Coalition Formation of Cognitive Agents

Myriam Abramson
Naval Research Laboratory
Washington, DC 20375
myriam.abramson@nrl.navy.mil

Abstract

Part of the knowledge and set of beliefs of a cognitive agent are its mental models of the world and of other agents. A mental model reflects the cultural context and past experiences of an adaptive agent. In interacting with other agents, we claim that a causal mental model, i.e. a cognitive map, which might be different for each agent, adapts and changes causing coalitions to emerge. Possible actions are reflected through the respective cognitive maps of the agents to determine whether a favorable steady state will emerge from a coalition. The problems addressed are which coalitions will be more likely to form and how to adapt the cognitive maps of the agents in order to reduce conflicts. The claim is that a group mental model can emerge from individual mental models through the learning and adaptation of cognitive maps. An algorithm based on Particle Swarm Optimization (PSO) for evolving cognitive maps for coalition formation is introduced and experiments on randomly generated cognitive maps are presented. We conclude with possible uses of this adaptive cognitive modeling approach to understand other cultures, predict coalitions/chaos and shifts of allegiance, and induce group formation through avatars in a virtual world.

Introduction

Cognitive agents are distinguished from state-based reactive agents by their capability to be proactive in the pursuit of their goals, desires, and intentions. In their interactions with other agents, beliefs on the indirect outcome of possible actions and not just the current situation, can influence the degree of cooperation or conflicts between agents. In this approach, agents adapt their belief system, represented in a cognitive map, in response to their social environment represented by the belief system of other agents they know about. Very few efforts have concentrated on team formation based on the inference and modification of beliefs. Teamwork of SOAR agents in STEAM (Tambe 1997), based on the theory of joint intentions (Cohen & Levesque 1991), does assume shared mental properties of choice and commitment, but relies on the explicit communication of those mental properties to achieve mutual belief rather than learning those mental properties through experience. Game-theoretical approaches on coalition formation do not take into account the inference and adaptation capability of cognitive agents from interactions with other agents but rather the deterministic position and influence of the agents in a social network (Efird 2008) evolving over time. In this paper, we claim that teams or coalitions are formed when the cognitive maps of the individuals are consonant or complementary with each other over time.

This paper is organized as follows. Cognitive maps are first introduced as a representation for the mental model of an agent and its inference mechanism. The PSO algorithm and our approach to optimize cognitive maps for team formation is then presented with some experimentation results. The approach is then related and contrasted to other works. Finally, we conclude with future work and possible uses of this approach.

Cognitive Maps

Cognitive maps are graphical models of perceived cause-and-effect relationships between concepts, events, and/or actions expressed as directed edges between nodes (Fig. 1). They differ from other graphical representations for problem solving, such as Bayesian belief nets and influence diagrams, mainly because feedback loops, i.e cycles, are possible. Cognitive maps are therefore well suited to represent complex models of interactions that evolve with time. Positive or negative causality are specified on the edges to indicate whether an increased strength in a causal node effect an increased or decreased strength in the related node (see Fig. 1). Fuzzy cognitive maps (Kosko 1994) further expand this representation by assigning a value to the edges in the fuzzy causal range $[-1,1]$ and a value to the nodes in the fuzzy range $[0,1]$.

Cognitive maps can be extracted from communications or questionnaires (Axelrod, Nozicka, & Shapiro 1976) and synthesized from various sources representing a collectivity (Abramson 2007). The additivity property of cognitive maps makes it an ideal representation for combining the knowledge of various experts (Kosko 1986). This combined knowledge is obtained as a linear combination of the individual cognitive maps by superimposing them at the nodes

---

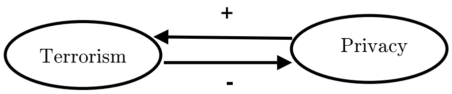

Figure 1: Simple cognitive map describing the opinion that terrorism dilutes privacy (negative causality) because of the necessary increased surveillance but that increased privacy, through anonymity for example, facilitates terrorism (positive causality).

