Emergence in Social Networks: Modeling the Intentional Properties of Multi-Agent Systems

Jorge Louçã[1], John Symons[2], David Rodrigues[1], André Morais[1]

[1] ISCTE / LabMAg, Av.das Forças Armadas, 1649-026 Lisboa, Portugal
Jorge.L@iscte.pt, {david.rodrigues, anoeee}@gmail.com

[2] University of Texas at El Paso, 500 W.University Ave., El Paso, Texas 79968, USA
jsymons@utep.edu

Abstract. This paper examines some methodological and conceptual challenges concerning the study of emergent patterns of behavior in communication networks. We discuss the notion of intentionality in collectives in the light of two case studies. The first one regards the mechanisms of communication in a community of insects. The second case study concerns the "smart mobs" phenomena in human societies. The focus of our discussion is the detection of what we call *emergent fingerprints*, which signal the presence of a system of agents behaving in a manner identical to a social structure whose intentional behavior under some known set of constraints we have modeled. These structures are used to attribute some intentionality to social systems, even in the presence of noisy data sets.

Keywords: agent-based model, pattern-oriented modelling, social simulation, communication patterns.

1 Introduction: Intentionality in Collectives

Social networking technologies have given rise to new social dynamics and distinctive kinds of collective behavior. This paper describes how collective behaviors can be modeled and examines some of the associated methodological and conceptual challenges.

We assume an anti-reductionist conception of collective behavior, recognizing, for instance, the effect of the social context on the behavior of its constituents. For instance, as Bibb Latané (8) argues, when people engage in group activity, their sense of individual responsibility attenuates and their behavior changes accordingly. Cases where we observe a downward causal effect of the social context on its constituents support a distinction between collective and individual intentionality (14). We therefore assume that socially networked behavior is not merely the composition of constituent individuals.

Understanding the emergent features of group behavior has obvious practical relevance. Take for instance Howard Rheingold's (10) notion of a smart mob. Rheingold introduced the expression "smart mob" to describe the concept of a "mobile ad hoc social network". Smart mobs are social networks where people communicate using mobile and wireless internet technologies. As a first introduction to the intentional character of smart mobs consider the relationship between a network of friends and some scarce and desirable resource. One important reason to have friends is that it allows one to share in a larger store of relevant information with respect to valuable resources. Technologically extended social networks amplify and alter the effect of familiar friendship relations in dramatic ways. For example, the press has reported several cases of smart mobs aiming to exploit scarce resources. The June 29th 2006 edition of *The Economist* described the example of Guangzhou, where 500 shoppers assembled in front of a popular electronics superstore. They arrived *en masse*, at a given moment previously agreed to online and by cellphone messages. Shoppers left the store having secured 10-30% discounts on cameras, DVD players and flat-screen televisions. Smart mobs of this kind play are an obvious and easily explained example of collective action in pursuit of bargaining power.

Smart mobs are becoming increasingly familiar for their role in social and political expression. Street demonstrators at the 1999 anti-WTO protests in Seattle used dynamically updated websites, cellphones, and swarming tactics (13). Demonstrators using SMS were able to quickly converge on a single place from highly dispersed positions. They overwhelmed the local police presence with brief but intense protests, quickly dispersing and blending back into the crowds before the police could reallocate their forces. More recently, SMS communication was used to organize mass protests all over Spain in the aftermath of the Madrid train bombings of March 11th 2004 (6). Again, viral communication strongly spread through social networks mainly composed of friends, where trust between members of the network is extremely high.

The increasing importance of networked forms of political, cultural and economic expression were noted by Tony Negri and Michael Hardt in their study of globalization *Empire* (7). While Negri and Hardt's work is highly impressionistic and imprecise, it correctly points to the possibility that networked activity (what they call 'the multitude') can give rise to radically new political and economic changes. Given the increasingly pervasive role of networking technologies, it is highly probable that collective activity like smart mobs will play an increasingly important role.

In our analysis, we will point towards the possibility that collective behavior exhibits properties which are not as evident as those exhibited by the shoppers in Guanghzhou or the demonstrators in Seattle. In those cases, both the emergent feature and the explanation of those features are relatively obvious. Investigating emergent properties can be more than simply attempting to explain obvious emergent features of groups. Instead, our approach opens the possibility that previously unnoticed emergent features can be discovered in complex systems and that these features can play a role in explanations at various levels of analysis.

