998)
6. Leake, D.V.: Case-Based Reasoning: Experiences, Lessons, & Future Directions. AAAI Press, Cambridge, MA, USA (1996)
7. Dechter, R.: Constraint Processing. Morgan Kaufmann, San Francisco, CA, USA (2003)
8. Stiny, G.: Introduction to Shape and Shape Grammars. Environment and Planning B: 7:3, 343-351 (1980)
9. Boden, M.A.: Dimensions of Creativity. MIT Press, Cambridge, MA, USA (1996)
10. Clancey, W.J.: Situated Cognition: on Human Knowledge and Computer Representations. Cambridge University Press, Cambridge, UK (1997)
11. Schön, D.: The Reflective Practitioner: How Professionals Think in Action. Basic Books, New York, NY, USA (1983)
12. Gero, J.S.: Situated design computing: Principles, in BHV Topping, (ed), Civil Engineering Computations: Tools and Techniques, Saxe-Coburg Publications, Stirlingshire, UK, pp. 25-35. (2007)
13. Epstein, J.M.: Generative Social Science: Studies in Agent-Based Computational Modeling. Princeton University Press, Princeton, NJ, USA (2007)
14. Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., Balan, G.: MASON: A Multi-agent Simulation Environment. Simulation 81:7, 517-527 (2005)
15. United States Census Bureau: Census 2000 Indexes of Industry and Occupation. United States Census Bureau, Washington, DC, USA (2000)
16. Kalos, M.H., Whitlock, P.A.: Monte Carlo Methods. Wiley-VCH, Hoboken, NJ, USA (2008)
17. Feldman, D.H., Csikszentmihalyi, M., Gardner, H.: Changing the World: A Framework for the Study of Creativity. Praeger, Westport, CT, USA (1994)
18. Sosa, R., Gero, J.S.: Social Structures that Promote Change in a Complex World. Futures 40:6, 577-585 (2008)
19. Gómez de Silva Garza, A., Gero, J.S.: Exploring Knowledge Transfer in a Society of Designers. In: 17th International Conference on Engineering Design (ICED), Vol. 8, pp. 99-110. Design Society (2009)
20. Schumpeter, J.A.: Capitalism, Socialism and Democracy. Harper, New York, NY, USA (1942)