

Chapter 1

Effect of Humans on Belief Propagation in Large Heterogeneous Teams

Praveen Paruchuri, Robin Grinton, Katia Sycara and Paul Scerri

Abstract Members of large, heterogeneous teams often need to interact with different kinds of teammates to accomplish their tasks, teammates with dramatically different capabilities to their own. While the role of humans in teams has progressively decreased with the deployment of increasingly intelligent systems, they still have a major role to play. In this chapter, we focus on the role of humans in large, heterogeneous teams that are faced with situations, where there is a large volume of incoming, conflicting data about some important fact. We use an abstract model of both humans and agents to investigate the dynamics and emergent behaviors of large teams trying to decide whether some fact is true. In particular, we focus on the role of humans in handling noisy information and their role in convergence of beliefs in large heterogeneous teams. Our simulation results show that systems involving humans exhibit an enabler-impeder effect, where if humans are present in low percentages, they aid in propagating information; however when the percentage of humans increase beyond a certain threshold, they seem to impede the information propagation.

1.1 Introduction

As intelligent systems get increasingly deployed in the real world, humans need to work in cooperation with various other entities such as agents and robots. The effective sharing and use of uncertain information in decision-making is the key to the success of such large heterogeneous teams in complex environments. Typically, noisy information is collected by some portion of the team and shared with others via a network [4]. Each team member will then use this uncertain information and the beliefs of those around them to develop their own beliefs about relevant facts. For example, consider the case of Network Centric Warfare [7]. In this domain, a

Carnegie Mellon University, Address of Institute, e-mail: {paruchur+, rglinton+, katia, pscerri}@cs.cmu.edu

large number of agents and humans must do a variety of tasks, such as planning, information fusion and resource allocation. All these tasks need the exchange of potentially uncertain information between the different team members to arrive at a decision required to achieve joint goals. In this chapter, we focus on the role of humans in handling noisy information and their role in convergence of beliefs in large heterogeneous teams that include machine agents.

We developed an abstracted model and simulator for studying the dynamics and interaction between the various members of large, heterogeneous teams. In our model, the team is connected via a network with some team members having direct access to sensors, while others relying solely on neighbors in the network to provide inputs for their beliefs. Each agent then uses Bayesian reasoning over the beliefs of direct neighbors and sensor data to maintain its own belief about a single fact which can be *true*, *false* or *unknown*. Simulation results indicate that the number of agents coming to a correct conclusion about a fact and the speed of their convergence to this belief depends on various factors including network structure, the percentage of humans in the team and the conditional probability of belief in neighbor's information.

Our experiments show that our belief propagation model exhibits the property that a single additional piece of data can cause a cascade of belief changes, referred to as *avalanches*. In particular, we observed that the frequency distribution of these avalanches follows a *power law* over a certain conditional probability range, a key feature of *SOC's* (Self-organizing critical systems). Our results also indicate that humans have a significant effect on the way these avalanches occur. In particular, we uncovered a new effect due to the presence of humans which we call the *enabler-impeder* effect which holds across a general range of experimental settings.

The rest of this chapter is organized as follows. Section 1.2 provides a brief description of the Self Organized Criticality property. Section 1.3 describes the *enabler-impeder* effect, a new effect that we uncovered in this work. Section 1.4 provides the actual model of a heterogeneous team that we use for purposes of simulation. Section 1.5 provides the results of the various experiments we performed using our simulator. In particular, we present extensive simulation results that shows the influence of humans under different system parameters on the behavior of large scale heterogeneous teams. Our results uncover the *enabler-impeder* effect due to the presence of humans and the *SOC* nature of such teams. Section 1.6 provides the related work while Section 1.7 summarizes the chapter and provides a list of possible extensions for this work.

1.2 Self-Organized Critical Systems

There has been a lot of focus on identifying systems that exhibit Self Organized Criticality (*SOC*) as a property [8, 2, 12, 6]. The main reason that *SOC* behavior is considered important is that most of the observed phenomenon in *SOC* systems occur for a large range of parameters [2], as opposed to the dependency on fine-tuning

of parameters in general complex systems. Systems which exhibit SOC share the following fundamental characteristics. All SOC systems consist of a large number of interacting constituents which interact with a relatively small percentage of the other system elements. The behavior of these constituents are dominated by three factors: (a) An external drive that changes the state of the individual; (b) Resistance of the individual to change; (c) Threshold in the local resistance at which the individual changes its belief or behavior. In our experimental section, we show that the system modeled here indeed exhibits these features over a large parameter range. A key feature of SOC's is that the avalanches caused by a single additional piece of data exhibits a power law property over a certain conditional probability range. The power law implies that the distribution of avalanche sizes is dominated by many small avalanches and exponentially fewer large ones. Identifying this behavior is important since the critical behavior of SOC's is not dependent on finely tuned parameters, and hence we can expect this criticality to occur often in real-world systems too. The power law suggests that large avalanches are relatively infrequent. However when they do occur, when sparked by incorrect data, the result can be that the entire team reaches a wrong conclusion despite exposure to primarily correct data. In many domains such as our Network Centric Warfare, this is an unacceptable outcome even if it does not occur often.