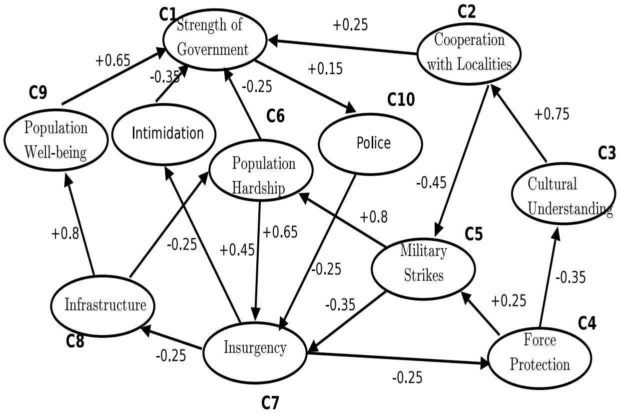

Figure 2: Counter-insurgency Cognitive Map manually extracted from article *To Defeat the Taliban – Fight Less, Win More* by N. Fick, The Washington Post, August 12, 2007.

and taking the average strength at the edges. Combined cognitive maps exhibit more feedback loops than can be found in a single cognitive map because of coherence constraints in expressing individual beliefs (Axelrod, Nozicka, & Shapiro 1976). The dynamical aspect of agent interactions in the real world is therefore stressed in a combined cognitive map. Figure 2 shows a cognitive map manually extracted from a recent newspaper article.

Agent Mental Model

Cognitive maps have been used to represent the reactive behavior of agents in a virtual world in response to their perceived environment state as a sequence of composite behavioral states. In this context, the specific agent interactions mediated by events in the virtual world are encoded as causal links between disjoint cognitive maps (Dickerson & Kosko 1998). In our approach, the cognitive map of an agent consists of concept nodes, utility nodes representing the desirable/undesirable states, and “policy” nodes that represent the possible actions of the agent (see Fig. 3) (Axelrod, Nozicka, & Shapiro 1976). “Policy” nodes trigger conceptual events represented by concept or utility nodes but are not themselves triggered by other nodes and are either on or off while concept and utility nodes are associated with real-valued strength in the [0,1] range. The structure of a cognitive map is a categorization indicating which nodes are “policy” nodes, which nodes are concept nodes, and which nodes are desirable/undesirable utility nodes. The causal links represent positive or negative effects of nodes on other nodes in the [-1,1] range. The cognitive maps of the different types of agents are not necessarily disjoint and concepts overlap in the augmented combined map of an interaction.

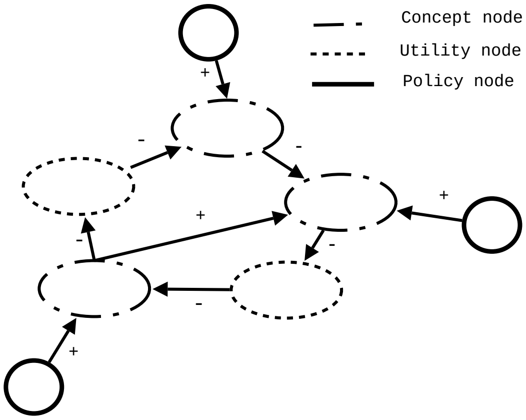

Figure 3: Agent mental model as a cognitive map

Inference Mechanism

The strength A at the nodes can be inferred as an adaptive function f of the sum of all incoming edges W times the value of the causal nodes as follows:

$$A_i(t+1) \leftarrow f(A_i(t) + \sum_{j=1, j\neq i}^{n} W_{ji} A_j(t)) \quad (1)$$

$$f(x) = \frac{1}{1+e^{-x}} \quad (2)$$

If the function is the sigmoid squashing function (Eq. 2), A is bounded within $[0,1]$ and can be evaluated comparatively with other value nodes. Fuzzy cognitive maps have been learned successfully as an associative neural network (Kosko 1992). It is also possible to learn a cognitive map through dynamic programming and Monte Carlo simulations using the conditional probabilities of the effect nodes given their causal nodes extracted through the data mining of opinion rules (Abramson 2007). The output of the learning process after several iterations indicates whether there is a convergence to a fixed set of values, to a cycle of values, or to chaos if no convergence was possible.

