WHAT'S COOL? - COLLECTIVE FASHION-LIKE BEHAVIOR EMERGES FROM NEURO-PSYCHOLOGICAL CONDITIONING

Jorge Simão[1] and Luís Moniz Pereira[2]
Centro de Inteligência Artificial – CENTRIA
Faculdade de Ciências e Tecnologia – Universidade Nova de Lisboa
2829 - 516 Caparica, Portugal
E-mail:[1] **jsimao@di.fct.unl.pt**; [2] **lmp@di.fct.unl.pt**

September 7, 2004

1 KEYWORDS

Fashion, conditioning, Neuro-psychology, Human social behavior, preference dynamics, culture change, agent-based simulation

2 ABSTRACT

We present an agent-based model of the emergence of fashion-like collective behavior, based on a simple abstraction of individuals neuro-psychological conditioning. Our results show that fashion-like collective behavior can emerge from social interaction and the working of individuals' neuro-psychological mechanisms, within a wide range of plausible assumptions about the levels of social stratification within a population and cognitive inertia.

3 INTRODUCTION

What do miniskirts, afro haircuts, and body tattoos have in common? They are all forms of body accessories that have had a characteristic fashion-like career. They emerge out of obscurity and spread through a population very fast, only to, shortly after they have reached their maximum popularity, vanish again from the cultural landscape, sometimes to surge again long after. For example, in the USA turtleneck shirts where first popularized in the 1920s, enjoyed a resurgence in the 1960s, and became popular again in the 1990s [19].

Although contemporary history of popular culture is riddled which such fashion-like phenomena, current scientific proposals to explain such forms of collective behavior are in many ways unsatisfactory or limited in scope. One classical explanation was proposed by Georg Simmel's [27]. According to Simmel, humans are endowed with two opposing drives or instincts: on the one hand they tend to *imitate* those individuals they admire, and on the other they tend to *distinguish* themselves from those to whom they are indifferent or who they despise [5]. From this it follows that, within a stratified society, characteristic traits of upper social classes should spread down to lower classes by imitation. As this occurs, inevitably individuals in the upper-classes will want to distinguish themselves from those in lower ones. This makes them drop previously desired traits and replace them by others.

While Simmel's effect is logically sufficient to make fashion-like collective behavior emerge (see [22, 23], for a computer simulation), its assumptions present some problems. First, postulating the existence of *instincts* raises an important ontological question [4]: What do we mean, exactly, by instinct? Can't it be used to label any behavioral regularity or pattern? If so, we should then further identify the neuro-psychological mechanisms that might generate the behavior in question. For example, while one can rationalize about the adaptive advantages of imitating somebody else [8], the advantages involved in distinguishing oneself are much less clear. Second, how do the imitation and distinguishing behavior relate to the set of values and norms shared by a population? Could the upper-classes embrace any characteristic trait, or are these individuals also somewhat constrained by what is acceptable? While, as observers, one can retrospectively find some cultural traits strange or even bizarre (given our own cultural background), it seems that even upper class individuals have been constrained in their choices by their peers and members of their social network (e.g. by virtue of negative social reinforcement that would discourage too extravagant a behavior [2]). Moreover, lower class individuals would certainty not imitate everything the upper-classes might envision. Clearly, some cultural traits are more likely to be found more often than others. That is, the existing values and norms of individuals in the lower classes also constrain what they imitate and what spreads through the population. Finally, how do individuals learn what cultural specific traits are characteristic, and how can they identify others as being members of a higher class? Without having an appropriate theory of the way in which preferences and meaning are associated with symbols and traits, one ends up with a circular explanation for the emergence of fashion.

