SAND2013-9802C

# A comparison of two methods to reduce diffusion time in a multi-agent, multi-attitude model

Kiran Lakkaraju

Sandia National Laboratories*
P.O. Box 5800, MS 1327, Albuquerque, NM 87185
klakkar@sandia.gov

**Abstract.** In this work we study the effects of agenda setting on attitude diffusion using the "Multi-Agent, Multi-Attitude Framework" (MAMA). MAMA captures the interaction between attitudes (through cognitive consistency effects) and interpersonal interaction. Agenda setting is when media's focus on certain stories increases their importance in the minds of the viewers. Using the MAMA model, we study the impact of Agenda Setting on time to diffusion. We show that agenda setting can significantly decrease diffusion time for a variety of network topologies. Secondly, we show that agenda setting plus strategic choice of seed nodes provides the fastest time to convergence.

## 1 Introduction

Attitudes have long been known to have an impact on individual behavior. Thus, many computational models of attitude diffusion have been proposed. However, it is also clear from the social psychology research that attitudes are not independent – sets of attitudes interact through, for instance, the drive for consistency among attitudes.

The first contribution of this work is to develop a multi-agent system that models attitudes, the interaction between attitudes, and interpersonal influence. We call this the "Multi-Agent, Multi-Attitude Framework" (MAMA). This framework captures the important concept of "cognitive consistency", which has a large effect on the dynamics of attitude change and attitude interaction.

Cognitive consistency is a hypothesized drive for individuals to have attitudes that are "consistent" with each other. Cognitive consistency has long been shown to be an important factor in attitude change [1, 2]. For instance, according to these theories, an individual holding a strong positive attitude towards environmentalism should also hold a strong positive attitude towards recycling; if they do not, the attitudes are inconsistent with each other and could cause an uncomfortable feeling (i.e. *cognitive dissonance*) which tends to result in either attitude or behavior change [3].

* Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Another aspect to consider in attitude diffusion is the impact of the media. Agenda setting is when the media's focus on certain stories increases their importance in the minds of the viewers [4, 5]. Agenda setting and interpersonal communication interact to influence individuals. In the early "two-step" model, media was thought to influence "opinion leaders" who then interacted and influenced others [6]. While the two-step model has lost support, the general idea of media influencing the agenda, which in turn spurs discussion and interpersonal influence, has some support [4].

Agenda setting becomes more important when considering interacting attitudes – how does the order of discussion topics influence attitude diffusion? Using the MAMA model, we explore this question through extensive simulations. We show that agenda setting can significantly decrease the time to convergence in the MAMA model.

A common problem is to identify "seed" nodes that initially have the attitude that should be diffused [7]. We show that agenda setting plus strategic choice of seed nodes provides the fastest time to convergence. Finally, the simulations show that for a scale-free network, seed node choice is still the most important factor in determine time to convergence – far outweighing the impact of agenda setting aspect. This highlights the important role of interpersonal influence on diffusion.

## 2 MAMA Model

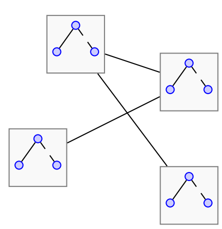

**Fig. 1.** Social network. See text for details.

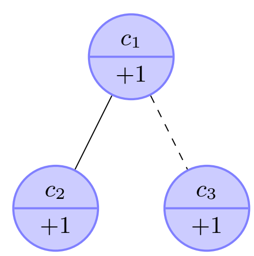


**Fig. 2.** Cognitive network. See text for details.

To explore the impact of multiple attitudes we have developed a multi-level agent based model that contains two levels, a *social* level – which captures interpersonal interaction between agents – and a *cognitive* level that captures the interactions of attitudes *within* an agent.

Let $G_s =< V_s, E_s >$ be a undirected graph that represents the social level of the model. Let $a_i \in V_s$ be the set of agents, and $(a_i, a_j) \in E_s$ represent a bidirectional influencing relationship between agents $i$ and $j$. Figure 1 depicts an example social network, where each rectangle is an agent

Each agent has a cognitive network associated with it. A cognitive network is a weighted undirected graph, $G_c =< V_c, E_c >$ that represents cognitions and the interactions between them. We use the term cognitions to refer to any entity towards which an individual can have an attitude, such as people, places and things; but also to more abstruse entities like values.