2 Outline and Assumptions

Our principal goals involve identifying emergent patterns of behavior in a social environment. To this end, we explore the practicality of adopting what Daniel Dennett has called 'the intentional stance' towards collective behaviors (1). We test several related research hypotheses via the modeling and simulation of case studies in a variety of social contexts, arguing that modeling such systems both reveals and explains emergent patterns in communication. The focus of our discussion is the detection of what we are calling *emergent fingerprints*, which signal the presence of a system of agents behaving in a manner identical to a social structure whose behavior under some known set of constraints we have modeled. We begin with relatively simple simulated social environments. From these simulations we extract emergent properties of the behavior of some group of agents which we then treat as a template for pattern-matching within the massive and noisy data sets which are derived from the study of natural systems.

To study this hypothesis, we propose to distinguish patterns or forms that arise by virtue of constraints on the flow of information. In our modeling projects, we identify sequences and combinations of those forms. This approach will lead us to interpret and predict emergent properties of various forms of social systems, even in the presence of massive and noisy data sets.

We build our models on the assumption that knowing only the links between the nodes of the communications network is enough to characterize the flow of information in a social environment. One advantage of the assumption of this hypothesis is its independence from the semantic properties of communication between agents. While most current treatments of emergent features of group knowers involve propositional or semantic features of the systems under consideration, our strategy is to examine patterns that appear at the level of behavioral patterns. We anticipate that social structures in which we can understand the characteristics and constraints on communication will exhibit behavioral features when modeled. We predict that some of these patterns will be robust enough to detect in real world settings.

Aiming to test these hypotheses, our research derives from two main scientific domains: the well known field of multi-agent based simulation (4) and the research that has recently been labelled pattern-oriented modeling (5).

This paper presents the broad trajectory of our research, starting with the case study *Cicada barbara lusitanica*, detailed in a previous text (9). From this initial study concerning the mechanisms of communication in a community of insects, we have developed another case study, concerning the emergence of self-organized social structures in human communication.

We provide a brief characterization of the theoretical framework supporting our research before discussing its relevance to the understanding of social phenomena. Next, our initial case study from zoology is discussed as the basis for the implementation of a multi-agent simulation environment illustrating some of the general features of our approach. In the simulations, we compare the input communication patterns and resulting emergent movement patterns from different species of cicadas. The second case study concerns a multi-agent model of communication and movement, illustrating the social dynamics of the "smart mobs" phenomena. This case study allows the identifying of patterns, both at the micro and macro levels of observation. Finally we discuss open questions and directions for further research.

3 Grasping Macro Intentionality

In our work, we take a commonsense attitude towards the relationship between individuals and groups. Contrary to a rigidly reductionist approach, we assume that the constituents of social structures and the social structures themselves mutually influence and transform one another. In this tangled and complex circumstance, some features will strike observers as stable and significant. Such features of human life or of complex systems are not necessarily (or generally) amenable to reductive explanation. Such features are sometimes called emergent. Just because patterns are not easily reduced to the behavior of their constituents, does not mean that they are somehow mysterious or that they cannot be explained scientifically. They can often be understood in terms of the kinds of mechanisms that computational modeling reveals (15). In fact, our approach supports the possibility that through modeling, not only are we in a position to explain previously puzzling emergent properties, but perhaps more interestingly, our models allow us to discover previously unknown emergent patterns in complex systems of various kinds.

In ordinary human affairs, we assume that other people have beliefs and desires and that their actions can be explained in terms of those beliefs and desires. Daniel Dennett (1) called this pattern of belief-desire-action explanation folk psychology and argued that when we engage in folk psychological reasoning, we are adopting the intentional stance towards our fellow human beings (14).