Economists too have hurdled with the problem of explaining the strange career of products that have their value change by fashion-like behavior [6, 3, 7]. According to standard micro-economic theory, the demand for a product should mostly depend on the supply and an individual's own evaluation of the product [9]. In general, however, such theoretical framework is insufficient to explain the rapid spread of product consumption and the rapid change in individual preference characteristic of fashions. Thus, to explain this kind of phenomena, additional factors (often designated as *externalities* [26]) are required. One such proposal is that the value of some product is proportional to the total number of its users (e.g. as in telephone or other communication systems). A high number generates a positive feedback loop that leads to rapid spread of products once an initial slow inception phase is surpassed [16]. A related proposal, employed to explain the introduction and spread of innovation in markets, is to assume that the likelihood individuals will buy a product is dependent not only on its price (and therefore, its intrinsic value), but also on the number of individuals having bought it already [12]. These explanations are however limited, because they do not explain why many products suddenly lose their value. As put by Margo Bergman: "[. . .] Most products are still useful to people even if they are the only ones who own them. [. . .], miniskirts are still clothing, and therefore useful, even if they aren't popular". To fix this, this author has proposed a model where in addition to considering imitation effects in product acquisition, she considers additionally that an individual's interest in a valued product decays over time [6]. While this generates a fashion-like characteristic behavior, it is no more than an ad-hoc fix in the basic model assumptions, as it does not explain why individuals' valuations decay in the first place, neither why they can then increase at a later time.

An alternative type of hypothesis, proposed by some economists, is based on the assumption that individuals often imitate others blindly, possibly ignoring information available to them. This can create *information cascades* [3, 7], that lead to the spreading of bad products if the initial choice is poor or ill-advised. Thus, when the product's "real" value is later discovered, and the attending information spreads, this creates a sudden drop in product use. The latter kind of explanation presents two shortcomings though. First, similarly to the previous explanations, it does not apply to cases where the product is still functional and its "real" value is high. Most importantly, is assumes that there is such thing as a "real" value that can be independently assessed by the isolated individual (at least given enough information, resources, and time), without really explaining where this value comes from.

The above discussion suggests that one possible avenue to explain the emergence of fashion-like collective behavior is to identify a plausible (neuro-psychological) mechanism that would explain how individuals attribute value to traits or products, and how this value might change over time. With this mechanism one could then model a population of agents with some such psychology, and examine what kind of macro patterns it would generate. Fortunately, we know of (at least) one such mechanism: classical (or Pavlovian) conditioning [18, **?**].

In classical conditioning, previously neutral stimuli (here "traits") are made to trigger emotional or behavioral responses, by recurrent association in space and/or time between them and other stimuli that, taken alone, already produced such responses [18, **?**]. A neutral stimulus is said to have become *conditioned* by other (unconditioned) stimuli. Although this kind of neuro-psychological process is often setup in controlled laboratory environments for experimental scientific studies, one can find evidence that the same or similar processes are also operating in more natural settings. For example, studies have shown that when physically attractive people are rated in other attributes, such as linguistic or speech skills, they receive from observers higher scores than individuals with below average attractiveness [25, 17]. And an all too well known example of the "application" of conditioning mechanisms is the marketing strategy of associating value charged stimuli (e.g. sex or status symbols) with the products or brands producers want to sell [24]. In fact, it is reasonably well understood that learning to trigger similar responses when presented with similar or spatiotemporal correlated stimuli is a basic organizational principle of cognitive systems — both in artificial AI systems [13], and in natural brains [1].

In this paper, we present an agent-based model where a population of agents observe others' traits, and change their preferences for those traits by a rule akin to that of classic conditioning. This permits us to study the macro-level dynamics of preferences change and trait usage in the population. Our results show that fashion-like collective behavior can emerge out of this individual neuro-psychological mechanism, within a wide range of plausible assumptions about the level of social stratification within a population. While agents themselves change their preferences in a very natural and reasonable way — making the plausible assumption that conditioning is in general of adaptive value [21, 29] — at the population level the pattern of change is found to be quite capricious and seemingly irrational.

The rest of this paper is organized as follows. First we describe our model design. Next we present the model's results. We then broaden the discussion to relate our results with those of previous models, and examine how our results and approach can be utilized in the wider context of cultural preference change. Finally, in the last section we set forth our conclusions.