Let $c_k \in V_c$ be the set of cognitions, and $(c_k, c_q) \in E_c$ represents a bidirectional influencing relationship between cognitions $k$ and $q$. $w(k, q)$ is the weight of edge $(c_k, c_q)$; the weight can either be $+1$, or $-1$: $w(k, l) \in \{1, -1\}$. The weight represents the relationship between cognitions, as we describe later on. For convenience, we let $n_c = |V_c|$.

An attitude towards a cognition is represented as a real number, called the *value* of the cognition, between $-1$ and $+1$. The sign of the value represents the *valence* of the attitude; positive values indicate positive attitudes and negative values represent negative attitudes. The size of the value represents how strongly the individuals holds the attitude. So a value of 0.5 would be a mildly positive attitude, whereas a value of $-1.0$ is a very strong positive attitude. In this work, we limit the values to be either $-1$ or $+1$. Let $v(i, k)$ be the value of cognition $k$ of agent $i$.

Figure 2 depicts a cognitive network. The lines represent relationships between cognitions; dashed lines are negative relationships, solid lines are positive relationships. The bottom of each cognition contains the currently assigned value.

Let $\chi_i(k, q)$ be the *consistency* of an edge $(c_k, c_q)$ in the cognitive network of agent $a_i$. The value of $\chi_i(k, q)$ is:

$$\chi_i(k,l) = \begin{cases} 1 & \text{if } w(k,q)v(i,k)v(i,q) > 0, \\ 0 & \text{Otherwise} \end{cases} \tag{1}$$

Intuitively, if an edge has a negative weight, the edge is consistent if the two cognitions have values with differing signs. If an edge has a positive weight, the edge is consistent if the two cognitions have values with the *same* sign.

Let the *state* of a cognitive network be an assignment of values to its cognitions. There are $m = 2^{n_c}$ states for a cognitive network, labelled: $s_1 \dots s_m$. $s_p(k)$ is the value of cognition $k$ in state $p$.

The consistency of a cognition $k$ for agent $i$ is:

$$\phi_i(k) = \frac{\sum_{(c_k,c_q)\in E_c} \chi_i(k,q)}{l_i} \tag{2}$$

where $l_i$ is the number of edges incident to concept $i$. Intuitively, consistency increases as a cognition has more edges that are consistent.

Cognitive networks can be viewed as *bi-valued, binary constraint satisfaction network* [8]. A significant body of work has been developed around binary constraint satisfaction. The problem there is finding the correct solution; our problem is understanding when a solution diffuses across a network.

### 2.1 Attitude change

In our model, attitude change is initiated by interpersonal interaction, but mediated by the state of the cognitive network. Given no cognitive influence, the baseline probability that an agents changes their attitude is indicated by $P_{base}$. In this work we assume that $P_{base}$ is the same for all agents. In the following we describe how the state of the cognitive network modifies $P_{base}$.

**Drive for cognitive consistency** In Figure 2, concept $c_1$ is in an inconsistent state with concept $c_3$ – the link between them is negative, so they should be opposite, however there are positive attitudes towards both concepts. Based on cognitive consistency theory, we should expect $c_1$ to be more likely to change[1]. On the other hand, concepts that are highly consistent should be less likely to change – since they are consistent with most of their neighbors.

We represent this drive to consistency as a multiplicative weight on the baseline probability.

Let $f_{con}(k,i)$ represent the *inclination* to change cognition $k$ of agent $i$ based on its consistency with other cognitions. Intuitively, the more consistent the concept is with its neighbors, the less likely it is to change.

We define $f_{con}(k,i)$ as a sigmoid curve:

$$f_{con}(k,i) = \epsilon + \frac{2}{1 + e^{-10((1-\phi_i(k))-.5)})} \tag{3}$$

For cognitions that have more than 50% of their neighbors in an inconsistent state, $f_{con}(k) > 1.0$, thus increasing the probability they will change (with a maximum multiplicative increase of 2). For those with less than 50% of their neighbors in an inconsistent state, $f_{con}(k) < 1.0$, decreasing the probability to change (with a minimum of $\epsilon$).