For Dennett, when we take the intentional stance towards something or someone, we project the virtual world of beliefs and desires onto the other person or animal in somewhat the same way a geographer might project lines of latitude and longitude onto the Earth's surface. In both cases, the projections permit us a means of manipulating the objects in questions and in both cases the question of whether these virtual objects *really exist* is misguided. The intentional stance is a strategy that begins with the assumption that other animals believe what they should believe given their perceptions and desire what they should desire given their needs. This is what Dennett calls the assumption of optimal design. We assume that other animals (including people) tend to pursue outcomes that serve their interests and that they have been equipped, by natural selection with suitable perceptual and cognitive capacities to manipulate their environments appropriately. Ascriptions of beliefs and desire are often objectively true, he grants, but not by virtue of describing inner mechanisms, any more than references to centers of gravity, vectors, equators and other useful virtual notions.

The most distinctive aspect of Dennett's approach to the mind is his choice of starting point. Rather than diving immediately into traditional questions about the nature of mental entities, Dennett begins by investigating what we ordinarily say or write about other minds. Conversation or language, Dennett writes, is the royal road to the knowledge of other minds (2, p13). So, rather than immediately focusing on the *terra incognita* of what the mind actually is, he focuses instead on the practical role played by the terms we use to talk about the mind, terms like 'belief', 'desire', 'the will', 'consciousness' etc. Dennett's interest lies in understanding the ordinary commonsense descriptions of other minds that we find so useful in our daily lives. What are we doing, for instance, when we say that another person believes, desires or imagines something? And why does our talk of what people think, want, or hope, seem to work so well in ordinary life?

Traditionally, philosophers have been interested in understanding the nature of “belief” rather than the process of interpretation by which we come to know the beliefs of another person or animal. Generally, they began with the assumption that words like ‘belief’, ‘desire’ and the like refer to mental entities of some kind. They were not overly concerned with the ways we come to know those mental entities, since presumably, even if we encounter some difficulty when it comes to other minds, our own minds are easy enough to examine. Dennett does not assume that these mental entities exist. Instead, his approach is to remain agnostic (at least initially) towards the meaning of psychological terms like ‘belief’, ‘desire’, etc. Rather than assuming that the word ‘belief’ refers to some kind of entity in the mind or brain, he focuses instead on the role such terms play in practice. In effect, this means that he begins by considering the reasons we use words like ‘belief’ and ‘desire’ rather than the nature of belief and desire *per se*. By focusing on what we do when we interpret or describe the mental life of another person or animal, Dennett redirects our attention from traditional philosophical puzzles and opens fertile new lines of inquiry for both scientists and philosophers.

This subtle shift in perspective changes the contour of philosophical problems considerably. Instead of worrying about the mysterious inner workings of other minds, or the problem of how the mind and body interact, we look instead to the publicly observable utterances we use to describe one another. These utterances, along with the contexts in which they appear, form an objective starting point for a more scientific approach to the problems posed by mental life. In principle, it should be possible to study the relationship between these patterns of utterances and the underlying processes that take place in the brain and central nervous system.

According to Dennett, our use of psychological terms arises primarily out of our interest in predicting the behavior of people and other animals. For obvious reasons, animals need a way to anticipate the behavior of possible mates, predators, prey and competition. This is accomplished in a variety of ways in different species. Naturally, we humans are most familiar with the techniques employed by mammals like ourselves. While mammals have an array of techniques for predicting behavior, humans tend to be especially biased towards visual cues of various kinds. By watching the eyes, posture and motion of other mammals we can generally predict their behavior into the very near future with some accuracy. So, how do we manage these remarkable prophecies? Is it a matter of knowing the mind of the other person or animal?

For Dennett, Darwin’s theory of evolution through natural selection provides the key to understanding our amazing ability to predict the behavior of other animals. Over the course of natural history, those who fail to reliably anticipate the future behavior of relevant others soon perish. As a result, a significant number of species have evolved to become masters in the art of predicting the futures of massively complex biological systems. For example, without knowing anything about the physiological processes taking place under the animal’s skin or in its mind, most of us can predict, with relative certainty, what a hungry dog is likely to do when we offer it a bowl of food. Our ability to predict the behavior of other biological systems is not the result of an ability to somehow see inside the minds of other animals. Instead, according to Dennett, our interpretations are the product of a skill that has been sculpted by a long process of natural selection. Our ability to accurately predict the future behavior of other creatures based on the evidence of their past behavior is a skill that we, along with many other animals, simply inherited.

