Interest Dynamics with Social Constraints and Exogenous Drivers

H. Van Dyke Parunak, Sven A. Brueckner, Elizabeth Downs, Andrew Yinger,

Vector Research Center, Jacobs Engineering Group,
3520 Green Court, Suite 250,
Ann Arbor, MI 48105 USA
{Van.Parunak, Sven.Brueckner, Liz.Downs, Andrew.Yinger}@jacobs.com

Abstract. An important dynamic among interacting knowledge workers is the transition from cognitive convergence (a positive group phenomenon) to collapse (which can lead to overlooking critical information). This paper extends previous studies of this subject in two ways. 1) We place agents in distinct social groups and vary within-group affinity. 2) We provide exogenous drivers of interest. We exhibit a metastable configuration of this system with three phases, and show how to distinguish convergence from collapse. Then we use this metric to explore the system's dynamics, over the space defined by social affinity and precision of queries, and under a range of different functions for the influence that an interaction partner has on an agent.

Keywords: Interest Dynamics, Emergent Behavior, Agent-Based Modeling.

1 Introduction

Humans form interests by interactions with peers (collaboration) and with information sources (search). The pattern of interactions emerges from a person's environment (social network, access to information sources) and current interests. In turn, interactions shape the environment in a complex feedback process. Such dynamics can lead to emergent collective cognitive effects at the system level (across groups of people) that can dominate individuals' interest evolution without a person's being aware of them. One common phenomenon is alignment of interests, sometimes called "consensus formation" [6] or "collective cognitive convergence" [12,13]. This phenomenon contributes to the power of collaboration, but it poses a threat if convergence turns into collapse, blinding the group to new ideas.

Such emergence of global features through feedback loops among autonomous actors in a shared environment is a common feature of stigmergic systems [15]. While stigmergy as a "design pattern" is best known in systems such as social insect colonies, humans also coordinate through simple interactions in a shared environment to accomplish a common goal [11]. A shared social and information environment couples individuals' interest formation. One agent may select an interaction partner (person or information source) based on her current interests. The interaction may change not only her interests, but also her relations, affecting subsequent interaction decisions by others.

We seek to measure cognitive convergence and modulate it by making appropriate changes to a group of knowledge worker using agent-based models. The models to date are driven solely by the initial interests of the agents, and social connections, if represented at all, form a connected graph among all agents. The specific setting that motivates our model requires two extensions. In many government and business settings, a population of *analysts* formulates recommendations for policy makers. While internal discussions are an important part of their work, they also consult exogenous information, in the form of a dynamic collection of *documents*. Furthermore, the population is divided into separate *communities*, within which analysts interact preferentially. Each community starts with a *tasking*, a description of the subject that they are to study. Exploring the dynamics of such a system requires two extensions: interaction of disjoint social groups, and the influence of exogenous information.

Section 2 surveys previous work on interest dynamics. Section 3 outlines our model and formal measures on its behavior. Section 4 reports experiments over varying levels of group affinity and varying query precision. Section 5 concludes.

2 Previous Research On Interest Dynamics

One recent review of computational studies of consensus formation [5] traces relevant studies back more than 50 years, including both analysis and simulation. These studies differ in the belief model and the topology, arity, and preference of agent interactions (Table 1).

An agent's **belief** can be a single variable or a vector, with real, binary, or nominal values. Vector models can be either *independent*, in which an agent can hold any combination of beliefs concurrently, or *correlated*, with consistency constraints beliefs.

Different **topologies** can constrain interactions. Some models constrain interactions by agent location in an incomplete graph, usually a lattice (though one study[8] considers scale-free networks). In others any agents can interact (the "choice" model).

Interaction **arity** can allow agents to interact only two at a time, or as larger groups.

The probability of interaction may be modulated by similarity-based **preference**.

Our work extends this field in two ways. 1) It supports multiple disjoint social networks. 2) It provides exogenous influences, in the form of a collection of documents that agents query. These extensions model groups of agents collectively analyzing information from a changing collection of sources.