Particle Swarm Optimization

The PSO algorithm (Alg. 1) is based on the cultural learning metaphor (Kennedy & Eberhart 2001) where an agent, represented as an n-dimensional feature vector, adapts its solution in the problem space from its social interactions. The search through the problem space is controlled by an n-dimensional velocity vector giving the learning agent a “particle” movement characteristic. Two types of interaction are usually distinguished: (1) a top-down type of interaction based on normative knowledge of the “global best”,

Algorithm 1 Basic PSO Algorithm

```
Initialize weight vectors $x$,
  parameters $vmax, w, \varphi_1, \varphi_2$
DO
  $p_g \leftarrow \arg\max_i f(p_i)$ % global best particle
  FOREACH particle $i$
    $p_l \leftarrow \arg\max_j f(p_j)$ % local best particle
    FOREACH dimension $d$
      $r_1 \leftarrow random(0,1)$
      $r_2 \leftarrow random(0,1)$
      $v_{id} \leftarrow wv_{id} + \varphi_1 r_1 (p_{ld} - x_{id}) + \varphi_2 r_2 (p_{gd} - x_{id})$
      $v_{id} \leftarrow sign(v_{id}) min(abs(v_{id}), vmax)$
      $x_{id} \leftarrow x_{id} + v_{id}$
    ENDFOR
  decrement $w$
  ENDFOR
UNTIL termination criterion
```

gbest, and (2) a bottom-up type of interaction based on internal and neighborhood knowledge of the “local best”, *lbest*. Additionally, lbest acts as the agent’s episodic memory of past performances. The cognitive and social influences are modulated by the stochastic parameters φ_1 and φ_2 respectively. The shape of neighborhood models affects the convergence of a swarm (Liu *et al.* 2006). An inertia parameter w, decreasing with time, acts as the momentum in neural networks in controlling the exploration-exploitation tradeoff of the search. The velocities v are bounded to a value $\pm vmax$ to control the search. As agents interact, subsets of the population become more similar and successful together. As in semantic maps (Kohonen 1997), the topology of the population determines the local best neighbor and not the similarity of the vectors themselves as is found in instance-based learning.

Particle Swarm Optimization of Cognitive Agents

Our scenario involves a population of N agents with associated cognitive maps represented as a matrix with a maximum of m nodes where the rows represent the causal nodes and the columns the effect nodes. Each cognitive map represents an agent mental model with u desired/undesired utility nodes. Actions, represented by clamped binary “policy” nodes, will trigger different inferences depending on the associated causal beliefs. The weights and sign of the causality factors have to be evolved from the interactions of the cognitive agents such that positive interactions will increase and negative interactions will decrease. The basic agent loop for social, cognitive agents is described as follows:

1. Select partner to interact with
2. Interact with partner
3. Change partner preference
4. Modify internal cognitive state

In this context, a *round* or *cycle* consists of executing steps 1 and 2 and then, assuming all interactions occur *simultaneously,* steps 3 and 4. In synchronous environments, those two alternate modes, *active* and *passive*, are executed in locksteps. In asynchronous environments, the time required for executing a round depends on the internal clock of the agent. Depending on how fast the internal clock of the agent is relative to the other agents and the environment, a discrepancy between the social environment and the internal cognitive state of the agent will occur.

Algorithm 2 Cognitive, social agent loop

```
initialize partner preferences
initialize cognitive map
lbest ←cognitive map
active ←true
WHILE (no termination condition) DO
   IF (active) THEN
      select partner from neighborhood
      interact with partner
      update lbest from interaction outcome
      update partner preference
      active ←false
   ELSE
      get gbest from neighborhood
      update cognitive map from gbest
              and lbest
      move closer or farther from partner
      active ←true
   ENDIF
END WHILE
```

Proposed Algorithm

Coalitions are formed when individual members of the population find partners with which they can cooperate with. There is a conflict between the social optimum achieved and the individual optimum in the short term. In the long term, team cooperation will lead to accumulated benefits through stable individual performance. Through trials and errors and adaptation, agents will know which partners to cooperate with and will self-organize into groups. The fitness of an interaction for an individual in the population is represented by the strength ratio of the sum of desirable outcomes over the sum of all outcomes in the final output vector. This fitness will drive the interaction preference of an agent toward another agent. Those preferences are visualized by moving closer or farther to each other on a 2-dimensional grid. Unlike other models of social interaction that assume a fixed neighborhood or a fixed social network (Schreiber, Singh, & Carley 2004; Latane 1996), the agents through this proposed algorithm will reduce the search space by progressively redefining their own neighborhood.