By contrast, we humans are not suited by natural selection to predict the behavior of collectives in complex, technologically-mediated social contexts. The simple fact that the new social context results from communications technologies which postdate the selective pressures on our species means that technological enhancement is needed in order for us to adopt something analogous to the intentional stance in such cases. It is also worth noting that language, Dennett's royal road, drops out of the picture in the cases that are of interest to us.

The case studies presented in the following sections exemplify two very different examples of social interaction. Nevertheless, they both illustrate that the attribution of intentionality to groups in a social network allow the explanation of social dynamics via the only detection of emergent behavioral fingerprints. Computational modeling allows us to understand how the collective behavior of the system is shaped by constraints on communication i.e., without any access to the semantic content of that information. As such, this perspective allows deep insight into the study of social behavior. In addition to offering insight which was previously unavailable to us, this approach creates many new questions and opens up many new lines of research.

4 *Cicada barbara lusitanica* case study

Our first case study rests on experimental results obtained at the Faculty of Sciences at the University of Lisbon, concerning live experiments on the stereotyped singing response behaviors of cicadas (3, 9). We used their results to design a multi-agent simulation platform which allowed us to represent the behavior of two cicada species, *Cicada barbara lusitanica* and *Cicada orni*.

4.1 Previous experimental results

Previous experimental results concerning the study of cicada communication (2) allowed us to conclude that insect songs encode specific information about the identity of species, which are used by individuals to discriminate conspecific from heterospecific sympatric species. Figures 1 and 2 depict examples of temporal calling song configurations. The *Cicada barbara* song, characterized by an uninterrupted pulse, illustrates a continuous pattern (figure 1). By contrast, the *Cicada orni* song shows a clear discontinuous pattern (figure 2).

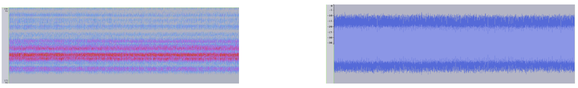

Fig. 1. *Cicada barbara* calling song analysis: spectrum and waveform db.

Figure 1 shows that the *Cicada barbara* calling song has a broad spectrum (image on the left), although some frequencies are more intense than others. The waveform analysis shows that the calling song is continuous, with no pauses.

Fig. 2. *Cicada orni* calling song analysis: spectrum and waveform db.

The *Cicada orni* spectrum depicts a clearly discontinuous calling song. This particular species of cicada uses a calling song characterized by pulses and pauses. The spectrum occupies a narrower band than that of *Cicada barbara*. Nevertheless, it is very intense in some frequencies. The waveform represents the discontinuity in the calling and regular pauses (e.g., silence) are evident.

Concerning frequency, the analysis of a pulse period in *Cicada orni* calling song reveals that very few frequencies are significant. By contrast, *Cicada barbara* might have to produce a broader spectrum of frequencies as it will not use temporal patterns (the song is continuous). Fonseca and Revez conclude that such configurations of calls are generally used by cicadas to identify members of their own species.

Fonseca and Revez stated that both frequency spectrum of the signal and temporal pattern carries information about the species-identity of a calling male, but the use of only one parameter might not be sufficient (3). The pre-copulatory isolating mechanism based on song analysis, used to maintain species integrity, uses one or/and another parameter according to the species environment.

4.2 Multi-agent simulation of cicada communication behaviors

Fonseca and Revez experiments and conclusions were used to implement a multi-agent simulation of cicada communication behaviors. The model considers the existence of two patterns used by *Cicada barbara* and *Cicada orni* to recognize conspecifics: temporal (pulse and pause duration) and frequency patterns. *Cicada barbara* and *Cicada orni* individuals are randomly placed in a finite environment. Males are static and females move, attracted by songs from their conspecific males. Each female cicada has a given initial energy, which she expends by traveling through the environment. When a female meets a male, two things can occur: either they belong to the same species and in this case her energy is set to a maximum level, or they belong to different species and then her energy is set to a minimum level. The key issue under consideration in the model is the pattern-matching mechanism used by females to recognize conspecific song. Several simulation experiments were modeled and executed using temporal patterns, frequency patterns separately and simultaneously.