Table 1. Representative studies in interest dynamics.

Study	Belief	Topology	Arity	Preference?
Krause [6]	Real scalar	Choice	Many	Yes
Sznajd-Weron [16]	Binary scalar	Lattice	Two	No
Malinchik [9]	Real scalar	Lattice, Random, or Hierarchy	Two	No
Deffuant [4]	Real scalar	Choice	Two	Yes
	Binary vector, independent	Choice	Two	Yes
Axelrod [1]	Nominal vector, independent	Lattice	Two	Yes
Bednar [2]	Nominal vector, correlated	Choice	Many	No
Lakkaraju [7]	Real vector, correlated	Complete, Lattice, Regular, Small-world	Two	No
Parunak [12]	Binary vector	Choice	Many	Yes
This paper	Real vector, independent	Arbitrary, Unconnected	Two	Yes

3 An Agent-Based Model

This section describes our model and the metrics we use to monitor its dynamics. A wide range of configuration parameters are available to configure a scenario (discussed under "model components") and govern its execution (discussed under "model execution").

3.1 Model Components

Our model has five components.

Topic Space.—Analysts and documents live in an abstract Euclidean space constructed from a set of topics. In the real world, a topic is a probability distribution over lexicographic terms (e.g., domain-relevant key words), constructed from a large collection of relevant documents (using, e.g., Latent Dirichlet Allocation[3]).The topic space has dimension equal to the number of topics, with a range of [0, 1] on each axis. A given location is a Topic Model Vector (TMV). A *theme* is a region in topic space. We generate analysts or documents associated with a theme by sampling a Gaussian with configurable mean and variance, resampling when the tails yield a location with a coordinate outside of [0, 1].

Social Network.—Analysts belong to (static) groups whose members are more likely to interact with each other than with members of other groups. This structure models organizational and geographical constraints that externally influence analyst interaction. Additional interaction preferences arise from similarity of analyst interest.

Document.—Our model represents exogenous influences on agent interests in the form of documents. A document is a TMV. Real-world document repositories typically contain documents from different sources and with different concerns. We model this clumping of documents with the notion of a theme, and generate a population of documents by sampling from several themes with specified means (locations in topic space) and variances. We model the arrival of new information during the runtime of the analyst agents as the delayed introduction of documents sampled from a new theme.

Analysts.—An analyst's current interest is also a TMV. A community's tasking is defined as a theme, and we generate a community of analysts working on a given tasking by sampled that theme. The central object of our study is the movement of the analyst's TMV through topic space, relative to the TMVs representing documents and other analysts.

Document Search.—Real-world analysts use search engines to select documents for review. Queries define topics of interest. To model Document Search, an analyst poses a subset of topics. The search weights documents by the strength of their entries on those topics, and probabilistically selects and returns a single document. Noise in this selection process models a real-world analyst's willingness to review results that were not ranked first.

3.2 Model Environment

Fig. 1 shows the documents (small circles) and analysts (large circles with links indicating the social groupings) for two themes and two groups. The rim color of document nodes identifies the theme from which they were sampled, and the rim of analyst nodes shows group membership. The color of the inner document/analyst circle is defined by aggregating the elements of the entity's TMV, and changes as analyst interest evolves.

We dynamically determine the drawing location of documents and analysts through force-based graph layout, a form of multidimensional scaling. This process does not affect the dynamics, which take place in the n dimensional topic space.

