Cognitive Modeling of Behavioral Experiments in Network Science Using ACT-R Architecture

Oscar J. Romero(✉) and Christian Lebiere(✉)

Cargenie Mellon University, 5000 Forbes, 15213 Pittsburgh, USA
oscarr@andrew.cmu.edu, cl@cmu.edu
http://fms.psy.cmu.edu/

Abstract. The Network Science has dedicated a considerable amount of effort to the study of many distributed collective decision-making processes which must balance diverse individual preferences with an expectation for collective unity. Several works have reported their results about behavioral experiments on biased voting in networks individuals, however we will focus on the results reported on [1] on which were run 81 experiments, on which participated 36 human subjects arranged in a virtual network who were financially motivated in a heterogeneous manner and whose goal was to reach global consensus to one of two opposing choices. Multiple experiments were performed using diverse topological network configurations, different schemes of financial incentives that created opposing tensions between personal preferences, and finally different ratios of both inter and intra-connectivity among the network nodes. The corresponding analysis of the results demonstrated that changing those features of the experiments produced different kind of social behavioral patterns as a result. Thus, the purpose of this work is manifold: on the one hand, it aims to describe the possible structures that underlie the decision-making process of these experiments through the modeling of symbolic cognitive prototypes supported by a robust and complex cognitive architecture so-called ACT-R and, on the other hand, by applying modifications in the ACT-R parameters to find those subtle aspects that can either influence both the performance and speed of convergence of the experiments or cause the total inability to reach a global consensus in a reasonable amount of time.

Keywords: Cognitive modeling · Social behavior · Network science · Multi-agent systems · Cognitive architecture ACT-R · Coloring problem

1 Introduction

Most studies in the field of network science focus on the analysis of both individual and global behaviors in the absence of incentives [2–4], however, in order to observe the complex interactions that all elements of the network can show when opposing incentives schemes are present, some scenarios in which individual preferences are present but subordinate to achieving unanimous global consensus are

F. Grimaldo and E. Norling (Eds.): MABS 2014, LNAI 9002, pp. 239–251, 2015.
DOI: 10.1007/978-3-319-14627-0_17

proposed. A real example of such phenomenon appears in decision-making and voting processes in politics, business and many other fields.

Due to the main objective of our work is to design a cognitive model of the experiments done by [1], we focused on both the methodological aspects and the behavioral results obtained in each experiment in order to recreate these behaviors through cognitive models that were then executed in the ACT- R cognitive architecture.

In Sect. 2, the methodological fundamentals that were considered for the execution of the human experiments will be explained. In Sect. 3 the basic design principles and modules of the ACT -R architecture will be introduced. In Sect. 4 the computational cognitive modeling that reflects the results obtained in the experiments with humans will be presented and the design aspects that were taken into account for the implementation of the multi-agent system that supported the interaction among nodes (cognitive agents) will be detailed. Section 5 will analyze and discuss the results and finally, in Sect. 6 the conclusions obtained from this research will be presented.

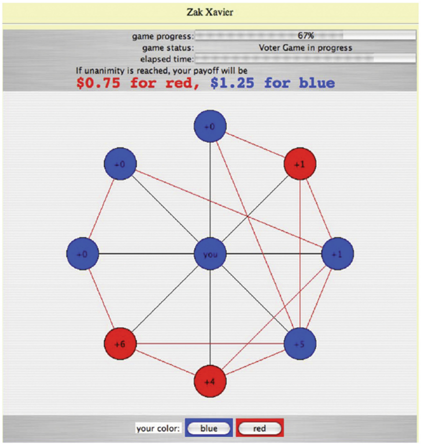

Fig. 1. Each subject sees only a local ("ego network") view of the global 36-vertex network, showing their own vertex at the center and their immediate neighbors surrounding. Edges between connected neighbors are also shown, as are integers denoting how many unseen neighbors each neighbor has. Vertex colors are the current color choices of the corresponding subjects, which can be changed at any time using the buttons at the bottom. The subjects payoffs for the experiment are shown (in this case \$0.75 for global red consensus, \$1.25 for blue), and simple bars show the elapsed time in the experiment and the "game progress" a simple global quantity measuring the fraction of edges in the network with the same color on each end. This progress bar is primarily intended to make subjects aware that there is activity elsewhere in the network to promote attention, and is uninformative regarding the current majority choice. This figure has been taken from [1] (Colour figure online)