1.3 The Enabler-Impeder Effect

Addition of humans to large heterogeneous teams results in changing the interaction dynamics of the existing system. In particular, we uncovered a new team dynamics which we refer to as the enabler-impeder effect. This effect means that up to a certain percentage, humans enable belief propagation and above that percentage they actually impede the propagation. Thus, big avalanches occur most often in teams with some humans, but less often in teams with few or many humans. The reason behind the enabler-impeder effect is as follows: The enabler effect is due to the fact that agents have a higher belief in humans while the impeder effect is due to the fact that humans have a high latency between successive belief calculations. More details of our model of humans is presented in section 1.4.

We now provide a brief discussion on the effect of the enabler-impeder effect on the formation of avalanches. In general, occurrence of avalanches is good as it implies that agents are able to convey their beliefs to many other agents in the network. Bigger avalanches imply that each agent is able to convey and influence the beliefs of a larger set of agents. However, really large avalanches can potentially be harmful since agents can start imposing their beliefs on others or manipulate the network. We mention the word potentially here because there can be situations where we may want one agent's opinion to be instantaneously propagated across the network. Such situations can include an emergency where we want an instantaneous propagation of beliefs or it can be that the new input is the truth and hence propagating it instantaneously is the correct thing to do. Considering the fact that avalanches are good

up to moderate sizes and potentially harmful thereafter, the enabler-impeder effect of humans can be classified as a positive effect on the team. This is because while a moderately high percentage of humans in a team helps in forming avalanches, addition of humans thereafter impedes the formation of big avalanches that are potentially harmful.

1.4 Model of Information Dissemination in a Network

We now introduce the network model similar to the one presented in [8]. In this model, there is a team of agents, $A = \{a_1, \dots, a_N\}$, where an agent is a generic term to imply humans, robots, or agents. We are interested in large teams, $|A| \geq 1000$, that are connected to each other via a network, K . K is represented by $|A| \times |A|$ matrix where $K_{i,j} = 1$, if i and j are neighbors and 0 otherwise. The network is assumed to be relatively sparse, with $\sum_j K_{i,j} \approx 4, \forall i$. There are six different network structures we consider in this chapter: (a) Scalefree (b) Grid (c) Random (d) Smallworld (e) Hierarchy and (f) HierarchySLO (A hierarchy network with sensors only at the leaf nodes). A comprehensive definition of these networks can be found in [11].

Some members of the team, $H \subset A$ are considered to be humans. Certain members of the team, $S \subset A$ with $|S| \ll |A|$, have direct access to a sensor. A sensor simply returns a value of *true* or *false* with a particular probability. Sensor readings are noisy, so the sensor could return incorrect values with some probability. The frequency with which a sensor returns the correct value is modeled as a random variable R_s which is normally distributed with a mean μ_s and a variance σ_s^2 . Agents that are directly connected to sensors incorporate new sensor readings according to the following belief update equation.

$$P'(b_{a_i} = \text{true}) = \frac{A}{B + C}$$

$$A = P(b_{a_i} = \text{true})P(s_{a_i} = \text{false}/F = \text{true})$$

$$B = (1.0 - P(b_{a_i} = \text{true}))P(s_{a_i} = \text{false}/F = \text{false})$$

$$C = P(b_{a_i} = \text{true})P(s_{a_i} = \text{false}/F = \text{true})$$

Agents use the following equation to incorporate the beliefs of neighbors:

$$P'(b_{a_i} = \text{true}) = \frac{D}{E + G}$$

$$D = P(b_{a_i} = \text{true})P(b_{a_j} = \text{false}/F = \text{true})$$

$$E = (1.0 - P(b_{a_i} = \text{true}))P(b_{a_j} = \text{false}/F = \text{false})$$

$$G = P(b_{a_i} = \text{true})P(b_{a_j} = \text{false}/F = \text{true})$$

where, $P(b_{a_i})$ gives the prior belief of agent a_i in the fact F and $P(s_{a_i}/F)$ gives the probability that the sensor will return an estimate of the fact (*true* or *false*) given the

actual truth value of the fact. We refer to, $P(b_{a_j}/F)$, interchangeably as the *belief probability* or the *conditional probability* or *CP*. $P(b_{a_j}/F)$ gives a measure of the credibility that an agent a_i assigns to an estimate of the value of F received from a neighbor a_j given F .