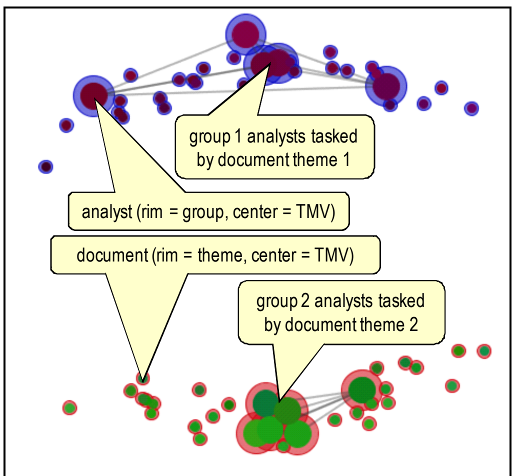

Fig. 1. Visualization

3.3 Model Execution

First we configure a scenario with specified dimensions, analysts, groups and themes. Then analysts repeatedly choose interaction type, assemble interaction options, select interaction target, and execute interaction. Analysts execute randomly with replacement.

Choose Interaction Type.—The analyst chooses probabilistically whether to interact with another analyst or with a document.

Assemble Interaction Options.—If an analyst is interacting with a document, its interaction options (possible targets for interaction in this cycle) are documents currently in the document space. If the agent is interacting with another analyst, its interaction options are other analysts, either in its own group (with probability defined by its Affinity) or in another randomly selected group.

Select Interaction Target.—The analyst selects one target from the options assembled in the previous step. As both document content and analyst interest are represented as TMVs, this step is identical for documents and analysts.

An analyst constructs a query by probabilistically selecting a subset of topics from its TMV, favoring high-interest topics, but adding noise (parameterized by temperature). Then an interaction target is selected based on the query, modeling search execution. The TMVs of the interaction options are sorted by their values in the topics of the query, thus ranking TMVs by their relevance to the given query. Based on that ranking, but again with noise, we select one TMV. Here the noise temperature models an analyst selecting one result.

Execute Interaction.—The analyst now updates its interest model. If the analyst interacts with another analyst, it samples the Learning Style parameter to determine the personality it should assume in this interaction. The selected personality sets the update rule for updating each topic's interest level as a function of the difference in interest on that topic between the agent and the selected interaction target. Section 0 explores this rule. In the standard personality, the agent shifts its interest level in updated topics toward the interest level in the interaction target. In the curmudgeon personality, it shifts away from the other. If the interaction target is a document, then the agent always uses the standard personality.

3.4 Performance Metrics

We define a set of component metrics, and a single aggregate metric.

Component Measures.—The topics in a given model span a high-dimensional metric space with TMVs limited to [0, 1] for each topic. As analysts update their TMVs through interactions, they move through this space. We measure aspects of individual analyst movement to detect dynamic characteristics that indicate cognitive collapse.

The most fundamental measure is the magnitude of a single TMV update, which is the length of the vector between the agent's prior and new location in each cycle.

The length of a step in topic space conveys the **absolute** magnitude of the impact a particular interaction had on the analyst's interest. It does not show the nature of the step **relative** to the other analysts. A second measure is the distance of the analyst's location (after the step) to the mean over the TMVs of all analysts regardless of group affiliation. The mean TMV may not be near any analyst. Movement of analysts shifts the location of the mean TMV, so successive values of this measure, unlike "step-length," are not statistically independent.

An Aggregate Measure.—Initial explorations based on these measures show that we also need to discover a directed walk, in which an agent's successive steps are correlated with one another. In previous work [14], we applied information-theoretic (entropy) measures to detect a directed walk, but encountered idiosyncrasies from the specific definition of the system states whose probabilities are measured in the entropy calculation. For the current research, we developed an aggregate metric that measures the "directedness" in an agent's movement through topic space without the complications of the entropy calculations. The delayed step length metric adds the step vectors (delta TMV) for a single agent over the most recent n cycles (configurable, 50 in the results reported here). The vector sum of steps of a **random walk** is on the order of $\sqrt{n}$, while the vector sum of steps that generally point in the same direction (**directed walk**) tends to be on the order of n.