4 A MODEL OF FASHION EMERGENCE BY CONDITIONING

We assume a population of P agents with two attributes and two preferences. One attribute is real-valued and remains fixed during the entire simulation runs. It represents some quality or cluster of qualities of the agent that is positively valued by all members of the population. For example, physical attractiveness, speech ability, or some aggregated (holistic) perception of several of such traits[1]. We dub this attribute the *quality attribute* or *quality* for short. Its value q_i is assumed to be randomly generated from a normal distribution with mean 0 and standard deviation σ^2. A second attribute t_i is a binary trait, and represents some agent trait or item that the agent either carries with him (when $t_i = 1$), or not (when $t_i = 0$). For example, t_i may represent a type or brand of clothes, or a body ornament. This second attribute can be changed by the agent at will, according to what is perceived by the agent as being the most valued option (from its subjective aesthetical point of view). To distinguish it from the quality attribute, we will call it the *trait attribute* or *trait* for short. Two additional *preference* attributes represent the agent's subjective perception of what is the value of not having the trait, v_i^0, or of having the trait, v_i^1. The values of v_i^0 and v_i^1 are updated according to a process of conditioning (or association) between observed traits and qualities, described next. Overall, an agent a_i can be represented formally as a four element vector: $< q_i, t_i, v_i^0, v_i^1 >$.

At each time step every agent observes a set of N different agent role models. This is the set of agents which influence an agent's current perception of the value of having or not the *trait*. As a very simple abstraction of the process of conditioning, we assume that the value v^1 of having the trait is changed by an amount proportional to the average of the qualities of the role models that have the trait. Additionally, since we want to explore the effects of memory within agents, we let an agent's new valuation to be influenced by its previous valuation. For this purpose, we define a learning rate parameter α, that specifies how insensitive agents are to new observations. That is, it specifies how slow or how fast agents are in forgetting previous valuations and changing to new ones. Specifically:

$$v_i^1(t) = v_i^1(t-1) \cdot \alpha + \frac{1}{N} \sum_{a_j : a_j \in M_i \wedge t_j = 1} q_j \cdot (1 - \alpha)$$

[1] On the specific topic of mating preference this as been abstracted by Donald Symons as *mate value* or *mate quality* [28].

In the equation above, M_i is the set of models for agent a_i ($\#(M_i) = N$), $\alpha \in [0, 1]$ is the learning parameter, and $v_i^1(t)$ and $v_i^1(t-1)$, are respectively the new and previous valuations of trait usage for agent a_i. In a dual manner, we define the value of not having the trait to be the average of the qualities of the models that do not have the trait, weighted by the learning parameter α. Formally:

$$v_i^0(t) = v_i^0(t-1) \cdot \alpha + \frac{1}{N} \sum_{a_j : a_j \in M_i \wedge t_j = 0} q_j \cdot (1 - \alpha)$$

The set of role models for an agent is selected stochastically at each time step, such that the probability of an agent being a model to another agent is proportional to how close their qualities are. This is intended to model a scenario of a socially stratified society, where agents are more likely to interact (or observe) agents that are closer to them in quality than they are to interact with average quality individuals. We implement this by using a roulette-wheel selection mechanism [10], and making the probability of agent q_j to be a model of q_i proportional to $max\{(3\sigma^2 - |q_i - q_j|)^{-E}, 0.001\}$, where the exponent parameter E specifies the degree of social stratification (assortment) in the population. This setting implies that, for positive values of E, individuals with extremely high and low qualities have little chance of interacting amongst them, while individuals with equal quality have maximum probability of interaction. When $E = 0$, the population is not stratified, and individuals have an equal probability of interacting with other individuals with regard to qualities. Thus, higher E values imply higher correlation coefficients between the quality of agents and their models (numerically illustrated in the next section).

At every time step each agent goes through two phases. In the first phase, the agent is assigned a model set of size N, and it updates its subjective value of having and not having the trait as described above. In the second phase, each agent a_i decides whether or not to switch its trait attribute value. If $v_i^1 > v_i^0$, the agent will start using the trait by setting $t_i = 1$ (if not set already). Conversely, if $v_i^0 > v_i^1$ the agent will set $t_i = 0$, indicating that it does not have the trait. Once an agent changes the trait usage, we assume that it takes at least D time steps until he can switch it again. This is intended to represent a cognitive or material inertia factor [14]. Although changes do not occur, the values v_i^0 and v_i^1 continue to be updated.

Additionally, to allow the model to escape stationary situations where all or none of the agents have the trait, we define a parameter θ that specifies the probability that an agent will choose a random value for t_i independent of its own valuation (e.g. due to some exogenous factor).