Each agent decides that F is either *true*, *false* or *unknown* by processing its belief using the following rule. If T is a threshold probability and ϵ an uncertainty interval, the agent decides that the fact is *true* if $P(b_{a_i}) > (T + \epsilon)$, *false* if $P(b_{a_i}) < (T - \epsilon)$, and *unknown* otherwise. Once the decision is made, if the agent's decision about F has changed, the agent reports this change to its neighbors. Note that in our model, neither $P(b_{a_i})$, the probability that F is *true* according to agent a_i , or the evidence used to calculate it, is transmitted to neighbors. Instead, the only parameter that is transmitted is the actual value of F when there is a change in belief (i.e., either *true*, *false* or *unknown*).

We use a distinct model for the humans in our simulator. In particular, we make the following two assumptions with respect to humans: (a) Humans are given much larger credibility than other agents. In other words, agents have a higher belief in human's estimation of the value of F than in the value of F given to them by other agents. That is, for $a_h \in H$ and $a_k \notin H$, $P(b_{a_h}/F) \gg P(b_{a_k}/F) \forall_{h,k}$. This is to abstract away the fact that humans are good at analyzing (given particular information) and hence agents believe in humans more. (b) Humans have a higher latency between successive belief calculations. Effectively, this means that humans are slower to sense the environment or update their beliefs. This is to account for the fact that humans do not just believe others at face value, but change their beliefs using more complicated reasoning models without getting into the detailed modeling issues. We believe that these two assumptions help in achieving our goal of building more realistic human models.

1.5 Simulation Results

We performed a series of seven experiments to understand the effects of humans on the key properties and predictions of the system. All our experiments were performed with networks of 1000 agents with six different network topologies [11]. These six different topologies as mentioned earlier are: (1) Scalefree (2) Grid (3) Random (4) Smallworld (5) Hierarchy and (6) HierarchySLO. Before presenting the details of our experiments, we introduce a metric named center of mass which is useful to concisely capture the dynamics of belief cascades. Center of mass is defined as:

$$\frac{\sum_{i=1 to N} (AvalancheFrequency_i * AvalancheSize_i)}{\sum_{i=1 to N} AvalancheFrequency_i}$$

where $AvalancheSize_i = i$, $AvalancheFrequency_i$ is the frequency of avalanche size i and N is the maximum possible avalanche size (which is 1000 in our experiments since there are 1000 agents in our simulation).




Fig. 1.1 Enabler-Impeder effect of humans.

Our first experiment studies the effect that humans have on belief propagation within the system. Recall that our modeling assumption about humans is that: (a) They have higher credibility, i.e., an agent believes a human with higher probability than it believes another agent and (b) Humans are slower to make decisions i.e., they have a higher latency between successive belief calculations. In our first experiment, we varied the percentage of humans in the team from 0 to 100 in increments of 20 across the six different network topologies and measured the center of mass of the avalanches obtained. The x-axis of Figure 1.1 shows the percentage of humans in the network increasing from 0 to 100% and the y-axis represents the center of mass. The figure shows the enabler-impeder effect due to the presence of humans in the team. In particular, we observe that the center of mass increases till the percentage of humans in the team reaches 20%, the enabler effect, and then decreases thereafter, the impeder effect. This effect was noted across all the network types. The higher credibility of humans tends to encourage a change of belief in others who receive human input and thus humans can encourage or enable the formation of avalanches. Conversely the long latency of humans in changing their beliefs has the effect of impeding the avalanche formation. When the percentage of humans is below 20% the avalanche enhancing effect of the high human *CP* predominates, while above this value the impedance effect caused by the long latency predominates.

In our second experiment, we show that the network convergence rate is a function of percentage of humans in the team for a Scalefree network. Figure 1.2 shows the varying percentages of humans on the x-axis and the frequency of convergence to the ground truth on the y-axis. The total number of runs for this experiment is 5000 and the ground truth is *True*. Each line in the graph corresponds to a particular percentage or greater number of agents converging to the ground truth. Note that greater than 90% of agents, never converged to the ground truth and hence the graph




Fig. 1.2 Network convergence rate as a function of human percentage.

shows only four lines corresponding to greater than or equal to 50, 60, 70 and 80 percent. An example data that can be obtained from this graph is: If the percentage of humans in the team is 20, then at least 50% of the agents have converged to the ground truth 3833 times out of the total 5000 runs. When there are no humans in the team more than 60% of the agents never converged to the ground truth.