4 Experiments with the Model

This section walks through an example scenario, exhibits the system's metastability, derives an objective way to measure the cognitive collapse of a knowledge community, and uses this measure to explore the space defined by community affiliation and interaction with exogenous information. Our experiments explore only a small portion of the space defined by these dimensions, but still suggest two practical principles for managing convergence and preventing collapse among knowledge workers.

4.1 A Representative Scenario

We consider a small scenario with two distinct document themes and two groups of analysts. We sample 25 documents for each theme. One theme is the tasking (sample initial interest vector) for all 6 analysts of the first group, and the other theme initializes the 5 analysts of the second group. The documents and analysts are embedded in a 10-dimensional topic space. The two document themes (and thus the two document subsets and the two analyst groups) are distinctly different, but somewhat overlapping in two topics. Fig. 2 (frame 1) shows the initial layout. Each group is associated with a subset of the documents. Analysts in the larger (bottom) group have higher affinities (preference for their own group) than in the smaller group.

As the simulation runs, we visualize the recent interactions of agents with other agents (red lines) and documents (blue lines) fading away into history (line transparency). The screen shots in Fig. 2 (frames 2-5)show the interactions in the four

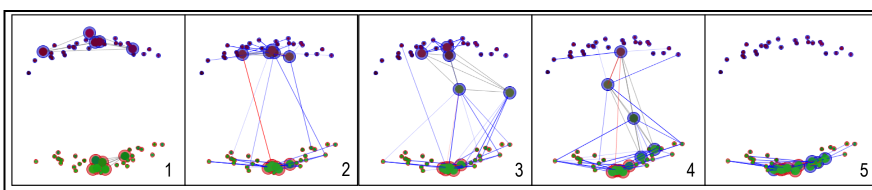

Fig. 2. Stages in model evolution.

most recent cycles.

In (2), the agents of each group that were initially spread out in their respective tasking theme converge on their common interest and thus form tighter clusters in their respective group. While there are also cross-theme/group interactions (blue/red lines crossing the gap in the center of the view), most interactions occur within the tightly clustered groups and their surrounding theme.

In (3), interactions of low-affinity analysts with the other group eventually lead to the defection of two analysts from the interest pattern of the smaller group and their transition towards the larger group.

In (4), once the first analyst defects from the smaller, low-affinity group, others follow rapidly. Eventually, all analysts abandon their interest in the upper document theme.

In (5), both groups of analysts have converged on interests exemplified by the document set in the lower part of the screen. Interactions among analysts within a group are no different from Out-of-Group interactions. Interactions with documents are (mostly) confined to the theme on which all analysts converged.

4.2 A Metastable Transition

Our metrics reveal three phases of interest evolution (Fig. 3, metrics applied to a single analyst from the upper, smaller group). The agents from this group first remain in their separate interest area (Phase 1), eventually defect one-by-one to the other group (Phase 2), and then explore the other interest area (Phase 3).

In **Phase 1**, the agent is far from the mean TMV, reflecting the initial separation of analysts' interests. Short steps (most frequent) are interactions with other analysts from the same group (and thus similar interest) or documents near the analyst's initial tasking. Medium length steps are interactions with documents from the other theme. We set the model parameters so that document interactions have less impact on analyst interest than analyst-to-analyst interactions. The longest steps, and the least frequent, are interactions with analysts from the other group. The roughly constant distance to the mean TMV shows that the agent's successive steps are not correlated.

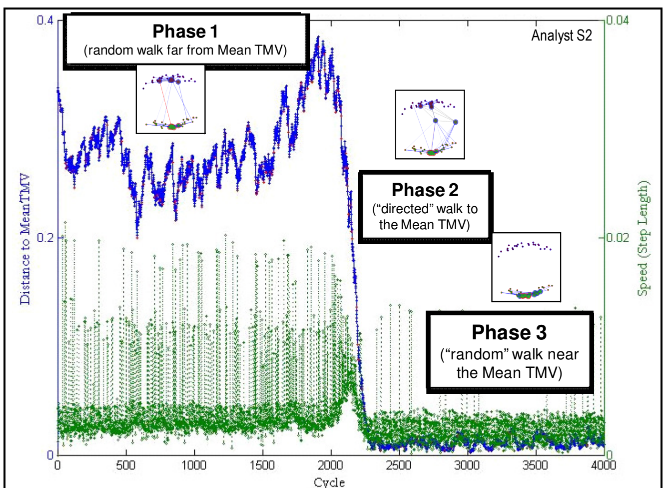

Fig. 3. Dynamics of one analyst from top group.