Each agent knows about other agents in its neighborhood ranked by preference. An interaction partner is selected through fitness proportional selection based on those preferences. The interaction outcome after the inference process on the combined cognitive maps will redefine the partner preference as a basis for the next encounter. The neighborhood is redefined by moving closer to partners with positive

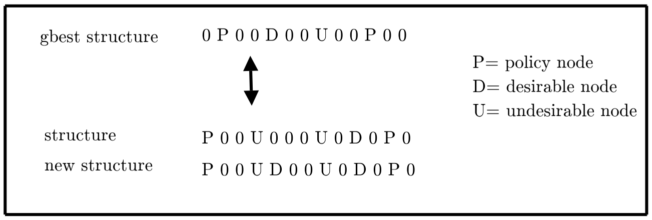

Figure 4: Cognitive map structure imitation from gbest. The categorization of the nodes is swapped with a fitness proportional selection based on the node strength from the last *gbest* interaction and the commonality of the node in the agent's neighborhood.

interaction outcomes and further from partners with negative interaction outcomes. The neighborhood size is slowly decreased as the simulation progresses successfully. The cognitive map of an agent is updated from the local and global best by adjusting the structure of the nodes and the strength of the causal beliefs between nodes. The structure of the nodes is represented as a multi-valued classification vector. Values of this node structure are swapped with a fitness proportional selection when different from gbest (Fig. 4) and where the fitness of a node from gbest is evaluated by its commonality in the neighborhood according to social impact theory (Latane 1996). The strength of the causal beliefs are updated with the basic PSO algorithm (Algorithm 1). Coalitions will form as the structural and functional similarity of the cognitive maps evolve. Further work should explore a unified PSO algorithm for evolving the form and function of a cognitive map in a social context. Because of the non-linearity of the interaction outcomes due to differences in structure, diversity between self-organized groups will be preserved.

Empirical Evaluation

The experiments were conducted with RePastJ (Collier 2001), an agent-based simulation and modeling tool where agents act concurrently in a decentralized manner.

Evaluation metrics

The current fitness of an individual in the population is represented by the strength ratio of desirable outcomes d_i over the sum of desirable d_i and undesirable u_j outcomes:

$$Fitness = \frac{\sum_i d_i}{\sum_i d_i + \sum_j u_j + \varepsilon} \tag{3}$$

where ε prevents a division-by-zero. A convergence exception is thrown to nullify a match in the absence of convergence within a pre-specified number of steps. Further work should detect/evaluate cyclical patterns of convergence and automatically determine the required number of steps. To avoid sudden changes, the moving average fitness of an individual is propagated as gbest. Comparisons for gbest are done within the local neighborhood according to the social comparison theory that individuals who are too different do not compare and influence themselves (Festinger 1954). The range of the neighborhood contracts as the agents' cognitive maps successfully evolve toward their peers.

The degree of coalition is measured as the partition quality (Fisher 1987) obtained with an iterative K-means clustering procedure in the 2D spatial space based on intra-class similarities and inter-class dissimilarities of the groups C_k and proceeds as follows. Let $P(p_i|c_k)$ be the probability of p_i belonging to group c_k :

$$P(p_i|c_k) = 1 - \frac{distance(p_i, c_k)}{\sum_j distance(p_j, c_k)} \tag{4}$$

The overall intra-class similarity is given as $\sum_k \sum_i P(p_i|c_k)$ and inter-class similarity as $\sum_k \sum_i P(c_k|p_i)$. Those probabilities can be combined in an overall measure of partition quality:

$$\sum_k \sum_i P(p_i) P(c_k|p_i) P(p_i|c_k) \tag{5}$$

Using Bayes' formula, Eq. 5 simplifies as follows:

$$\sum_k \sum_i P(p_i) \frac{P(p_i|c_k) P(c_k)}{P(p_i)} P(p_i|c_k) \tag{6}$$

$$\sum_k P(c_k) \sum_i P(p_i|c_k)^2 \tag{7}$$

and where $P(c_k)=\frac{n_k}{N}$ with n_k as the number of elements of c_k and N the total number of elements.