This two-phase procedure can be anecdotally interpreted as follows: individuals go out every day sporting the trait or otherwise (e.g. wearing a certain type of clothes). During the day they interact and/or observe other people and, based on their observations, they update their mental subjective perception of the value of possessing or not the trait (possibly sub-consciously [18]). At night, back home, they sleep on the subject. In the morning, with their new sense of aesthetics properly in place, they make a fresh new decision on whether or not to exhibit the trait during the day. Since in this version of the model updates and actions of agents are asynchronous, to keep to this metaphor, we should assume that each agent performs its updates and trait changes on different days.

Parameter	Description	Value(s)	Note
P	population size	50	small sample
N	number of models	5	small
E	assortment	4	$r \approx 0.75$
α	1 - learning rate	0.2	fast learning
σ	standard deviation	2	
D	delay	2	cognitive or material

Table 1: Parameter settings.

5 RESULTS

To study the model dynamics, we set the population size parameter $P = 50$, and the model set size $N = 5$. The standard deviation of agents' quality σ^2 is set to 2. We start by setting the assortment parameter $E = 4$, which generates a correlation coefficient r between the quality of agents and their role models to be approximately 0.75. We set learning rate parameter $\alpha = 0.2$, and the random change parameter $\theta = 0.04$. We start all simulation runs with no agent having the trait ($t_i = 0$). And we make the initial value that agents attribute to having or not the trait equal to $-3{*}\sigma$. This corresponds to a situation where trait usage is initially neutral for all agents, and therefore it's indifferent for them whether or not they possess it. We make each simulation run from $T = 100$ to $T = 500$ time steps. In Table 1, we present a summary of the model's parameters and the base values we assign them here.

In Figure 1, we depict a bit map (Wolfram diagram) representing trait usage across time. In the y-axis agents are sorted by quality, with higher quality agents towards the top of the figure and lower quality ones towards the bottom. Time runs left to right. As can be seen in the figure, with the parameter values set above the simulation produces a wave-like pattern. This pattern can be more clearly demonstrated by plotting in Figure 2 the frequency of trait usage across time. This wave-like pattern emerges because individuals tend to switch to the trait value of higher quality individuals. This feature is responsible for the diffusion of a trait value, but also makes the opposite trait value more preferable. However, due to the inertia of agents in switching traits, as specified by parameter D, it takes some time until agents are allowed to switch back.

Contrary to our initial intuition, the diffusion of the trait tends to occur very rapidly. Just a few time steps are required for a large wave of a particular trait value to emerge. We found out that this is caused largely due to the low (but realistic) degree of assortment. To test under what scenarios diffusion of a trait value would be slower, we made the model selection deterministic such that agents only observe the immediately higher and lower quality individuals. In Figure 3, we present a typical simulation result for this case. As can be gathered from the figure, chance mutations of trait values in high quality individuals produce more ordered spreading of trait usage. However, in this case the correlation between agents and models raises to a value near 1. This contradicts empirical evidence about modern societies. For example, when considering assortment among individuals of the opposite sex in mating relationships, correlation values between .3 and .8 have been found in attributes such as IQ, physical attractive-

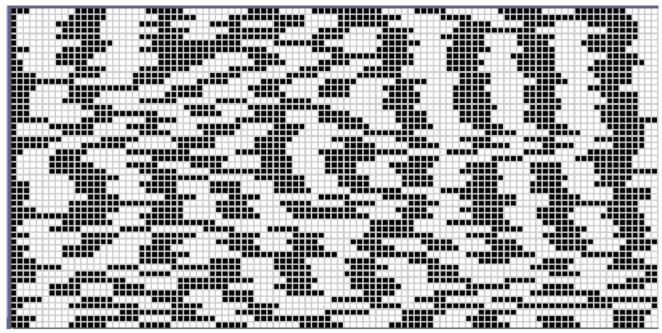

Figure 1: Bit map of trait usage across time ($D = 4$).

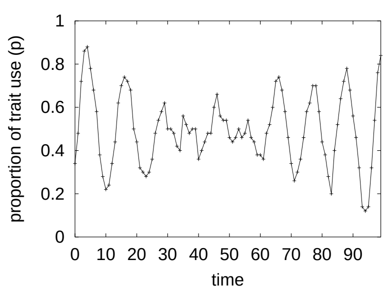

Figure 2: Frequency of trait usage across time ($D = 4$).