For example, consider Figure 2; $f_{con}(1) = 0.5$ because concept $c_3$ is inconsistent with $c_1$ but consistent with $c_2$.

On the other hand, $f_{con}(2) = \epsilon$, since all of its neighbors are consistent with it. Finally, $f_{con}(3) \approx 2.0$, since all of its neighbors are inconsistent with it.

**Embeddedness** The *embeddedness* of a cognition refers to how well it is connected to other cognitions in the cognitive network. Embeddedness is related to a resistance to change (see [9, Chap. 12] for a review).

Consider an individual who initially has a negative attitude towards environmentalism. If the individual were to suddenly have a positive attitude, there would be serious dissonance with past decisions and their current state; they may be dissatisfied with their car, worried about how they consume energy, etc. Thus, more connected cognitions are resistant to change.

[1] There are a host of other factors that play a role, such as the type of attitude (implicit or explicit), the persuasion route, etc. See [9, 10]. In this model we focus on the core concepts of cognitive consistency and leave further refinement to future work

In Figure 2, concept $c_1$ is connected to two other concepts vs. $c_2$, which is only connected to 1 other concept. Thus, we would expect $c_1$ to have more resistance to change than $c_2$ or $c_3$.

We represent this resistance to change as a multiplicative weight on $P_{base}$.

Let $f_{deg}(k, i)$ be the *resistance* to change cognition $k$ of agent $i$ based on the cognitions embeddedness, which we measure through its degree ($deg(k)$). Intuitively, we want $f_{deg}(k, i)$ to decrease as we increase the degree of the concept.

$$f_{deg}(k, i) = \begin{cases} 1.0 & \text{if } deg(k) < deg_{max,i}/2, \\ 0.5 & \text{else} \end{cases} \tag{4}$$

where $deg_{max,i}$ is the highest degree in the cognitive network of agent $i$.

For example, Figure 2; $f_{deg}(1) = 0.5$, since cognition $c_1$ has the highest degree. $f_{deg}(2) = f_{deg}(3) = 1.0$, since they only have one neighbor.

**Probability of Change** Bringing everything together, let $P_{change}(k, i)$ be the probability of cognition $k$ of agent $i$ changing value, given that $i$ is interacting with another agent with the opposite value for cognition $k$. Then:

$$P_{change}(k, i) = P_{baseline} \cdot f_{degree}(k, i) \cdot f_{con}(k, i) \tag{5}$$

### 2.2 Model Dynamics

Algorithm 1 specifies the dynamics of the model. Similar to other work ([7]) we study the progressive case. Since we have multiple concepts in our model, we designate a single state $s^*$ as the *goal* state. Once a cognitions switches to the value in the progressive state, it cannot switch back.

**Algorithm 1:** Model Dynamics

**for** $t \leftarrow 1$ **to** $t_{max}$ **do**
  Choose a strategy $\pi$
  **for** $t_s \leftarrow 1$ **to** $N$ **do**
    Choose a random agent $a_i$
    Choose $a_j$ a random neighbor of $a_i$
    Choose a *topic* cognition $\tau$ according to $\pi$.
    **if** $v(a_i, \tau) \neq s^*(\tau)$ **and** $v(a_i, \tau) \neq v(a_j, \tau)$ **then**
      Set $v(a_i, \tau) = v(a_j, \tau)$ with probability $P_{change}(\tau)$
    **end**
  **end**
**end**

Each iteration of $t$ is called a single *timestep*, and within each timestep we randomly and with replacement, sample $N = |V_s|$ agents. Thus on average, every agent is chosen once per timestep.

Note that these dynamics correspond closely to a voter model [11], except the probability of switching varies over the length of the simulation. Voter models have a long history as a simple tool to study diffusion (see Section 4.

### 2.3 Agenda Setting

In the classic study on agenda setting individuals were asked to indicate issues they thought were important, which then were compared to news coverage of those issues. It was found that what people found important was what was covered most [5].