Diversity from the structure of the cognitive maps in a population is evaluated based on the average commonality of each node in the cognitive map of a random agent in each cluster. Let $P(i)$ be the frequency of a node i from a random cognitive map of m nodes in a cluster. The diversity D of cognitive map structures in a population with k clusters and four node types (concept, policy, desirable, undesirable) is computed as follows using the entropy metric:

$$D = -\frac{1}{k} \sum_{i=0}^{m} P(i) log_4 P(i) \tag{8}$$

Experimental Setup

We experimented with a $n \times n$ toroidal grid, a maximum number of cycles C, a number N of agents (particles), a cognitive map of m nodes, p policy nodes, u utility nodes divided randomly between desirable and non-desirable outcomes and PSO parameters described above. Agents interact in their neighborhood range. If no partners were found in their neighborhood, agents moved randomly. If the fitness (Eq. 3) obtained from an interaction was greater or equal than the fitness obtained from a previous interaction (possibly with a different partner), the interaction was deemed successful. If the interaction was successful, the agent moved to an available adjacent free cell closer (measured by the Euclidean distance) to its selected partner and decreased its neighborhood range of interaction. If the interaction was unsuccessful, the agent moved to an available adjacent free cell farther away. Figure 5 shows the clustering obtained by superimposing the final state to the initial state with adaptation. Figure 6 shows the polarization at the corners of the search space without adaptation. The PSO parameters were set according to common values found successful in various implementations (Kennedy & Eberhart 2001). The partner preferences of the agent were randomly initialized and then set to the fitness obtained from an interaction. The absence of convergence in an interaction resulted in a decrease in the partner's preference. Figure 7 shows the coalition quality averaged over 10 runs obtained with the adaptive algorithm in a population of 50 agents from a baseline obtained with no adaptation (t-test p-value of 0.008 in the last cycle). Figure 8 shows the correlated fitness (see Eq. 3) of the agents obtained with adaptation compared with the baseline. Figure 9 illustrates how the average diversity of the cognitive maps in the population drops rapidly with adaptive cognitive agents and then levels off. No significant statistical difference was found when varying the cognitive and social influence parameters, φ_1 and φ_2 respectively. While the experiments are canonical and do not apply to any particular context, the results indicate how the adaptation of cognitive agents enhances group formation over time and is faster than no adaptation.

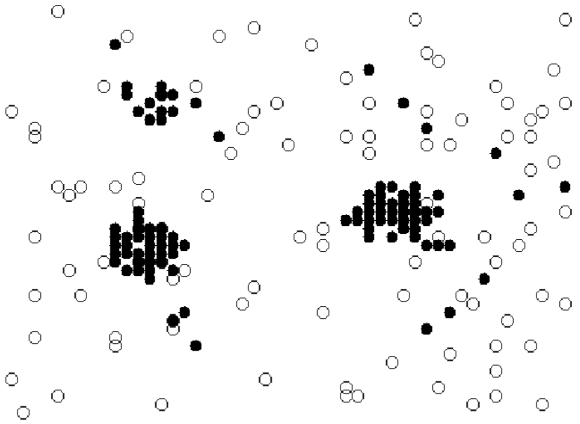

Figure 5: Group formation: the filled circles indicate the final state of the simulation after adaptation, $n = 50, N = 100, C = 300, m = 10, p = 3, u = 3, w = 1.2, \varphi_1 = 2$, $\varphi_2 = 2, vmax = 2$

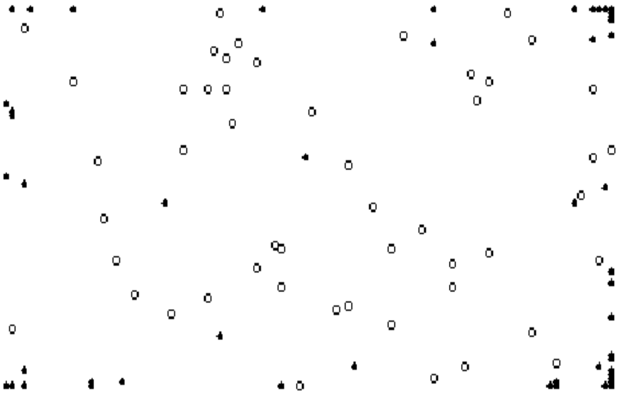

Figure 6: Polarization: the filled circles indicate the final state of the simulation without adaptation, $n = 50, N = 100, C = 300, m = 10, p = 3, u = 3, w = 1.2, \varphi_1 = 2, \varphi_2 = 2, vmax = 2$