Figure 1.2 confirms our observation that there is an enabler-impeder effect. In particular, addition of humans to the team helped in formation of avalanches, thus implying that humans play a key role in propagating the ground truth. This observation is true here since more than 50% of the team converges to the ground truth in almost all the settings of the graph. However, once the percentage of humans crosses 60, the frequency of convergence to the ground truth decreases since humans impede the propagation of information and hence the propagation of the ground truth.

Our third experiment shows that our system exhibits SOC behavior with a varying CP and percentage of humans in the team. In particular, we performed the experiment varying the CP values from 0.52 to 0.58 in increments of 0.02 and the human percentage from 0 to 100 in increments of 20. The results of this experiment are shown in Figure 1.3. The five plots in the figure correspond to varying the human percentage in the system. The x-axis in all the plots denotes the size of the avalanches and the y-axis denotes the frequency with which avalanches of various sizes occur. Both the axes are plotted on a log-log scale and hence we expect a power law to be evident as a straight line. The underlying network used to produce this figure was a scale-free network. The 4 lines in the figure correspond to CP values of 0.52, 0.54, 0.56 and 0.58. Each point in the graph is an average over 5000 runs. In general, CP values greater than or equal to 0.6 did not exhibit the power law. This is because if agents give their neighbor's estimates high credibility (large values of CP), all the agents change their beliefs immediately upon receiving input

from a neighbor. Hence, the bigger avalanches occur frequently without having a distribution over the small and medium size avalanches which led to the power law. While we show this figure solely for a scalefree network, our experiments show that this trend is true across all the network topologies we examined.

From Figure 1.3, we notice that most of the plots follow the power law in the CP range from 0.52 to 0.58. However, the addition of humans has a significant effect. When the percentage of humans was zero, the plot was more spread out i.e., CP had a greater effect and hence there is a greater spread between the lines even with small CP increments. However, the effect of CP decreases even with the addition of a small percentage of humans in the team i.e., there is not much spread between the various CP lines when the percentage of humans in the team is greater than zero. The explanation for this is as follows. When there are no humans in the team, the main parameter affecting the avalanche sizes was the belief in neighbors. Hence, when CP increases there is a great variation in the occurrence of various avalanche sizes. However as humans get added to the team (even by a small percentage), given that agents have a higher belief in humans, they have a much greater effect on the dynamics of avalanche formation. Thus, increasing the CP in teams with humans has a small effect on the dynamics of avalanches formed as opposed to the large effects seen when there are no humans.

Our fourth experiment looks at the relationship between the number of sensors and the center of mass for a Scalefree network as the percentage of humans in the team is varied. A similar trend holds for the other network topologies considered in this chapter. Figure 1.4 shows the number of sensors varying from 10 to 485 on the x-axis and the center of mass on the y-axis. We assume that sensors return observations randomly on an average of every second step. The six lines in the graph correspond to human percentages of 0, 20, 40, 60, 80 and 100. Each point in the graph is generated using 1000 agents and averaged over 1000 runs. The graph shows that as the number of sensors increase the center of mass of avalanches decreases (or increases by a negligible amount), across the various human percentages. This is because agents with sensors have access to more information and are thus more certain in their own beliefs. Hence, they are less likely to change their beliefs in response to information from neighbors and hence the cascades of belief changes are likely to be smaller. This effect of low probability of beliefs in neighbors can be a problem if some sensors in the system give out particularly bad data since such agents are likely to ignore potentially correct information from their neighbors.

Another interesting point to note from this graph is the enabler-impeder effect described in our earlier experiment. In Figure 1.4, let us consider the center of mass values when the number of sensors is fixed to 210. The center of mass values at this point are $< 1.29, 1.39, 1.40, 1.33, 1.21, 1.02 >$ corresponding to the following percentages of humans: $< 0, 20, 40, 60, 80, 100 >$. The enabler effect is predominantly visible here since the center of mass increases until the percentage of humans reaches 40. However, above 40% the center of mass values decrease. This shows that varying the human percentage has a strong effect on the outcome even if the other parameters of the experiment change significantly. A similar pattern is observed in all our experiments presented below even though a variety of parame-