We assume that issues are represented by cognitions, and that increasing issue importance leads to more discussion. Thus we define agenda setting in terms of the choice of topic cognition. We define an *agenda* $\pi = [P(c_1), \ldots, P(c_{n_c})]$ as a probability distribution over cognitions. The agenda defines which cognition is chosen in 1.

A *Time-Independent Agenda* (TIA) is a fixed probability distribution over the cognitions. A special case is the uniform distribution, where each cognition has a probability of $1/n_c$ of being chosen.

A *Time Varying Agenda* (TVA) is an agenda that changes over time. Essentially, it is some number of agendas which are active at certain times. For instance, for the cognitive network from Figure 2, we can define a time varying agenda by specifying multiple agendas that span the timestep range from $(0, \infty)$. From timestep 0 to timestep 1000, the agenda may be $\pi_{1000} = [1/3, 1/3, 1/3]$, but from 1000 onwards, the agenda may be: $\pi_\infty = [1/9, 1/9, 7/9]$.

## 3 Experiments

Our goal is to study how agenda setting vs. node choice influences the propagation of attitudes in a population. Thus, the metric we will use is the *mean diffusion time* – the number of timesteps the system takes for 90% of the population to reach the goal state, averaged over some number of runs (we also call this time to convergence).

We assume that all agents have the same type of cognitive network, depicted in Figure 2, and only vary in their initial state. We call this the "3-Fan" network, because it has one central cognition ($c_1$) and two ancillary cognitions that connect to it ($c_2, c_3$). More generally, a "$n$-Fan" network would have 1 central cognition and $n-1$ ancillary nodes connected to it.

The fan network, while simple, represents the interaction between attitudes. Consider the central cognition as an attitude towards a specific decision, such as purchasing an energy efficient lightbulb. The ancillary nodes represent other cognitions that can impact this decision, say political preference [12]. The bidirectional link between the central node and the ancillary nodes illustrate bidirectionality of influence – your attitudes can influence your behavior, but your behavior can influence your attitudes.

The goal state is set to $s^* =< +1, +1, -1 >$. Initially, a random 10% of the population is assigned the goal state. The rest of the agents are assigned the state of $s^{'} =< -1, -1, +1 >$. Note that these are the only two fully consistent states in the 3-Fan network.

We use two social networks. The first is a *k-regular graph* – a graph where each vertex has degree $k$.

The second social network is a small-world network, which is a network that features a high clustering coefficient and short average path length.

Three social networks were studied:

**$k$-regular Graph** A graph in which each vertex has degree $k$. We use a network with 1000 agents, and $k = 4$.

**Small World Graph** A graph which features a high clustering coefficient and short average path length. Small world networks appear in many real-world domains [13]. We used the algorithm defined in [13], implemented in [14]. We use a network with 1000 agents, and following [13], we set the initial number of neighbors to 10, and the rewiring probability to 0.01. This produces a small world network with clustering coefficient of 0.668519 and average path length of 6.289300.

**Scale Free Graph** A scale free network has a degree distribution described by $P(k)\ k^{-\alpha}$. We used the software tool SNAP to generate a scale free graph by the Barabasi-Albert preferential attachment method, with $n = 2000$ and $m = 3$ (see [15, 16] for details on the algorithm)

**Facebook Circles** We used the Facebook ego dataset collected by [17] which contains $n = 4039$ vertices and 88234 edges. This data set contains the (anonymized) ego networks of participants of a Facebook Survey. See [17] for more details.

We are interested in two conditions, whether agenda setting is on or off, and how the seeds are chosen.

When agenda setting was on, we used a time varying agenda. Let $\pi_b^k$ be an agenda that was used from timestep 0 to timestep $b$ which sets cognition $k$ to $p = .9$ and the other two cognitions to $(1 - p)/2$. After timestep $b$, we set the agenda to $\pi = [1/3, 1/3, 1/3]$. Previous work has shown that a time varying agenda with a high $p$ for cognitions 2 or 3 performs better than a time-independent agenda [18].

Without agenda setting, we assume that each cognition can be a topic with equal probability. This corresponds to an agenda of $[1/3, 1/3, 1/3]$.