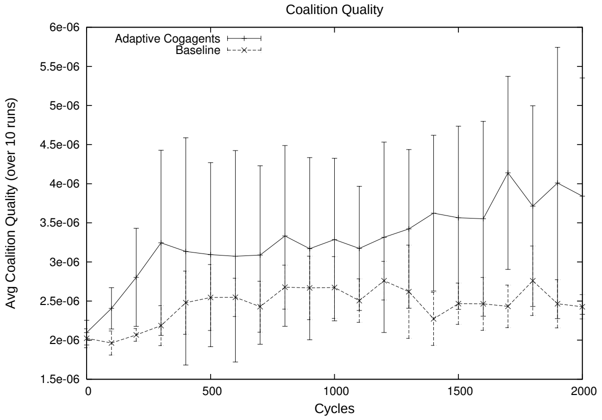

Figure 7: Coalition Quality,$n = 100, N = 50, C = 2000, m = 10, p = 3, u = 3, w = 1.2, \varphi_1 = 2, \varphi_2 = 3, vmax = 2$

Related Work

The goal of Construct (Schreiber, Singh, & Carley 2004) is to evaluate the mechanism of information diffusion in an organization based on how well the agents perform in a canonical classification task. As in Construct, our agents coevolve through peer-to-peer interactions and combined mental models. However, information about other agents is not directly disseminated since the selection of an interaction partner is based solely on subjective preferences. Information about other agents is indirectly propagated through the influence of the neighborhood global best. Our agents are not explicitly situated in an organization but are individuals trying to self-organize through the emulation of peers. It is however possible to initially set the spatial configuration of the agents to test whether the groups are sustainable.

Nexus agents (Duong *et al.* 2007) have a mental model represented as a Boltzmann machine neural network through which they interpret current social situations such as beliefs of trustworthiness, support or blame of other social groups. Based on the theory of cognitive dissonance (Festinger 1957) and narrative paradigm (Fisher 1984), the Boltzmann machine deconflicts those beliefs to indicate final support to certain groups. As in Nexus, our agents interpret current situations through the inference prism of their mental model. Unlike Nexus, there are no direct beliefs characterizing other agents or groups but beliefs on the cause and effect of certain

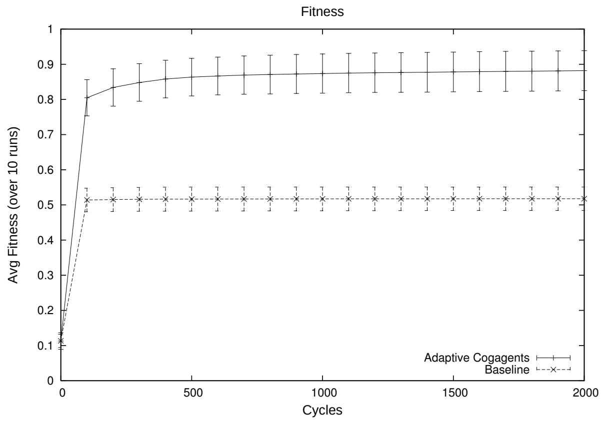

Figure 8: Fitness, $n=100, N=50, C=2000, m=10, p=3, u=3, w=1.2, \varphi_1=2, \varphi_2=3, vmax=2$

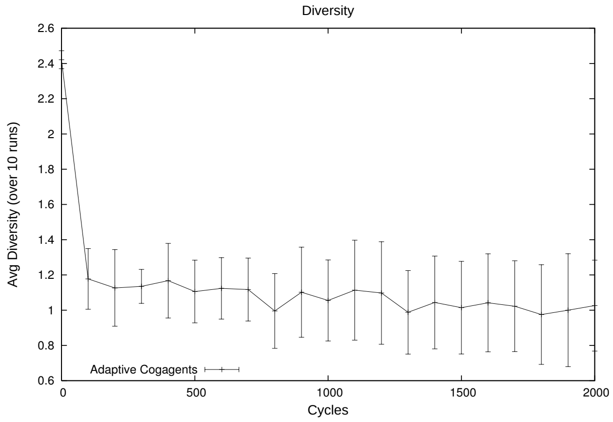

Figure 9: Diversity, $n=100, N=50, C=2000, m=10, p=3, u=3, w=1.2, \varphi_1=2, \varphi_2=3, vmax=2$

concepts that indirectly affect group formation.

PSO has been applied to optimize the weights of a fuzzy cognitive map representing a process control system (E.I. Papageorgiou & Vrahatis 2004). The output is the *best-so-far* cognitive map in the population defined by maximizing the strength of the output concepts within their allowed bounds. Unlike our cognitive agents, there is no differentiation between desirable and undesirable output concepts and no structural difference between agents.