In **Phase 2**, the agent defects from

the region of its original tasking to the region occupied by the larger group. The distance to the mean TMV rapidly shrinks. High-frequency interactions increase in step length as the agent moves away from its own group, due to the TMV update rule, which computes larger changes for larger differences between the analyst's TMV and the TMV of its interaction partner. Similarly, lower-frequency steps that correspond to interactions with the other document set and out-of-group analysts decrease in length, as the analyst moves closer to those entities. In this phase, the agent's successive steps are correlated, as the rapidly falling distance to the mean TMV shows.

In **Phase 3**, as in Phase 1, successive steps are again uncorrelated. All analysts occupy the same region in topic space and share a common interest in documents of the second document theme. The analysts' distance to the mean TMV is small as they are now all tightly clustered. The analysts still interact (infrequently) with documents from the other theme (larger steps), but those interactions have no lasting effect on their relative locations.

Fig. 4 compares "distance to the mean TMV" (blue) for a single analyst with "delayed step length" (green). Phase 2 is indeed characterized by a directed walk while Phases 1 and 3 are (generally) less directed.

4.3 Defining and Measuring Collapse

Informally, "cognitive collapse" is the inability of an agent or a group of agents to respond to new information. We can now operationalize this definition: An agent is cognitively collapsed if it is not in Phase 2 dynamics and if it does not return to Phase 2 dynamics when offered qualitatively new information.

We introduce *qualitatively new information* by adding documents to the model at a time when all agents show Phase 3 dynamics. This new document theme (located between the two initial themes in topic space) probes the analyst dynamics. If analysts in Phase 3 are converged but not collapsed, they should return to Phase 2, as indicated by delayed step length.

We automate this detection with the nonparametric Mann-Whitney U test [10] to compare step lengths before and after probe insertion. The baseline configuration has a query temperature of 0.12 and an analyst affinity of 0.4. We explore parameter space from that point up to a query temperature of 0.32 and an affinity of 0.9,

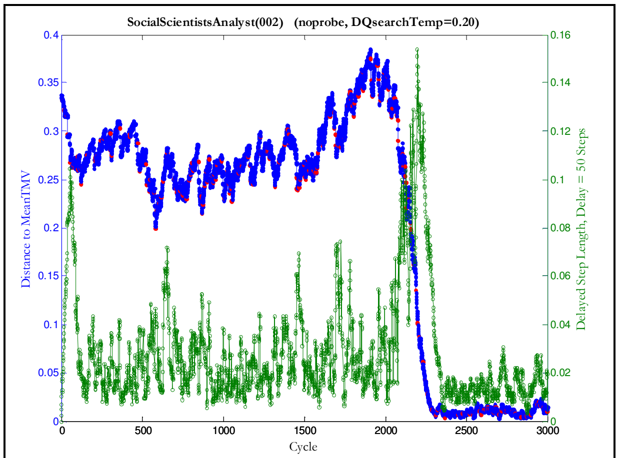

Fig. 4. Delayed step length marks Phase 2.

and run 20 replications at each point. When each configuration reaches Phase 3, we insert a probe, and compute the percentage of analysts who respond (and so are not collapsed). Increasing the query temperature, thus exposing analysts to unexpected documents, dramatically reduces collapse. Surprisingly, the probability of collapse does not vary systematically as we change analysts' affinity for their own group.