2 Consensus Problem

The aim of the original experiment with humans in [1] was a typical problem from the graph coloring category which used a 36-vertex network. Every single node was only able to see a portion of the network: their closest neighbors as shown in Fig. 1.

Three different network configurations were used for the experiments: Preferential Attachment, Erdos-Renyi model and Minority Power model. Preferential Attachment model (PA) is an stochastic process in which additional nodes are added continuously to the system and are distributed among the network as an increasing function of the number of neighbors that every node already has. One of the most known scale-free networks of PA is the Barabási model [5]. Erdos-Renyi model (ER) generates random graphs, including one that sets an edge between each pair of nodes with equal probability, independently of the other edges [6], and Minority Power model (MP) is a principle mainly derivated from politics and voting models, which says that a minority of persons could attract members from the other parties, in order to do anything at all. In the following, the parameters of the experiments are described:

- Number of experiments (81): 27 with PA, 27 with ER and 27 with MP.
- Vertices (36): Initially, half of vertices (18) were randomly selected and said that red color receives the highest incentive whereas the remaining 18 were said the opposite. The exception to this rule is the MP model, on which a minority of the vertices with the highest number of neighbors were then assigned incentives preferring red global consensus to blue, whereas the remaining majority were assigned the opposite. The size of the chosen minority was varied (6, 9, or 14).
- Edge count (101 ± 1): All of the networks had 36 vertices and nearly identical edge counts. Only the arrangement of connectivity varied.
- Connectivity (inter or intra): It controls whether local neighborhoods were comprised primarily of individuals with aligned incentives (high cohesion $\rightarrow$ 1:2 inter:intra ratio), competing incentives (low cohesion $\rightarrow$ 2:1 inter:intra ratio), or approximately balanced incentives (1:1 inter:intra ratio)
- Financial incentive ($0.25–$2.25): It specifies the incentive for reaching a consensus on blue or red color which was arbitrary assigned to every node. "Strong symmetric incentive": $1.50 vs. $0.50; "Weak symmentric incentive": $1.25 vs. $0.75; and "Asymmetric incentive": $2.25 vs. $0.25.

3 Cognitive Architecture

The cognitive model was developed using the ACT-R cognitive architecture [7], [8]. Cognitive architectures are computational representations of invariant cognitive mechanisms specified by unified theories of cognition. ACT-R is a modular architecture, reflecting neural constraints, composed of asynchronous modules coordinated through a central procedural system as depicted in Fig. 2.