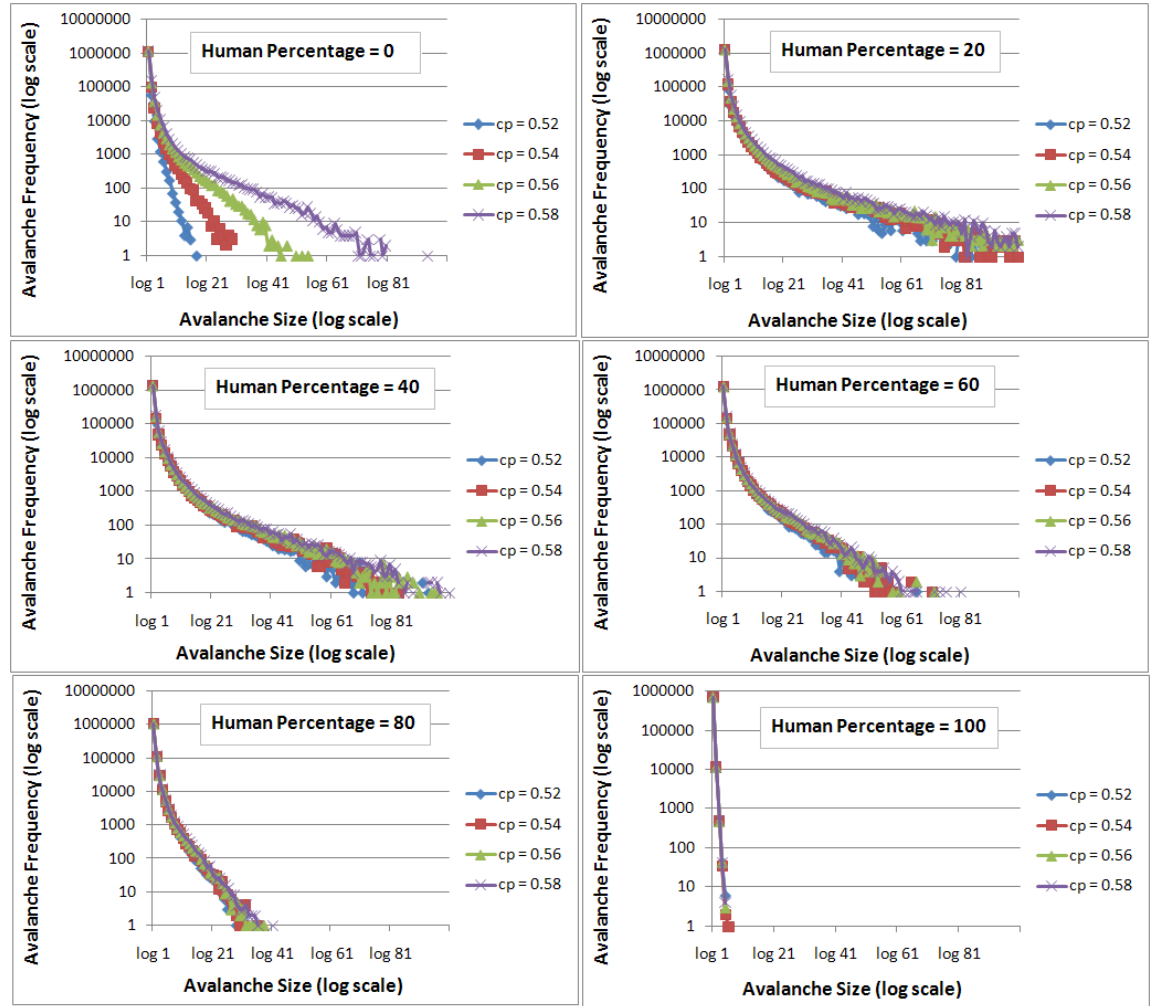


Fig. 1.3 Avalanche size versus frequency across various percentages of humans.

ters are changed, thus showing the generality of this effect. To avoid repetition, we indicate the occurrence of this effect in the rest of the experiments without repeating the description.

Our fifth experiment looks at the relationship between the network density and the center of mass values across six different network topologies as the percentage of humans in the team is varied. Figure 1.5 shows six plots each corresponding to a different network type. In each plot, the network density is plotted on the x-axis and the center of mass on the y-axis. The six lines in each plot correspond to the six different percentages of humans in the team. Each point in the graph is generated using 1000 agents averaged over 1000 runs. The plots show that as the




Fig. 1.4 Number of sensors versus Center of Mass.

network density increases the center of mass in general increases except for the grid network where the center of mass mostly remains constant. This is due to the fact that each agent can influence more neighbors and hence lead to the formation of bigger avalanches resulting in a higher or at least equal center of mass. As mentioned earlier, the enabler-impeder effect of humans remains a predominant effect in this experiment too across all the networks except the smallworld network.