The seed nodes were set according to four different methods:

**Random** Seeds were selected uniformly randomly from the vertices of the graph.

**Degree** Vertices were ordered by degree (from highest to lowest) and the top 10% were chosen to be seeds.

**Betweenness** Vertices were ordered by betweenness (from highest to lowest) and the top 10% were chosen to be seeds. A vertex with high betweenness is one that appears in many of the geodesics connecting nodes in the graph [19].

**Eigenvector Centrality** Vertices were ordered by eigenvector centrality (Google's PageRank measure is a variant of this centrality measure) and the top 10% were chosen to be seeds.

Previous work has shown that the boundary value $b$ has a significant impact on time to convergence, and thus we varied $b$ from 0 to 1500

### 3.1 Experiment 1: Impact of Agenda Setting

In this experiment we study the impact of time-varying agendas on mean time to convergence. Previous work [18] has shown that choosing a strategy with $k = 2$ or $k = 3$ and a high p ($\approx .90$) significantly reduces mean time to convergence. Thus, in the following experiments we set $k = 2$ and $p = .90$.

Figure 3 shows the results for this experiment.

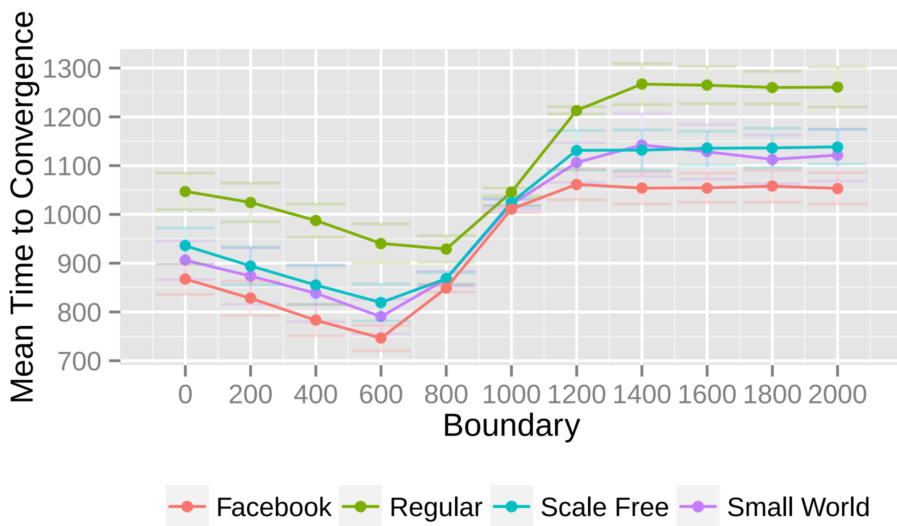


**Fig. 3.** Mean time to convergence for all graph types

At $b = 0$, the time varying strategy $\pi_b^2$ equals the strategy $[1/3, 1/3, 1/3]$. We consider this the "Agenda Off" case. As we increase the boundary value $b$, the impact of the non-uniform agenda is stronger. We can see that initially the non-uniform strategy helps, but eventually the over-focus on it actually increases the time to convergence (for a more thorough explanation, see [18]).

### 3.2 Experiment 2: Strategic seed choice vs. Agenda Setting

Figure 4 shows the mean convergence time (over 100 runs) for different choice of seeds and for varying boundary values.