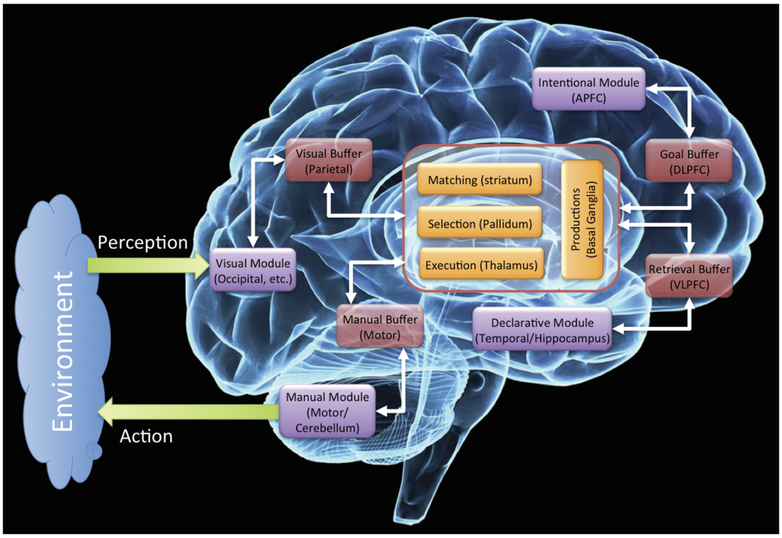

Fig. 2. ACT-R Cognitive Architecture

The procedural system is in charge of behavior selection and more generally the synchronization of the flow of information between the other modules. It is implemented as a production system where competing production rules are selected based on their utilities, learning through a reinforcement mechanism from the rewards and costs associated with their actions. The production system conditions are matched against limited-capacity buffers that control the interaction with the other modules by enabling a single command (e.g., retrieval of information, focus of visual attention) to be given at a time to a given module, and a single result to be returned (e.g., chunk retrieved from memory, visual item encoded). A declarative memory module holds both short-term information, such as the details of the current situation, as well as long-term knowledge, such as the procedural rules to follow. Access to memory is controlled by an activation calculus that determines the availability of chunks of information according to their history of use such as recency, frequency, and degree of semantic match. Learning mechanisms control both the automatic acquisition of symbolic structures such as production rules and declarative chunks, and the tuning of their subsymbolic parameters (utility and activation) to the structure of the environment. The perceptual-motor modules reflect human factor limitations such as attentional bottlenecks. Individual differences can be represented both in terms of differences in procedural skills and declarative knowledge, as well as in terms of architectural parameters controlling basic cognitive processes such as spreading of activation.

4 Cognitive Modeling

In the following, the computational cognitive mechanisms used for simulating the results of the social experiments obtained in [1] will be described. These mechanisms will be detailed in 4 subsections as follows:

4.1 Symbolic Productions

As mentioned before, the procedural system uses production rules which interact with different kind of buffers (retrieval, imaginal, declarative, visual, motor, and others) in order to carry the reasoning and inference process out according to the information that every node of the network (from now, the terms ‘cognitive agent’ and ‘node’ are equivalent) senses from its environment and acts over it. We have implemented several strategies to model the social experiments, and every one of these has a set of different productions as described below:

Follow-the-majority scenario: In this scenario, every cognitive agent senses the environment through the features in its visual buffer: (1) which color is the majority in the neighborhood (blue or red); (2) which is the current higher payoff for changing to either blue or red; and (3) whether or not is the majority color increasing through the time. As response, the cognitive agent can perform an action changing its color to either “blue” or “red”. Thus, there are 4 variables (3 inputs and 1 output) and they are binary, so we have $2^4 = 16$ possible states which were reflected as productions. An example is shown in Fig. 3.

```
(p majority-red-payoff-red-majority-decrease-then-red
  =visual-location>
    isa visual-location
    !eval! (equal =increasing-majority false)
    !eval! (>= =red-color =blue-color)
    !eval! (> =payoff-red =payoff-blue)
==>
   +vocal>
    isa speak
    string "red")
```

Fig. 3. Production that change to red when: the majority color is red, the higher payoff is given when choosing red color and the amount of nodes of the dominant color is decreasing over time (Colour figure online)

Follow-the-most-influential scenario: In this scenario, some opposing productions compete against the others to obtain the global control over the decision-making process that performs the agent. Some productions follow the majority, some others follow that agent which is the most influential over the neighborhood (that is, that node which has more unseen connections, keeps its color for more time and when it changes its color there are a considerable amount of seen connected agents that do the same), as shown in Fig. 4; some other productions just keep the same color no matter if the environmental conditions are not favorable for that, that is, the stubborn productions. The productions are depicted in Fig. 5:

```
(p increasing-then-change-to-majority
    =visual-location>
      isa visual-location
      !eval!   (eq =increasing true)
==>
    +vocal>
      isa speak
      string "change-to-majority")
```

(a) if the dominant color is increasing then change to the dominant color

```
(p decreasing-same-majority-then-change-influential
    =visual-location>
      isa visual-location
      !eval!   (eq =decreasing true)
      !eval!   (eq =current-color =majority-color)
==>
    +vocal>
      isa speak
      string "change-to-influential")
```