Our next experiment looks at the effect of the increase in agent's belief in humans over the center of mass as the percentage of humans in the team is varied. Figure 1.6 shows the agent's belief in humans on the x-axis and the center of mass on the y-axis. The six lines in the figure correspond to the human percentage in the team varied from 0 to 100. From the plot we can see that, as the agent's belief in humans increase, the center of mass in general increases across all percentages of humans in the team except when the percentage is 0 or 100. When the percentage of humans is 0, since there are no humans in the team the parameter belief probability in humans has no effect as expected. Similarly, when the percentage of humans is 100, i.e., when there are no agents in the team, the size of the center of mass remains constant since the parameter cannot have any effect. Similar to all our earlier plots, the enabler-impeder effect is observed in this plot too.

Our last experiment looks at the effect of increase in human's belief in other humans on the center of mass as the percentage of humans in the team is varied. Figure 1.7 shows the belief probability of humans in humans varying from 0 to 1 on the x-axis and the center of mass on the y-axis. The six lines in the figure correspond to the percentage of humans in the team varying from 0 to 100. When there are no humans in the team the belief probability parameter does not affect the simulation and hence a straight line is obtained as expected. For the rest of the lines, when the belief probability is zero it means that humans do not believe other humans and

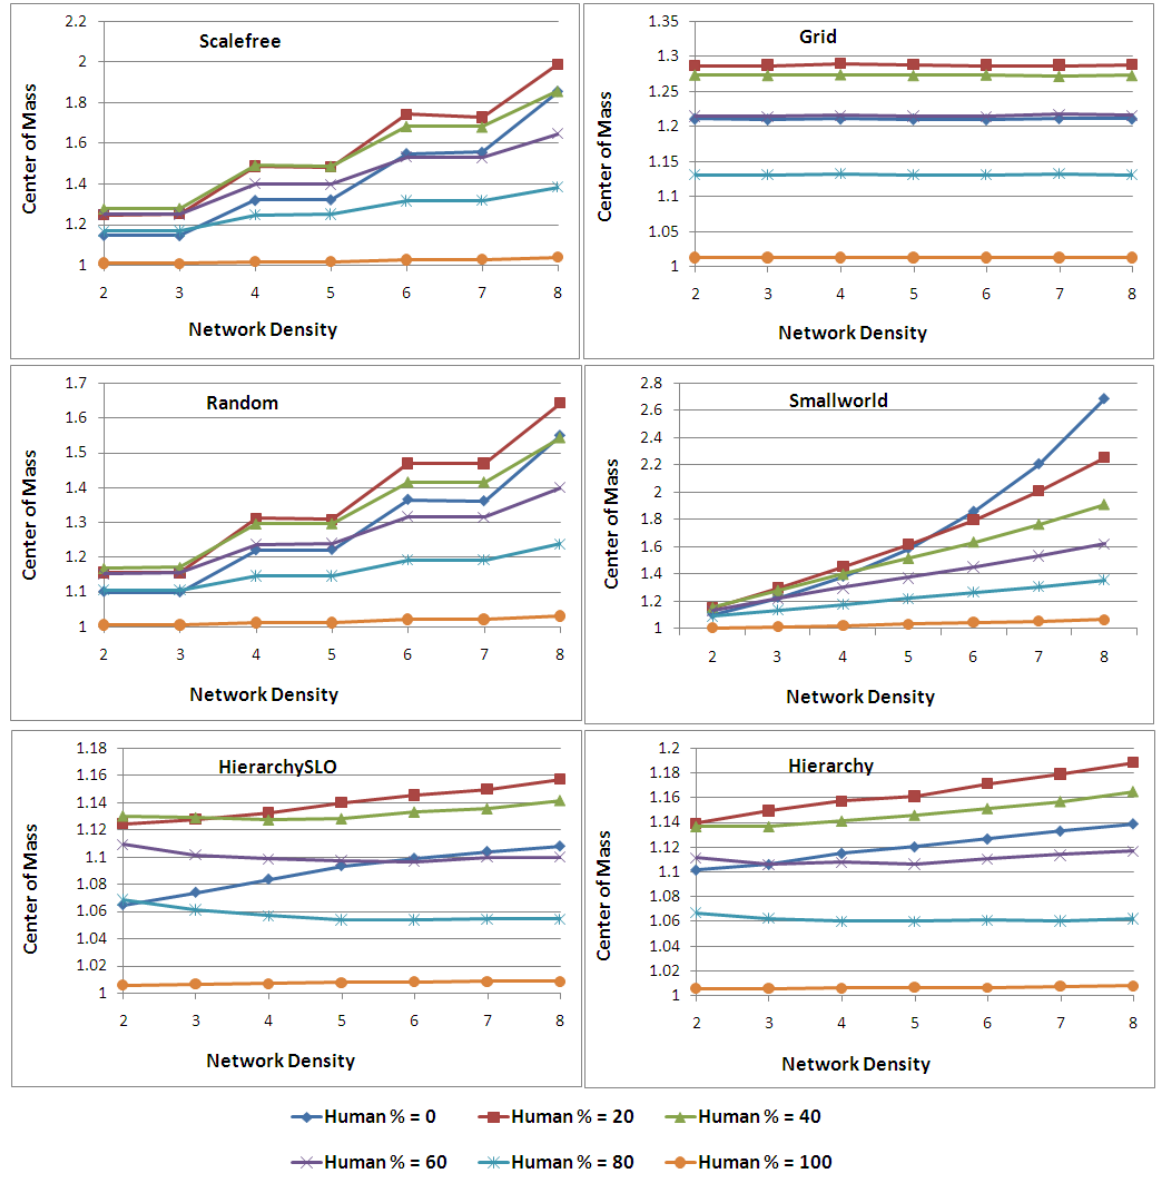


Fig. 1.5 Density versus Center of Mass.

hence disregard other human's inputs. As the belief probability increases, humans start considering others beliefs, but still remain skeptical. This trend continues till 0.5 where the situation becomes equivalent to information coming from an unbiased coin toss. Due to this skeptical nature of humans in other humans, they effectively