4.4 Exploring the Update Rule

Our TMV update rule translates the difference in TMV elements between an analyst and an interaction partner into the length of the analyst's step (toward the partner for an ordinary analyst, and away for a curmudgeon). In all experiments thus far, the magnitude of the change in the interest in a particular topic in the TMV is proportional to the difference in that topic between the agent and the interaction target. This assumption reflects **curiosity**, assigning more impact to an interest very different from mine. An alternative model is **homophily**: I am more likely to move toward ideas that are close to my own. In real-world analysts, the correct model is likely to be a **mixture** of these two effects: interests too far from mine are threatening, and interests too close to mine are boring, so my response will be greatest somewhere in the middle.

Fig. 5 summarizes these options. The parameter s is the hypothesized difference that will lead to maximum movement. When $s = 0$, we recover homophily, while $s = 1$ yields curiosity.

To explore the effect of the update rule, we focus on two observables in our model:

- The convergence of the two groups;
- The response to a probe.

We instantiate a large number of systems for each value of s and measure the probabilities $p(a)$ of convergence and $p(b)$ that the system responds to a probe (i.e., is not collapsed). Fig. 6 shows the results for a low affinity and high query temperature. We again plot the average over 20 random seeds for each data point. Low s (homophily) yields no interest convergence between the groups but decreasing probability of collapse (increasing probability of non-collapse). Interestingly, group convergence has a critical threshold. Thus, instead of a gradual rise of the convergence probability, the plot remains at 0% until a critical value of s is reached. Then the probability of all analysts converging on the same interest region rapidly

Cognitive Rationale	$g(x = \lvert d_j - a_j \rvert)$	Sketch
Curiosity	x	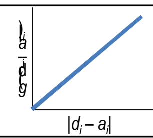
Homophily	$1 - (1 - x)$	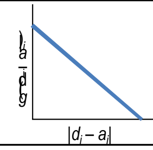
Mixture	$s \in [0, 1]$ = "sweet spot" parameter $x \leq s$: x/s $x > s$: $(1 - x)/(1 - s)$	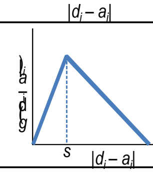

Fig. 5. Alternative TMV update rules g

increases, but the analysts remain equally open to new information (probability of collapse remains low). Finally, as curiosity dominates (large), collapse becomes more common. The increase in risk of cognitive collapse towards the end of the "sweet spot" sweep suggests another practical lesson for real-world knowledge workers: a mixed learning strategy that is most sensitive to information that is neither completely novel nor entirely familiar is less vulnerable to collapse than either extreme.

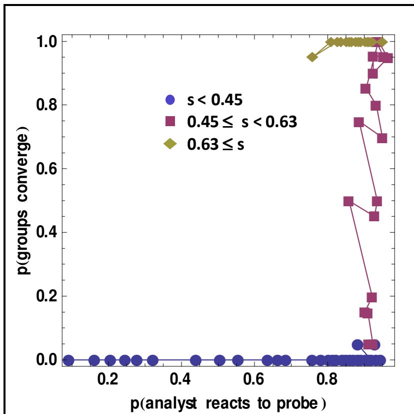

Fig. 6. Sweep over *s* for configuration with high search noise but low affinity.

5 Conclusion

The interest dynamics of multiple interacting knowledge workers are complex, often counter-intuitive, and yet critical for much collaborative work in the modern world, and enjoy the attention of a significant research community. Previous simulation studies focus on the evolution of an initial distribution of interests across agents. While suggestive, such studies do not account for two critical features of knowledge workers in the real world.

- Their *social environment* is highly clustered, and they are more likely to interact with another agent in their cluster than with an agent in another cluster.
- Their *information environment* includes exogenous knowledge sources ("documents") in addition to other agents, and they seek out these documents with a query process.

Our model implements both of these features. The resulting system exhibits a metastability that allows us to formulate an operational measure of cognitive collapse. A preliminary exploration of the parameter space of social affinity and query precision with this measure yields two (very provisional) practical lessons.