Results of the simulations are illustrated by the images in figure 3. The image on the left compares the number of contacts achieved by *Cicada barbara* individuals in three situations: using the temporal pattern uniquely, using the frequency pattern and using both temporal and frequency patterns. The image on the right presents the same comparison concerning the number of contacts achieved by *Cicada orni* individuals, using its species specific pattern values.

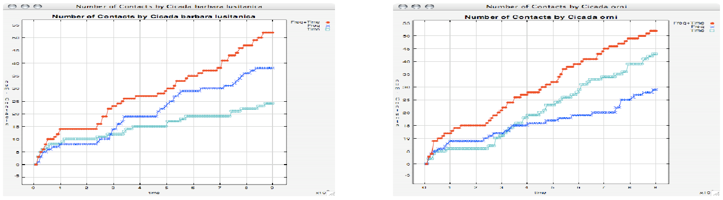

Fig. 3. Running the simulation with different singing recognition patterns used by *Cicada barbara* (left) and *Cicada orni* (right).

From the observation of graphs, we can conclude that, in *Cicada barbara* species the frequency pattern leads to a better performance, in *Cicada orni* is the contrary, and in both species the combination of the two patterns result in a more accurate perception about the origin of some calling song.

Let's consider now the coexistence, in a cicada's field, of the following behavior types:

Type I – cicadas do not use a recognition pattern and follow a song randomly;
Type II – cicadas use a temporal recognition pattern;
Type III – cicadas use a frequency recognition pattern;
Type IV – cicadas use both recognition patterns.

Cicada's paths are identified by gray marks in cells. Initially all cells are white, and they became darker when crossed by females traveling in search of conspecific males.

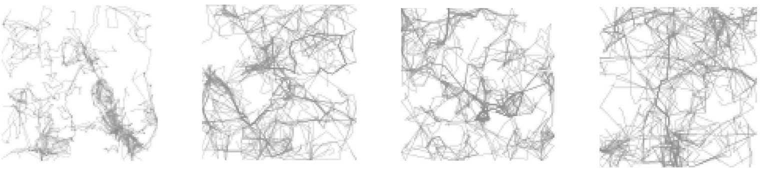

Fig. 4. From left to right: results of simulations showing Type I, II, III and IV paths.

Figure 4 represents the result of simulations, characterized by marks from Type I, II, III and IV, e.g. cicadas not using a detection pattern (on the left), using temporal or recognition pattern or using both (on the right). The extreme left image (Type I) shows that each female cicada is not moving far from a small radius, and her direction changes randomly. The result is a set of small and diffuse zones of dark cells, with no clear tracks. Males act like nodes of a network. Dark cells surround males. The center left image (Type II), concerning cicadas using the frequency pattern, shows a longer dislocation of females and the existence of clear tracks. In what concerns Type III (center right image), where cicadas use the temporal pattern, tracks are also well defined. Type IV paths (left image), regarding cicadas using both patterns, represent a thinner network, composed by a diversity of tracks. Some kind of star-like pattern concerning males positions is visible in the networks.

These experiments support the conclusion that, when cicadas use song recognition patterns, females dislocation is structured, covering the field with clear tracks. The particular case of simultaneously using both recognition patterns results in a thin network, with more tracks and more rapid access to males of the corresponding species. On another hand, when there is no use of recognition patterns, dislocations are short and direction random, with no clear existence of tracks. On another hand, the results of field and simulation experiences, support correlating the behavior of cicadas with their need of differentiating species trough communication patterns. This seems to be a clear intentional mark of cicadas social behaviors, aiming to preserve their species.

The assumption of emergent intentionality characterizing the actions observed in a group can be very useful in a diversity of social contexts, as illustrated by the following case study.

5 Smart mob case study

In Section One we introduced the notion of "smart mobs" and discussed a number of putative cases. These cases of smart mobs were characterized by viral propagation of messages through the social network of each individual. The objective was generally to achieve the gathering of a significant number of individuals at a given moment and in a given place. The result of a smart mob, whatever its goal, is on the one hand the grouping of a large number of individuals and, on the other hand, the surprise of some unintended collective action which was not previously broadcast by the media.

These basic characteristics allow the design of a generic model of smart mob dynamics, where the viral propagation of communication through the social networks of individuals coexists with the coordinated dislocation of individuals to some meeting point. The technologically-mediated cases are, in some respects, more easily modeled than the biological cases because the constraints on communication can be specified or described quite precisely.