(b) if current color is the same as the majority and its amount is decreasing then change to the most influential agent

```
(p increasing-then-keep-color
    =visual-location>
      isa visual-location
      !eval!   (eq =increasing true)
==>
    +vocal>
      isa speak
      string "keep-color")
```

(c) if global consensus is increasing then keep the same color

```
(p decreasing-different-majority-then-change-majority
    =visual-location>
      isa visual-location
      !eval!   (eq =decreasing true)
      !eval!   (not (eq =current-color =majority-color))
==>
    +vocal>
      isa speak
      string "change-to-majority")
```

(d) if current color is not the same as majority and its amount is decreasing then change to majority

Fig. 4. Productions related to the color changes in the cognitive agent's neighborhood

```
(p stuck-then-change-to-influential
    =visual-location>
      isa visual-location
      !eval!   (eq =increasing false)
==>
    +vocal>
      isa speak
      string "change-to-influential")
```

(a) if global consensus is stuck then choose stochastically the color of an influential cognitive agent in the neighborhood

```
(p stuck-then-change-another-majority
    =visual-location>
      isa visual-location
      !eval!   (eq =increasing false)
    ==>
    +vocal>
      isa speak
      string "change-another-majority")
```

(b) if global consensus is stuck then choose stochastically the color of another majority in the neighborhood

```
(p stable
    =visual-location>
      isa visual-location
      !eval!   (eq =increasing true)
==>
    +vocal>
      isa speak
      string "keep-color")
```

(c) if the global consensus is increasing then keeps the same current color

```
(p stubborn
    =visual-location>
      isa visual-location
      !eval!   (eq =decreasing true)
      !eval!   (not (eq =current-color =majority-color))
==>
    +vocal>
      isa speak
      string "keep-color")
```

(d) if current color is not the same as majority and this is decreasing then keeps the same color

Fig. 5. Productions related to the cognitive agent's internal motivations

As you can infer from Figs. 4 and 5, all the productions generate opposing tensions and a continuous competence for being the production to be fired. For example, "increasing-then-change-majority" production senses the same information as "increasing-then-keep-color" production but they trigger different actions: the first one will change its color in order to follow the majority and the second one will keep its current color if the majority is increasing (whichever its color is). Similar antagonisms are observed in the rest of productions. Due to the fact that multiple productions may match the same sensory input or that sometimes there is not a production which has a perfect match with the sensory input, a selection process based on production utilities and partial matching is required.

It is important to remark that both scenarios have another meaningful difference related to the production selection process: *follow-the-majority* scenario is completely deterministic, thus it calculates the number of seen/unseen nodes, the global consensus and the majority using always the highest number, whereas in the *follow-the-most-influential* scenario uses a stochastic selection process based on the Boltzman equation [8] as shown in Eq. 1.

$$P_i = \frac{e^{\frac{M_i}{t}}}{\sum_j e^{\frac{M_j}{t}}} \tag{1}$$

P_i is the probability that cognitive agent i follows agent j according to the function M, which can be either the majority or the most influential agent. t is the temperature which determines the randomness of the process and it is set at 0.35 for convenience.

4.2 Reinforcement Learning

The reinforcement model of ACT-R supports the utility learning mechanism of the architecture. The utilities of productions can be learned as the model runs based on rewards that are received from the environment. The utility of every production is updated according to a simple integrator model. If $U_i(n-1)$ is the utility of a production i after its $n-1st$ application and $R_i(n)$ is the reward the production receives for its nth application, then its utility $U_i(n)$ after its nth application will be as in Eq. 2 (typically, the learning rate α is set at 0.2).

$$U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)] \tag{2}$$

In our experiments, cognitive agents were requested to maximize their expected total reward over a given number of trials and learn about the structure of the environment by taking into account the reward associated with each choice. Due to the fact that we used two modeling scenarios for the productions, we proposed two different reinforcement algorithms for each one of these.