Fig. 1.6 Effect of increase in agent's belief in humans.

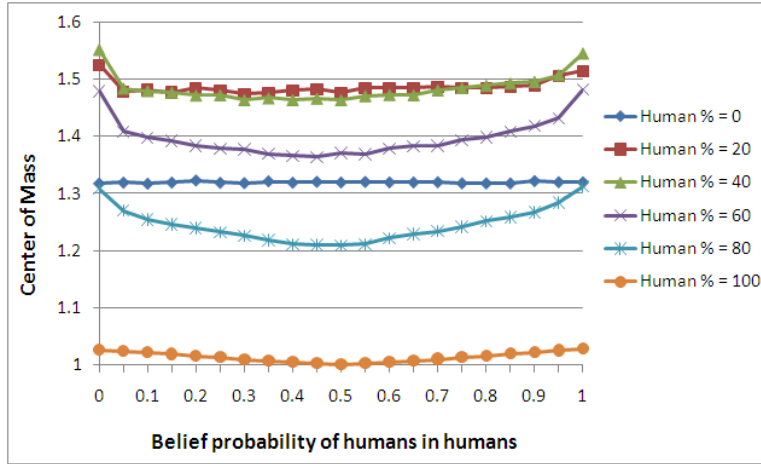


Fig. 1.7 Effect of increase in agent's belief in humans.

impede the formation of avalanches and hence lower the center of mass. As the belief probability in other humans crosses 0.5 and moves towards 1, humans start believing other humans more, enabling the formation of bigger avalanches and hence increase the center of mass. As obtained in all our earlier plots the enabler-impeder effect is observable across the various human percentages in the team.

1.6 Related Work

Many studies have focused on the belief propagation models for large heterogeneous teams. While most of these studies focus on the cascade effects of belief propagation, not much effort has been put into studying the specific effect of humans on such teams. For example, in [14] a belief propagation model is used for a large heterogeneous team. Using this model the paper shows that the cascades generated by their system follows a power law and hence exhibits SOC dynamics. Our model is similar to theirs, in the sense that they too use a heterogeneous nodes model with state change based on a distribution of thresholds. However, they do not introduce an explicit model of the human. The main difference introduced due to the presence of humans is that humans have a higher credibility and increased latency for changing as compared to other agents. We then explored the effects of these changes and uncovered two key properties: (a) The presence of humans results in the enabler-impeder effect, a new effect that has not been uncovered earlier. (b) The belief dynamics of our system across various human percentages exhibits the properties of a Self-Organizing Critical system. In [1], Motter discusses cascading failures in electrical networks. In particular, Motter investigates cascades due to removal of nodes as opposed to cascades of decisions as in our work. Furthermore, there is no explicit model of humans considered in their work.