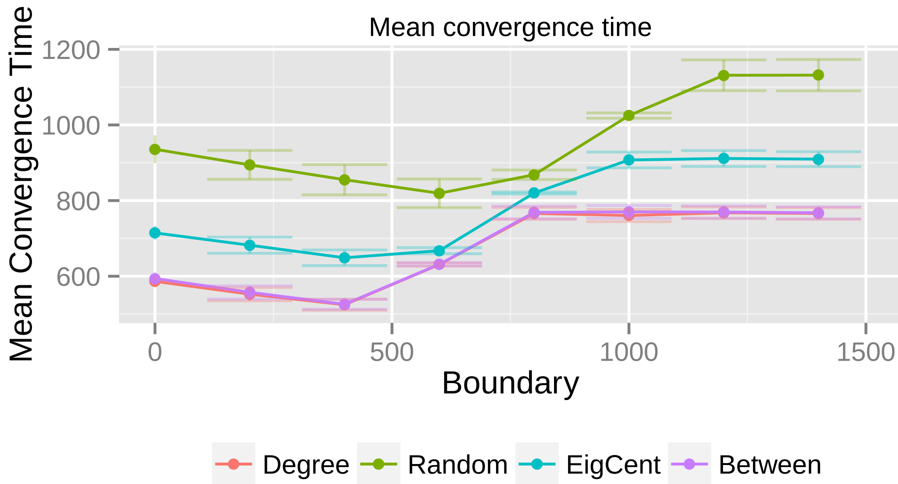


**Fig. 4.** Mean time to convergence +/- 1 standard deviation. Betweenness and Degree are overlapping. For the Scale Free network.

In the "Agenda off" situation, as expected, we see that the random setting performs the worst by more than 200 timesteps, than any of the other seed setting options. This matches expectations; existing work in the influence maximization setting indicates that strategically choosing the seed nodes can increase the influence spread under a linear cascade model [20, 7].

As we increase the boundary value, the influence of the non-uniform strategy become more prevalent. The minimum mean time of convergence occurs when the agenda setting is on, and the seed nodes are set (the results for Between and Degree almost entirely overlap).

These results indicate the the best performing option is to include agenda setting and strategic node choice. For the scale free network, choosing the degree or betweenness measure provides the best results.

Figures 5 and 6 show the mean convergence time for the Facebook network. Figure 5 shows all the node choice metrics. Surprisingly, the eigenvector centrality measure performs very poorly.

Figure 6 only shows the Degree, Random and Betweenness node choice metrics. Interestingly, the random node choice performs better than the degree based node choice. This is probably due to the topology of the Facebook graph.

## 4 Related Work

To our knowledge, there has been no work that computationally explores the impact of agenda setting on diffusion.

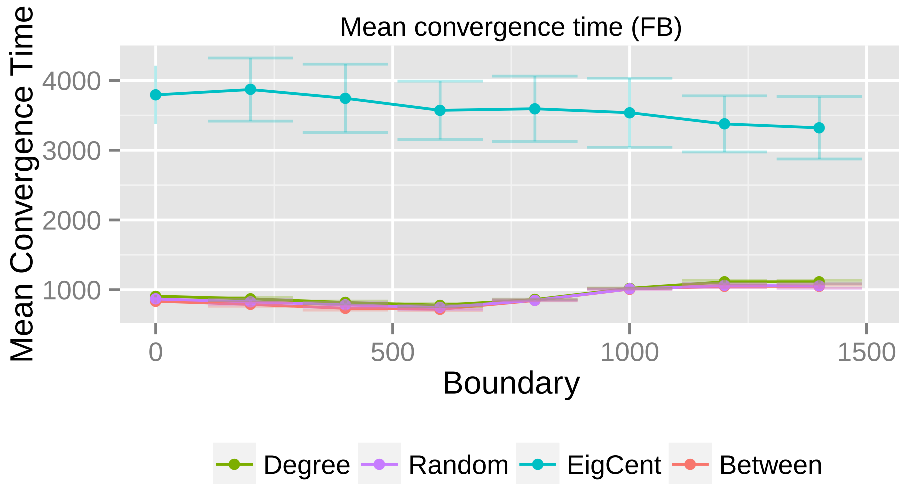


**Fig. 5.** Mean time to convergence +/- 1 standard deviation for the Facebook network.

The interpersonal interaction dynamics of our model are closely related to the *voter model* a well explored model from the physics domain [21] in which nodes can take on the values 0 or 1. At each timestep a random nodes takes on the value of one of its neighbors. While the choice of agents is similar, few voter models capture multiple interacting values within an agent.

Chapter 10 of [22] describes the "consensus = coherence + communication" (CCC) model. In this model each agent has a *parallel constraint satisfaction* network where concepts represent hypotheses and evidence, and links represent explanatory relations. The purpose of the model was to explain the diffusion of scientific theories in a population, so pairs of agents could interact (simulating a complete network) and agents could execute "lectures" for many other agents.

The main difference between our model and the CCC model is in the interaction. In the CCC model agents exchange concepts with others, thus changing the structure of the network. We have focused on a different perspective, given that individuals already know the links between concepts, how do persuasive messages between agents lead to attitude change.