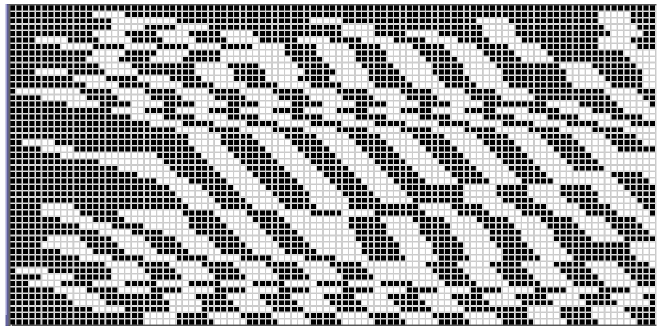

Figure 3: Bit map of trait usage across time ($D = 4$) with deterministic selection of model.

ness, and others [15, 20]. Similarly, it has been found in a series of empirical studies that the same pattern of assortment does occur amongst friends of the same-sex [11]. This result called our attention because some researchers have proposed to use spatial metaphors to represent social status in deterministic small neighborhoods [22, 23]. By comparing Figure 3 with Figure 1, it becomes apparent that the way social assortment is modelled can be relevant in obtaining simulation results and deriving theoretical conclusions.

It is useful to compare the model's behavior with the fixed parameter setting, with standard results from gene-culture co-evolutionary theory. According to R. Boyd and P. Richerson's model of two (discrete) traits' cultural transmission, the trait frequency of those traits with no bias should converge at equilibrium to the ratio of the mutation rates of the two trait values [8]. Since in our model the probability of random introduction of any of the trait variants is equal, the frequency of trait usage should converge to 50%. This does not happen in our model because in it we also introduce time varying bias for the preference of trait usage or avoidance. In another model, R. Boyd and P. Richerson showed that if the bias is fixed, the population should converge to the trait with the highest bias [8]. This also does not occur in our model because the bias changes with time. It gets higher whenever medium or low quality individuals do not have the trait of high quality individuals. This drives the model behavior into a continuum oscillation between near full trait usage to near full trait avoidance. This can be verified in Figure 2.

To test the robustness of this result we start by varying the value of the delay parameter D. In Figure 4, we present a bit map of trait usage across time with D set to 10, and in Figure 5 with D set to 0. As can be seen upon inspecting Figure 4, with large value of D the behavior of the model is qualitatively very similar for most values — it produces wave-like patterns across all ranges of agents' quality. On the other hand, in Figure 4 we can see that when $D = 0$ (no enforced delay between trait switches) the patterns are not so clear and appear almost random. This is so because agents respond faster to devaluations of trait usage. We also systematically varied the learning rate parameter, and

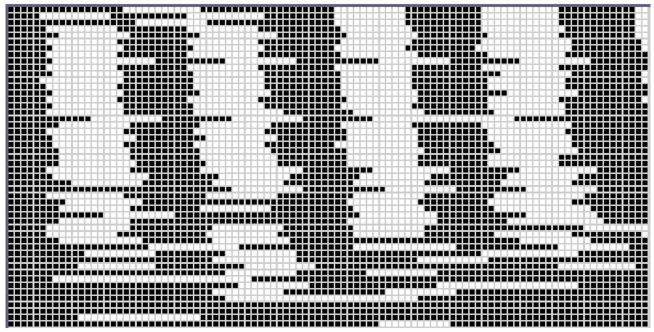

Figure 4: Bit map of trait usage across time ($D = 10$).

the number of models, and we found no important changes in model behavior. Thus, a tentative conclusion that we take from our simulation results is that cognitive inertia or material resources is an important factor in establishing the duration of a fashion-like phenomena. Although, in a real setting many more factors will be involved.