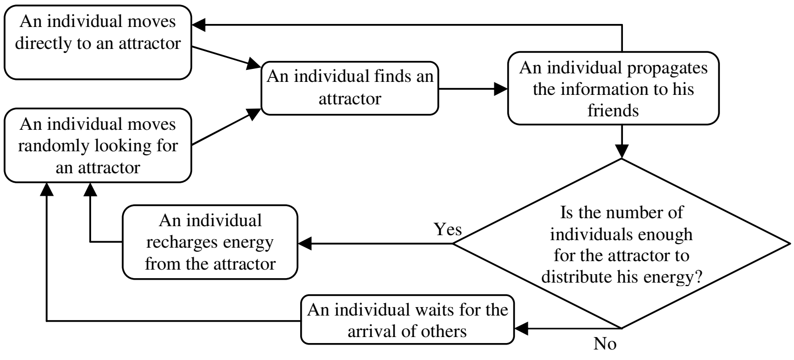

Fig. 5. Flux diagram of the behavior of individuals

Our model of smart mobs comprises two types of agents, individuals and attractors, initially placed in a bi-dimensional space. An individual circulates in the space while he has energy left; when he runs out of energy, he disappears. His goal is to recharge his energy. To achieve his goal, the individual searches for an attractor, moving randomly throughout the space. When an individual finds an attractor, he propagates this information to all his friends; consequently, they will then move towards the attractor, aiming to recharge their own energy. Also, when an individual is notified of the existence of an attractor, he will propagate this information to his friends. When the individuals surrounding an attractor exceed a given number, the attractor distributes his energy to the attending individuals; afterwards, each one reinitiates his random movement, until finding or being notified of the existence of a new attractor.

Besides the viral propagation mechanism through his social network, an individual can decide to move as a result of observing the dislocation of his neighbors. If he verifies that more than a given number of individuals is moving in the same direction, he concludes the existence of an attractor in that region of the space. On their side, each attractor accumulates energy all through the time of the simulation, and distributes his energy when a given number of individuals surrounds him.

The results of the simulation can be observed in the following images:

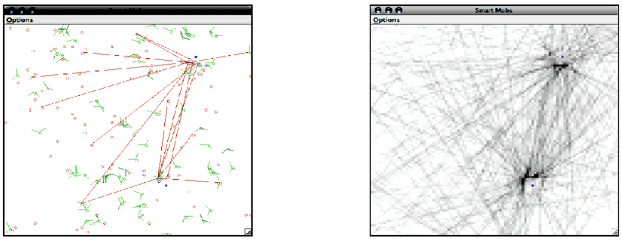

Fig. 6. The Smart Mobs model: macro observation of patterns

The image on the left of Figure 6 represents two fixed attractors and several individuals moving randomly. When an individual finds an attractor, he informs the members of his social network. These messages are represented by the traces connecting the individuals surrounding the attractor to the others dispersed throughout the space. The image on the right of Figure 10 represents a trace of the dislocation of all individuals. It is patent that straight-line movement is followed by those informed of the place of the attractors. The pattern characterizing the joining movements made during a smart mob is clearly visible. Meanwhile, individuals also move as a consequence of the dislocation they observe among their neighbors. This aspect is depicted by the following images.

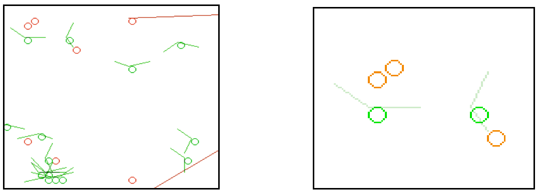

Fig. 7. The Smart Mobs model: micro observation of patterns

The image on the left of Figure 7 represents several moving individuals. Some of them are specifically directed to attain an attractor – these individuals are graphically represented by simple circles. Others move randomly – these have a cone-shaped vision, graphically represented by two traits connected to an individual (a circle), such as can be more closely observed in the image on the right.