The parallel constraint satisfaction model has been explored in a wide variety of contexts, such as impression formation [23], legal inference [24–26], and as a model of change in attitude to the persian gulf war [27].

In Axelrod's model of social dynamics agents have multiple cultural features where each feature can take on a value from a small set of "traits" [28],[29]. Agents can interact only The model uses a bounded confidence in which only agents that match on a certain percentage of features will interact.

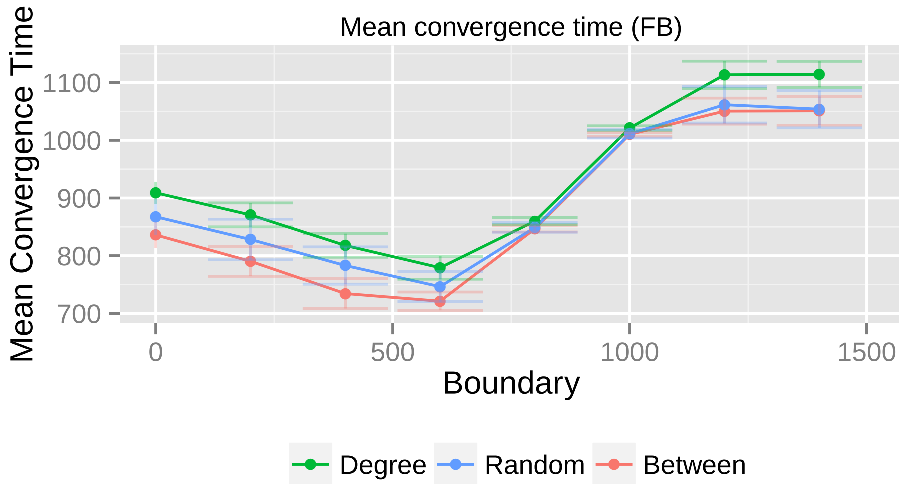


**Fig. 6.** Mean time to convergence +/- 1 standard deviation for the Facebook network. Excluding the Eigenvectgor Centrality measure.

Several extensions to Axelrod's model have been proposed that incorporate mass media. Often, this is incorporated through a virtual agent that represents the media and which has edges connecting it to all other agents in the population [29]. This work studies the influence effect of media, and not the agenda setting effect.

## 5 Conclusions

Agent-based simulation is an important tool that allows empirically study of complex interactions, in our case between interpersonal influence and attitudes. In this work, we developed a novel agent based model that captures social and cognitive factors that affect decision making (the MAMA model). We used agent-based simulation to study the impact of agenda setting on diffusion time within the MAMA model.

We found that:

1. Agenda setting can significantly reduce diffusion time.
2. Agenda setting paired with strategic seed choice results in the quickest diffusion time.

Surprisingly, merely setting the topic of discussion (not even influencing attitudes) can have an impact on diffusion time.

These results suggest that for rapid diffusion, influencers ad campaigns should focus on identifying the "influencers" in a network, and also understand the underlying attitudes towards a product.

## 6 Future Work

A critical part of the MAMA model is the definition of the cognitive network. While we have focused solely on the "3-Fan" network it's clear that other topologies could reasonably exist. Variance in cognitive network topology may dramatically increase diffusion time, or preclude convergence entirely (for instance, in the case where there are no consistent states).

While we can hypothesize many different cognitive networks, identifying real world cognitive networks may be more difficult. There has been some work for specific domains (e.g., health decisions [30]), but a general procedure to learn the links between attitudes does not seem to exist yet.