First, query precision has much more influence on collapse than does social affinity. To modulate the convergence of a community of analysts, managing the amount of noise added to their queries is a more promising method than changing their group membership.

Second, motivating analysts to prefer interests that are neither completely new nor completely familiar will lead to more robust convergence without collapse than the extremes.

Acknowledgment

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via US Army Research Office contract number W911NF-10-C-0108. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, ARO, or the U.S. Government.

References

1. Axelrod, R.: The Dissemination of Culture: A Model with Local Convergence and Global Polarization. Journal of Conflict Resolution, 41(2 (April)):203-226 (1997)
2. Bednar, J., Page, S., Bramson, A., Jones-Rooy, A.: Conformity, Consistency, and Cultural Heterogeneity. In Proceedings of Annual Meeting of the American Political Science Association, American Political Science Association (2006)
3. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet Allocation. Journal of Machine Learning Research, 3:993-1022 (2003)
4. Deffuant, G., Neau, D., Amblard, F., Weisbuch, G.: Mixing beliefs among interacting agents. Advances in Complex Systems, 3(1-4):87-98 (2000)
5. Hegselmann, R., Krause, U.: Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. Journal of Artifical Societies and Social Simulation, 5(3):none (2002)
6. Krause, U.: A Discrete Nonlinear and Non-autonomous Model of Consensus Formation. In Elaydi, S., Ladas, G., Popenda, J., Rakowski, J. (eds.) Communications in Difference Equations, Gordon and Breach Science Publ., 2000)
7. Lakkaraju, K., Speed, A.: A Cognitive-Consistency Based Model of Population Wide Attitude Change. 2010 AAAI Fall Symposium Series, 2010)
8. Long, G., Yun-feng, C., Xu, C.: The evolution of opinions on scale-free networks. Frontiers of Physics in China, 1(4):506-509 (2006)
9. Malinchik, S., Rosenbluth, D.: Paradoxical Dynamics of Population Opinion in Response to Influence of Moderate Leader. IEEE Symposium Series on Computational Intelligence (SSCI 2011), IEEE, Paris, France (2011)
10. Mosteller, F., Rourke, R.E.K.: Sturdy Statistics: Nonparametrics and Order Statistics. Reading, MA, Addison-Wesley (1973)
11. Parunak, H.V.D.: A Survey of Environments and Mechanisms for Human-Human Stigmergy. In Weyns, D., Michel, F., Parunak, H.V.D. (eds.) Proceedings of E4MAS 2005, vol. LNAI 3830, Lecture Notes on AI, pp. 163-186. Springer, 2006)
12. Parunak, H.V.D.: A Mathematical Analysis of Collective Cognitive Convergence. In Proceedings of the Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS09), pp. 473-480, (2009)
13. Parunak, H.V.D., Belding, T.C., Hilscher, R., Brueckner, S.: Understanding Collective Cognitive Convergence. In David, N., Sichman, J.S. (eds.) Multi-Agent-Based Simulation IX: International Workshop, MABS 2008, Estoril, Portugal, May 2008, Selected Papers, vol. LNAI 5269, pp. 46-59. Springer, Berlin (2009)
14. Parunak, H.V.D., Brueckner, S.: Entropy and Self-Organization in Multi-Agent Systems. In Proceedings of The Fifth International Conference on Autonomous Agents (Agents 2001), pp. 124-130, ACM (2001)
15. Parunak, H.V.D., Brueckner, S.A.: Engineering Swarming Systems. In Bergenti, F., Gleizes, M.-P., Zambonelli, F. (eds.) Methodologies and Software Engineering for Agent Systems, pp. 341-376. Kluwer, 2004)
16. Sznajd-Weron, K., Sznajd, J.O.: Opinion Evolution in Closed Community. International Journal of Modern Physics C, 11(6):1157-1165 (2000)