Figures 6 and 7 illustrate the existence of patterns at different abstraction levels: the trace of the movement of the individuals indicates specific macro communication and dislocation patterns. The observer can assign the intentionality of a collective movement of the smart mob kind to these patterns. From another side, the observation by an individual of a coherent movement of his neighbors will trigger the detection of a dislocation pattern. Also, the observing individual will attribute some intentionality to this pattern, i.e., dislocation towards some attractor.

6 Discussion

The case studies described above can be discussed in the light of the research hypothesis initially announced. In both cases, we have identified a clear fingerprint that can serve as the basis for predicting and explaining the behavior of groups in natural settings. In addition, since fingerprints are identified and group intentions assigned to the structures, we can conclude that knowing only the links between the nodes of the communications network is enough to characterize the flow of information in these specific social situations.

Another research hypothesis concerns the characterization of social structures through the combination of several elementary patterns. Both case studies suggest that the combination of patterns, identified at different abstraction levels, is a way of identifying social structures determined by communication processes. In the case of *Cicada barbara*, different input (pulse and frequency) and output (structured star-like dislocation) patterns are considered. In the case of the smart mob, communication and dislocation macro patterns are also considered, simultaneously with micro patterns of dislocation of the neighbors of an individual.

On the other hand, the assignment of a given intentionality to collective behavior patterns uniquely from identifying these patterns at a macro level and without knowing the semantic content of the communication allows the attribution of specific intentions to a group of individuals. For instance, observing star-like communication patterns centered on an attractor suggests a social dynamic with the intentionality of a movement with the configuration of a smart mob.

7 Further research

The work presented in this paper incorporates a larger research program, based on the study of the mechanisms and patterns of communication in different social contexts. We consider that our approach can be advantageously applied to several application domains. This way, the research will continue through modeling and simulating other case studies, namely concerning communication domains at a micro-cellular level, as well as new domains relating to sensor networks.

The tasks of designing models and programming multi-agent simulations will be supported by the methodology and programming library developed to the case studies described in this paper. These tools, as well as other resources such as texts, images and videos, are freely available to the community on our web page: <http://www.listaweb.com.pt/projects/cells>.

Acknowledgments. We express our gratitude to Professor Paulo Fonseca, from the Faculty of Sciences of the University of Lisbon, for allowing the use of the results of his experiences in the development of the *Cicada barbara lusitanica* case study.

References

1. Dennett, D.: The Intentional Stance. MIT Press, ISBN 0262540533 (1989)
2. Dennett, D.: Kinds of Minds: Towards an Understanding of Consciousness. Science Masters Series, ISBN 0465073514 (1997)
3. Fonseca, P. J., Revez, M. A.: Song discrimination by male Cicada barbara lusitanica (Homoptera, Cicadidae). The Journal of Experimental Biology No.205, 1285-1292, Ed. The Company of Biologists Limited, U.K. (2002)
4. Goldstone, R., Janssen, M.: Computational models of collective behaviour. TRENDS in Cognitive Sciences, Vol.9, No.9 (2005)
5. Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand T., DeAngelis, D. L.: Pattern-oriented modeling of agent-based complex systems: lessons from ecology. Science, No.310, 987-991 (2005)
6. The Guardian. April 7th edition (2004)
7. Hardt, M. and Negri, A.: Empire. Cambridge MA.: Harvard University Press (2000)
8. Latané, B.: The Psychology of Social Impact. American Psychologist 36: 343–56 (1981)
9. Louçã, J., Symons, J., Rodrigues, D., Morais, A.: Pattern-oriented analysis of communication flow: the case study of *Cicada barbara lusitanica*. 21st European Conference on Modelling and Simulation - ECMS 2007, Prague, Czech Republic (2007)
10. Rheingold, H.: Smart Mobs: the next social revolution. Basic Books, ISBN 0-7382-0608-3 (2002)
11. Sallach, D., Macal, C., North, M.: Proceedings of the Agent 2006 Conference on Social Agents: Results and Prospects. University of Chicago (2006)
12. Sawyer, K.: Social Emergence: Societies as Complex Systems. Cambridge University Press (2005)
13. Seattle Post-Intelligencer. December 3rd on-line edition (1999)
14. Symons, J.: *On Dennett.* Belmont, CA.: Wadsworth (2002)
15. Symons, J.: Computational Models and Emergence. *Minds and Machines* (forthcoming)