6 DISCUSSION AND FUTURE WORK

Our model shows how a constant oscillation between a near full trait usage and near full avoidance can occur. It postulates only very simple psychological mechanisms of valuation or conditioning, a (normal) quality distribution of agents, some criteria for the selection of model agents, and most importantly a cognitive delay or material cost between two successive switches. The act of switching traits can be interpreted as imitation, but to the extent that one assumes the switching of traits corresponds to the invocation of pre-existing behavior routines, then postulating imitation as a psychological mechanism is not absolutely required [**?**]. In any case, we do not need to postulate the existence of two drives, as in Simmel's effect, one for imitation and another for wanting to be different.

Thus, we put forward a new mechanism based on conditioning to explain fashion-like collective behavior. This is an alternative to other standard explanations such as the Simmel effect, or postulating irrational agents, as discussed in the introduction section. Still, much work needs to be done. In our simple model the imposed delay was the most important factor in establishing the period of fashion waves. It would be interesting to see under what conditions the period of the waves can be decoupled from direct influence of model parameters and instead emerge from the model dynamics.

An important addition to the model is to consider agents with multiple changeable traits, and traits with multiple values. This would most probably change the model dynamics for each individual trait. (See [**?**] for a similar model design.) Another aspect to consider would be to add differential abilities

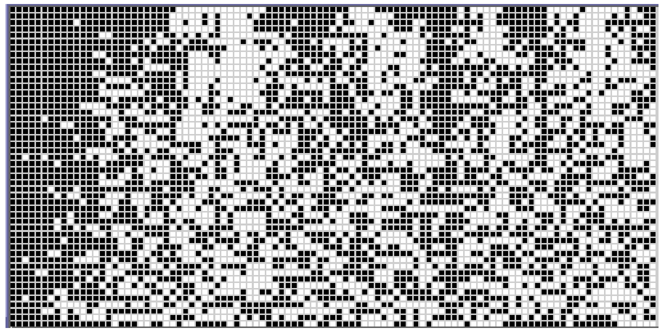

Figure 5: Bit map of trait usage across time ($D = 0$).

to switch traits based on an agent's quality, representing some form of social class empowering.

7 CONCLUSIONS

In this article we have presented a simple model that shows how mechanisms of neuro-psychological conditioning at the individual level can generate the emergence of fashion-like collective behavior. The model shows that for a wide range of parameter values, trait usage can oscillate continuously between periods of near full trait use to near total trait avoidance.

8 ACKNOWLEDGMENTS

This work was partially supported by a PRAXIS XXI, FCT/FSE, Ph.D. scholarship, and project FLUX, funded by FCT/MCES, Portugal. We would like to acknowledge the scientific and intellectual support and/or comments to earlier draft versions of this paper to Paulo Gama Mota, Nuno Preguiça, and João Sousa.

References

[1] Michael A. Arbib, Péter Érdi, and János Szentágothai. *Neural Organization: Structure, Function, and Dynamics.* Cambridge, MA: MIT Press/Bradford Books, 1998.

[2] Albert Bandura. *Social Foundations of Thought and Action: A Social Cognitive Theory.* N.J.: Prentice-Hall, Inc., 1985.

[3] Abhijit V. Banerjee. Simple model of herd behavior. *The Quarterly Journal of Economics*, cvii(3):797–817, 1992.

[4] Greogory Bateson. *Steps to an Ecology of Mind.* Chicago: University of Chicago Press, 2000.

[5] Sergio Benvenuto. Fashion: Georg simmel. *Journal of Artificial Societies and Social Simulation*, 3(2), 2000. http://www.soc.surrey.ac.uk/JASSS/3/2/forum/2.html.

[6] Margo Bergman. When a fad ends: An agent-based model of imitative behavior. In *Preproceeding of the Workshop on Self-organization, Evolution and Social Behavior.* In Press, 2002.

[7] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as information cascades. *The Quarterly Journal of Economics*, 100(5):992–1025, 1992.

[8] Robert Boyd and Peter J. Richerson. *Culture and the Evolutionary Process.* Chicago: University of Chicago Press, 1985.

[9] Robert H. Frank. *Micro-Economics and Behavior (fifth edition).* New York: McGraw-Hill, 2003.

[10] David E. Goldberg. *Genetic Algorithms: in Search, Optimization & Machine Learning.* Addison Wesley, 1989.

[11] David B. Grusky, editor. *Social Stratification: Class, Race, and Gender in Sociological Perspective.* Westview Press (2nd edition), 2000.