In contrast, [3], [5] and [10] study the properties of heterogeneous human teams. While each of these papers presents a different way in which the team members update their beliefs, all these works uncover the cascading effects in large scale teams with humans. For example, [10] studies the cascading effects in heterogeneous teams where the individuals are modeled as naive Bayesian norm followers. Using this model they uncovered the herding behavior in human societies, a direct result of the cascading effect of belief propagation. However, none of these works uncover the enabler-impeder effect that humans bring in nor the SOC nature of such teams.

1.7 Conclusions and Future Work

Our conclusions from this work are the following. We built an abstract model and a simulator that studies the dynamics and interactions between the various members of a heterogeneous team. We then performed various experiments to study the properties of the simulator. In particular, we focused on the role of humans in all our experiments. Our main results are as follows: (a) Humans can have a dramatic effect on information propagation which we characterized as the enabler-impeder effect. In particular, the effect means that small percentages of humans encourage information propagation in the system while large percentages inhibit the information propagation. (b) We also found that the enabler-impeder effect is consistent even if other parameters of the domain change significantly. (c) We demonstrated that our system of belief sharing agents with varying percentages of humans in the team ex-

hibits SOC dynamics for certain parameter ranges, thus providing a basis for why our results might hold in a real situation.

In the future, we plan to build more realistic models of humans. In particular, we would like to model the fact that humans may not maintain exact numbers for beliefs but a more complicated model. This complicated model could be maintaining beliefs as intervals or as probability distributions. We also plan to attribute psychological factors for humans such as optimism, pessimism etc. and other personal attributes such as motivations and emotions. Effectively, such a modeling can bring out the notion that humans are decision makers rather than just accumulators of belief. We would also like to model the fact that humans do not just communicate belief probabilities but also the reasoning behind them. This would require us to develop richer communication models.

Acknowledgments

This research was funded by AFOSR grants FA95500810356 and FA95500710039.

References

1. Adilson E. Motter, Ying-Cheng Lai (2002) Cascade-based attacks on complex networks. *Physical Review E*, Vol. 66, Issue 6.
2. Bak, P. and Tang, C. and Wiesenfeld, K. (1983) Self-organized criticality: an explanation of $1/f$ noise. *Physical Review Letters*, 381–384.
3. Banerjee, A. V. (1992) A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 107: 797–817.
4. Bellur, B. and Lewis, M. and Templin, F. (2002) An Ad-hoc Network for Teams of Autonomous Vehicles. *Proceedings of the First Annual Symposium on Autonomous Intelligence Networks and Systems*.
5. Bikhchandani, S. (1992) A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades, *Journal of Political Economy*, 100: 992–1026.
6. Clar, S. and Drossel, B. and Schwabl, F. (1994) Scaling laws and simulation results for the self-organized critical forest-fire model. *Phys Rev E*, 50, 1009–1018.
7. Dekker, A. H. (2003) Centralisation and Decentralisation in Network Centric Warfare. *Journal of Battlefield Technology*, 6(2), 23–28.
8. Grinton, R. and Paruchuri, P. and Scerri, P. and Sycara, K. (2009) Self Organized Criticality of Belief Propagation in Large Heterogeneous Teams. *Dynamics of Information Systems: Theory and Applications*, Hirsch, M. J. and Pardalos, P. and Murphey R. (Eds.), Springer Optimization and its Applications, 2010.
9. Malte-Sorensen, A. and Feder, J. and Christensen, K. and Frette, V. and Josang, T. and Meakin, P. (1999) Surface fluctuations and correlations in a pile of rice. *Physical Review Letters*, 83, 764–767.
10. Neill, D. B. (2005) Cascade Effects in Heterogeneous Populations. *Rationality and Society* 2005; 17; 191.
11. Newman, M.E.J. (2003) The structure and function of complex networks. *SIAM Review*, 45, 167–256.

12. Olami, Z. and Feder, H and Christensen, K. (1992) Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes. *Physical Review Letters*. 68, 1244–1247.
13. Sukthankar, G. and Sycara, K. and Giampapa, J. and Burnett, C. (2008) A Model of Human Teamwork for Agent-Assisted Search Operations. NATO Research and Technology Organization: Human Factors and Medicine Panel Symposium (NATO HFM-142/RTO Symposium).
14. Watts, D.J. (2002) A simple model of global cascades on random networks. *Proceedings of the National Academy of Science*, 99, 5766–5771.
15. Zhang, Y. (1989) Scaling theory of self-organized criticality. *Physical Review Letters*, 470–473.