Follow-the-majority Reinforcement: the reward is equal to the $higher - payoff \times k$ (with k as a constant) when the unanimity is reached by all individuals. Otherwise, if the current color of a node is equal to the majority color (observed in its neighborhood) then the reward is equal to $payoff/60$. If the current color is equal to the global consensus color the node receives and extra reward of $payoff \times 0.5$. If the node change the color when the majority decreases or keep the same color when the majority increases then the reward increases $payoff \times 0.01$. Otherwise, the payoff increases $payoff/20$.

Follow-the-most-influential Reinforcement: if the current color of the node is equal to the observed majority, then the reward is equal to: $payoff/100$. An additional reward is received from the influence that node i has over its neighbors, so if color of node i is equal to neighbor j then node i receives and increment of $k_1 = 5$. An additional reward comes from the validation whether the current number of nodes in consensus is higher than the previous number of nodes in consensus, in that case the reward would increment $payoff/60$ Otherwise would decrease $payoff/20$.

4.3 Multi-agent Approach

The experiments were run over a multi-agent platform on which multiple cognitive models interact through perception and action processes as shown on Fig.6

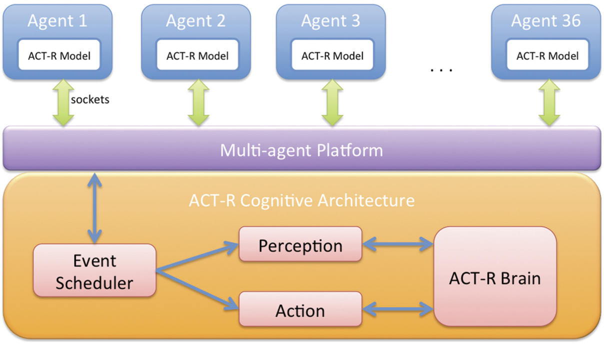

Fig. 6. Multiple cognitive models interact with the ACT-R architecture through a multi-agent platform. Socket channels make easier the communication between layers and an event scheduler is in charge of triggering the events of perception and action for every cognitive agent. Even though all the models share the same cognitive architecture (ACT-R), each one has its own separate set of productions, declarative memories, partial matching selection process and buffer contents; which are carried out through a multithreading approach.

5 Experimentation

The purpose of the experiments is manifold: firstly it aims to evaluate both the performance and convergence speed to reach a global consensus, secondly it expects to get close similar results between the cognitive simulation and the real

Fig. 7. Visualization of the collective dynamics for all 27 experiments

experiments with humans, thirdly it seeks to find meaningful variances regarding to the different network configurations and, finally, it aims to evaluate the differences between the different scenarios and modeling strategies proposed.

Results According to the Network Configuration: each one of the 27 experiments (3 network models $\times$ 3 intra-inter connectivity topologies $\times$ 3 incentive schemes) were run 100 times and the results were averaged using the harmonic mean value: 2700 experiments in total. The results are shown in Fig. 7.

A general remark about the results is that we had to include opposing and competing productions (such as making the wrong decision, following the minority, or being stubborn and never change the original color selection) in order to produce interesting convergence curves within the 60 seconds that lasted every experiment, otherwise the convergence was promptly reached and no meaningful differences were observed between all network configurations. In addition to that, only the results from the 'follow-the-most-influential' strategy are shown in Fig. 7 because of their similarity with the original experiments with humans.

From Fig. 7 we can remark some issues: In the Babarási experiments, the network configuration with a strong-symmetric incentive and a 2:1 inter-intra connectivity had the quickest and most stable (that is, the one that had the least fluctuations) curve of convergence. In a similar manner, both the Erdos-Renyi and the Minority Power experiments evidenced a quickest and most stable unanimous consensus whit a network configuration composed of a weak-symmetric incentive scheme and a 2:1 inter-intra connectivity. From left to right we can observe that Cohesion experiments (Babarási and Erdos-Renyi) gradually increase the degree of convergence when the number of inter-connections are augmented and symmetric incentives are kept. Asymmetric incentives worked better with both 1:2 and 1:1 inter-intra connectivity topologies. On the other hand, Minority Power experiments reached an unanimous consensus quicker than the Cohesion experiments. Specifically, an asymmetric incentive scheme seemed to work better when the minority was smaller (9 and 6 cognitive agents) in which case the majority was quickly influenced by the well-connected minority (at $\cong$ 30 s) despite of the fact that appreciable fluctuations were present at the early seconds.