Conclusion

This approach has shown how to incorporate belief changes in group formation and provides insight on what causes groups to emerge or diverge at a more fundamental level than the stated position and influence of key actors. This approach can be used to model the interactions of agents from different cultures represented by a set of prescriptive rules (e.g. proverbs). For example, the impact of a foreign presence in a multi-ethnic society can be modeled, quantified and evaluated over several time cycles based on the interaction of cognitive map agents. While an assumption of this approach is that policy nodes trigger certain known effects, it is useful to quantify the causal beliefs, or political will, that need to be associated for a given policy (set of policy nodes) to be successful. Comparisons between initial and final cognitive map variants can provide structural content insights in addition to predictive trends. Avatars with predefined cognitive maps based on simulation results can be introduced in a virtual world to induce certain desired social effects. Further work remains in grounding structural and functional cognitive map adaptation in a social context.

References

Abramson, M. 2007. Causal Emergence of "Soft" Events. In *AAAI Fall Symposium, Workshop on Emergent Agents and Socialities*.

Axelrod, R.; Nozicka, G. J.; and Shapiro, M. J., eds. 1976. *The Structure of Decision: The Cognitive Maps of Political Elites*. Princeton University Press.

Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Nous* 25(4):487–512.

Collier, N. 2001. Repast: the recursive porous agent simulation toolkit. Retrieved from http://repast.sourceforge.net.

Dickerson, J., and Kosko, B. 1998. *Technology for Multimedia*. New York: IEEE Press. chapter Virtual Worlds as Fuzzy Dynamical Systems.

Duong, D.; Marling, R.; Murphy, L.; Johnson, J.; Ottenberg, M.; Sheldon, B.; and Stephens, S. 2007. Nexus: An intelligent agent model of support between social groups. In *Agent 2007 Conference on Complex Interaction and Social Emergence*. Argonne National Laboratory.

Efird, B. 2008. The Senturion model and its utility in forecasting the future of US/Iranian relations. In *17th Conference on Behavior Representation in Modeling and Simulation (BRIMS)*.

E.I. Papageorgiou, K.E. Parsopoulos, P. G., and Vrahatis, M. 2004. Fuzzy cognitive maps learning through swarm intelligence. In *Fuzzy Systems and their Applications*, volume 3070 of *Lecture Notes in Computer Science*. Springer-Verlag. 344–349.

Festinger, L. 1954. A theory of social comparison process. *Human Relations* 117–140.

Festinger, L. 1957. *A Theory of Cognitive Dissonance*. Stanford University Press.

Fisher, W. 1984. Narration as human communication paradigm: the case of public moral argument. *Communication Monographs* 51:1–22.

Fisher, D. H. 1987. Knowledge acquisition via incremental conceptual clustering. *Machine Learning* 2:139–172.

Kennedy, J., and Eberhart, R. C. 2001. *Swarm Intelligence*. Morgan Kaufmann.

Kohonen, T. 1997. *Self-Organizing Maps*. Springer, 2nd edition.

Kosko, B. 1986. Fuzzy cognitive maps. *International Journal of Man-Machine Studies* 24:65–75.

Kosko, B. 1992. *Neural Networks and Fuzzy Systems.* Prentice Hall, Englewood Cliffs, NJ.

Kosko, B. 1994. *Fuzzy Thinking: The New Science of Fuzzy Logic*. Flamingo.

Latane, B. 1996. Dynamic social impact: The creation of culture by communication. *Journal of Communication* 46(4):13–25.

Liu, H.; Li, B.; Ji, Y.; and Sun, T. 2006. *Applied Soft Computing Technologies: The Challenge of Complexity*. Heidelberg: Springer-Berlin. chapter Particle Swarm Optimization: from lbest to gbest.

Schreiber, C.; Singh, S.; and Carley, K. 2004. Construct - a multi-agent network model for the co-evolution of agents and socio-cultural environments. Technical Report CMU-ISRI-04-109, CASOS - Center for Computational Analysis of Social and Organizational Systems, Carnegie Mellon University.

Tambe, M. 1997. Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7:83–124.

ICCCD 2008

ICCCD 2008
Proceedings

This paper was published in the *Proceedings of the Second International Conference on Computational Cultural Dynamics,* edited by V. S. Subrahmanian and Arie Kruglanski (Menlo Park, California: AAAI Press).

The 2008 ICCCD conference was held at the University of Maryland, College Park, Maryland, USA, 15–16 September, 2008.