[12] Giovanni M. Holanda, Ana L. C. Bazzan, Gustavo P. B. Gerolamo, Jo ao H. A. Franco, and Ricardo B. Martins. Modeling the bass diffusion process using an agent-based approach. In *Proceeding of the 4th Workshop on Agent-Based Simulation*, pages 147–152. SCS European Publishing House, 2003.

[13] John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery.* Cambridge, MA: MIT Press, 1986.

[14] Wander Jager. *Modelling Consumer Behavior.* Universal Press, 2000.

[15] S. Michael Kalick and Thomas E. Hamilton. The matching hypothesis reexamined. *Journal of Personality and Social Psychology*, 51(4):673–682, 1986.

[16] Michael Katz and Carl Shapiro. Network externalities, competition and compatibility. *American Economic Review*, 75(3):?–?, 1985.

[17] Ziva Kunda. *Social Cognition: Making Sense of People.* Cambridge, MA: MIT Press/Bradford Books, 1999.

[18] Joseph LeDoux. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life.* Touchstone Books, 1998.

[19] Mark A. Long. *Bad Fads.* ECW Press, 2002.

[20] Freddy B. Christiansen Marcus W. Feldman, Sarah R. Otto. Genes, culture, and inequality. In Kenneth Arrow, Samuel Bowles, and Steven Durlauf, editors, *Meritocracy and Economic Inequality*, pages ?–? Princeton University Press, 2000.

[21] Geoffrey F. Miller and Peter M. Todd. Exploring adaptive agency i: Theory and methods for simulating the evolution of learning. In D.S. Touretzky, J.L. Elman, T.J. Sejnowski, and G.E. Hinton, editors, *Proceedings of the 1990 Connectionist Models Summer School*, pages 65–80. San Mateo, CA: Morgan Kaufmann, 1990.

[22] Roberto Pedone and Rosaria Conte. The simmel effect: Imitation and avoidance in social hierarchies. In Scott Moss and Paul Davidsson, editors, *Proceeding of the 2th Workshop on Agent-Based Simulation*, pages ?–? Berlin: Springer, 2000.

[23] Roberto Pedone and Rosaria Conte. Dynamics of status symbols and social complexity. *Social Science Computer Review: Special Issue on the Simulation of Social Agents*, 19(3):?–?, 2001.

[24] Alissa Quart. *Branded: The Buying and Selling of Teenagers.* Perseus Publishing, 2003.

[25] Ted L. Rosenthal and B. Simmerman. *Social Learning and Cognition.* Academic Press, 1978.

[26] Thomas C. Schelling. *Micromotives and Macrobehavior.* New York: W.W. Norton & Company, Inc., 1978.

[27] Georg Simmel. Fashion. *American Journal of Sociology*, 62(6):541–558, 1957.

[28] Donald Symons. *The Evolution of Human Sexuality.* New York: Oxford University Press, 1979.

[29] Peter M. Todd and Geoffrey F. Miller. Exploring adaptive agency ii: Simulating the evolution of associative learning. In Jean-Arcady Meyer and Stewart W. Wilson, editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, pages 306–315. Cambridge, MA: MIT Press/Bradford Books, 1991.

9 AUTHOR BIOGRAPHY

Jorge Simão completed his graduate, masters and doctoral degree in computer science and engineering at the Faculty of Sciences and Technology, of the New University of Lisbon. His research interests include studying aspects of human (social) behavior at the neuropsychological, behavior, and societal level, and connecting the three levels within a bio-computational and evolutionary framework. E-mail: **jsimao@di.fct.unl.pt**. HomePage: http://centria.di.fct.unl.pt/~jsimao/.

Luís Moniz Pereira is full professor since 1985 at the Department of Computer Science, Faculty of Sciences and Technology, of the New University of Lisbon (FCT/UNL), the Director of its Center for Artificial Intelligence (CENTRIA), and the Coordinator of its M.Sc. in Applied AI. His research interests include Knowledge Representation and Reasoning, Logic Programming, and Cognitive Science. E-mail: **lmp@di.fct.unl.pt**. HomePage: http://centria.di.fct.unl.pt/∼lmp/.