Comparison Between Agent-Based vs. Human-Based Experiments: Table 1 summarizes the similarities and differences between both kind of experiments. Agent-based experiments were divided into 2 strategies: follow-the-majority (strategy-1) and follow-the-most-influential (strategy-2). Rather than looking for the perfect accuracy of the obtained results between human and agent-based experiments, the most pertinent question concerns whether the results are proportional and whether they are significantly different from each others. From Table 1 is possible to infer, for example, that in human-based experiments, the Minority consensus (88.89) got a proportional increment of success of 16.67 % over the Babarási consensus (74.07) and 54.16 % over the Erdos-Renyi consensus (40.74), whereas for the strategy-1 this proportional increment was 19.34 % and 24.19 % and for strategy-2 was 17.73 % and 53.3 % (the latter strategy was very close to the results of the human-based experiments). In general, Strategy-2 (*follow-the-most-influential*) obtained more similar results in relation to the human-based experiments than Strategy-1, suggesting that the more an agent supports its decision on the most influential agent in the local neighborhood instead of the local majority, the quicker the network reach an unanimous consensus and the more similar the results are in comparison with human-based experiments.

In order to estimate whether there is a significant difference between the results, we have performed an analysis of the variance of the three groups (the human-based and the two agent-based experiments) through an Anova test. Starting from defining as the null hypothesis: *"h_0: all the three groups of experiments are equal and do not reflect meaningful differences"*, we executed an Anova single factor test ($P < 0.01$) which analyzed both between and within group variances. As result, we got that $F < F_{crit}$ ($0.25 < 2.38$) for the variance analysis between human-based experiments and strategy-2 experiments, which means that null hypothesis may not be rejected and implying that strategy-2 does model in a more similar way the global and individual behaviors observed in the

Table 1. Comparison between agent-based vs. human-based experiments. St-1 means Strategy-1, PI is the Proportional Increment (%) of the highest value marked with an ∗ in relation to the other two values on the left column.

Feature	**Human**	**PI %**	**St-1**	**PI %**	**St-2**	**PI %**
Global Consensus all exp.(% succ.)	67.90	–	58.89	–	63.15	–
Averaged Convergence (sec.)	43.90	–	55.60	–	32.26	–
Standard Deviation	9.60	–	2.76	–	2.81	–
Mean Square Error	–	–	16.12	–	4.90	–
Consensus Babarási (% succ.)	74.07	16.67	55.56	19.34	56.67	17.73
Consensus Erdos-Renyi (% succ.)	40.74	54.16	52.22	24.19	32.17	53.30
Consensus Minority (% succ.)	88.89	∗	68.89	∗	68.89	∗
2:1 inter:intra (% success)	77.78	∗	64.44	∗	70.00	∗
1:1 inter:intra (% success)	44.44	42.86	56.67	12.05	42.23	39.67
1:2 inter:intra (% success)	50.00	35.71	55.56	13.78	62.22	11.11
Weak-symmetric (% success)	70.37	13.93	61.11	-1.85	60.00	8.48
Strong-symmetric (% success)	51.85	36.58	55.56	7.4	48.89	25.42
Asymmetric (% success)	81.76	∗	60.00	∗	65.56	∗

social experiment. On the other hand, we got that $F > F_{crit}$ (1.31 > 0.98) for the variance analysis between human-based experiments and strategy-1 experiments, reflecting that there is a significant difference between the experiments and therefore strategy-1 does not properly model the social behavior experiment.