## References


1. Russo, J.E., Carlson, K.A., Meloy, M.G., Yong, K.: The goal of consistency as a cause of information distortion. Journal of Experimental Psychology: General (2008)
2. Simon, D., Snow, C.J., Read, S.J.: The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. Journal of Personality and Social Psychology **86**(6) (2004) 814–837
3. Visser, P.S., Cooper, J.: Attitude change. In Hogg, M., Cooper, J., eds.: Sage Handbook of Social Psychology. Sage Publications (2003)
4. Vu, H.N.N., Gehrau, V.: Agenda diffusion: An integrated model of agenda setting and interpersonal communication. Journalism & Mass Communication Quarterly **87**(1) (March 2010) 100–116
5. McCombs, M.: A look at agenda-setting: past, present and future. Journalism Studies **6**(4) (2005) 543–557
6. Rogers, E.M.: Diffusion of Innovations. Fifth edn. Free Press (2003)
7. Kempe, D., Kleinberg, J., Éva Tardos: Maximizing the spread of influence through a social network. In: Proceedings of 2003 SIGKDD. (2003)
8. Dechter, R.: From local to global consistency. Artificial Intelligence **55**(1) (May 1992) 87–107
9. Eagly, A.H., Chaiken, S.: The Psychology of Attitudes. Wadsworth Cengage Learning (1993)
10. O'Keefe, D.J.: Persuasion: theory & research. Sage, Thousand Oaks, Calif. [u.a. (2002)
11. Sood, V., Redner, S.: Voter models on heterogenous graphs. Physical Review Letters **94** (2005) 178701.
12. Gromet, D.M., Kunreuther, H., Larrick, R.P.: Political ideology affects energy-efficiency attitudes and choices. Proceedings of the National Academy of Sciences (April 2013) 201218453 PMID: 23630266.
13. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature **393**(6684) (June 1998) 440–442

14. Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal **Complex Systems** (2006) 1695
15. Leskovec, J.: Stanford network analysis project
16. Barabási, A.L., Albert, R.: Emergence of scaling in random networks. Science **286**(5439) (October 1999) 509–512 PMID: 10521342.
17. McAuley, J., Leskovec, J.: Learning to discover social circles in ego networks. In: Proceedings of the 2012 Neural Information Processing Systems Conference. (2012)
18. Lakkaraju, K.: Reducing diffusion time through agenda setting in a multi-agent multi-attitude model. Submitted for review at the 2014 Autonomous Agents and Multiagent Systems Conference (2014)
19. Wasserman, S., Faust, K.: Social Network Analysis: Methods and Applications. Structural Analysis in the Social Sciences. Cambridge University Press (1994)
20. Chen, W., Collins, A., Cummings, R., Ke, T., Liu, Z., Rincon, D., Sun, X., Wang, Y., Wei, W., Yuan, Y.: Influence maximization in social networks when negative opinions may emerge and propagate. In: Proceedings of the 11th SIAM International Conference on Data Mining (SDM'2011). (2011)
21. Sood, V., Antal, T., Redner, S.: Voter models on heterogenous networks. Physical Review E **77** (2008)
22. Thagard, P.: Coherence in Thought and Action. Life and Mind: Philosophical Issues in Biology and Psychology. MIT Press (2000)
23. Kunda, Z., Thagard, P.: Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. Psychological Review **103**(2) (1996) 284–308
24. Simon, D., Krawczyk, D.C., Bleicher, A., Holyoak, K.J.: The transience of constructed preferences. Journal of Behavioral Decision Making **21**(1) (2008) 1–14
25. Simon, D.: A third view of the black box: Cognitive coherence in legal decision making. University of Chicago Law Review **71**(511) (2004)
26. Simon, D., Krawczyk, D.C., Holyoak, K.J.: Construction of preferences by constraint satisfaction. Psychological Science **15**(5) (2004) 331
27. Spellman, B.A., Ullman, J.B., Holyoak, K.J.: A coherence model of cognitive consistency: Dynamics of attitude change during the persian gulf war. Journal of Social Issues **49**(4) (1993) 147–165
28. Axelrod, R.: The dissemination of culture: A model with local convergence and global polarization. The Journal of Conflict Resolution **41**(2) (1997)
29. Peres, L.R., Fontanari, J.F.: The media effect in axelrod's model explained. EPL (Europhysics Letters) **96**(3) (November 2011) 38004
30. Shoda, Y., Mischel, W.: Personality as a stable cognitive-affective activation network: characteristic patterns of behavior variation emerge from a stable personality structure. In Read, S.J., Miller, L.C., eds.: Connectionist Models of Social Reasoning and Social Behavior. Lawrence Erlbaum (1998)