6 Conclusions and Future Work

According to the results discussed on Sect. 5, the agent-based strategy that follow-the-most-influential agent in the local neighborhood seems to simulate better the results obtained by the social experiments with humans. However, there are some issues that are worth being discussed: The strategy-2 (following-the-most-influential-agent – that agent who has more seen and unseen connections with other agents) turns out to be a key differentiator compared to strategy-1 (follow-the-majority) because the former plays with the uncertainty of what color decisions are operating behind these unseen connections instead of just following the current state of the local neighborhood which in most of the cases comes up either in arbitrary fluctuations or in local sub-networks that do not want to change their preferred color.

Nevertheless, it is important to remark that strategy-2 used a stochastic selection process instead of a deterministic one as in strategy-1. The stochastic approach avoids both falling in recurrent states on which different parts do not come to an unanimous agreement and forming sub-regions with different color choices. Due the stochastic nature of the selection process, agents can follow sometimes the most influential agent, sometimes the local majority or sometimes just becoming in a stubborn agent. However, because of the stochastic process uses a temperature factor that controls the randomness during the experiment execution, better decisions are more likely to be made instead of bad or non-conciliator decisions.

We have chosen these two main strategies because they were the strongest ones identified in the social experiments with humans. Humans used another kind of strategies, of course, but they were irrelevant for the cognitive modeling work because they mostly rely on exogeneous conditions such as boredom because of the way the experiments were carried out, changes in temperature in the room, misunderstandings about the purpose of the experiment and so on.

Regarding to production utilities, they play an important role because they reflect which decisions may favor a global consensus in the future by adjusting their current relevance through a learning process, which will determine which productions will be more likely to trigger in future situations. In spite of the fact that strategy-2 got similar results in comparison with the human-based experiments, the results could be improved by modifying the reinforcement functions of the utility learning process in order to better reflect how the stubbornness of some agents may affect both the convergence towards an unanimous consensus and the global performance. Stubbornness was a key factor for the simulation because it reflected natural social phenomena such as the indecision of some humans, the conflicting interests or simply the absence of attention of some humans who miss the dynamics behind the scenes during the experiment execution and then kept always the same preferred color.

Along these lines, stubbornness should be consider as a key aspect that reflect a natural aspect of human decision-making that should be considered in depth when carrying out cognitive modeling. It might be worth mentioning that cognitive architectures provide a principled framework to model individual differences in both knowledge and capacity, and that a population of individual agents with variations in capacity might provide some results that are fundamentally different from any that can be generated with a uniform population, i.e., the network aspect of the domain provides some non-linear dynamics that the typical averaging in cognitive experiments doesn't address. As a final remark, modeling social behaviors is an complex task which should have into account some other aspects of human decision-making such as mood states, intentions, expectations, game strategies and so and so forth. Modeling these aspects probably would improve the accuracy of our experiments.

Acknowledgments. This work was conducted through collaborative participation in the Robotics Consortium, Agreement W911NF-10-2-0016.

References

1. Kearn, M., Judd, S., Tan, J., Wortman, J.: Behavioral experiments on biased voting in networks. Nat. Acad. Sci. **106**(5), 1347–1352 (2009)
2. Kleinberg, J.: Cascading behavior in networks: algorithmic and economic issues. Algorithmic Game Theory **24**, 613–632 (2007)
3. Granovetter, M.: Threshold models of collective behavior. Am. J. Sociol. **83**, 1420–1443 (1978)

4. Schelling, T.: Micromotives and Macrobehavior. Norton, New York (1978)
5. Barabási, A.: Emergence of scaling in random networks. Science **286**, 509–512 (1999)
6. Bollabass, B.: Random Graphs. Cambridge University Press, Cambridge (2001)
7. Anderson, J.R., Lebiere, C.: The Atomic Components of Thought. Lawrence Erlbaum Associates, Mahwah (1998)
8. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., Qin, Y.: An integrated theory of the mind. Psychol. Rev. **111**(4), 1036–1060 (2004)