How Teams Benefit from Communication Policies: Information Flow in Human Peer-to-Peer Networks

David Reitter
Katia Sycara
Christian Lebiere
Yury Vinokurov
Antonio Juarez
Carnegie Mellon University
Pittsburgh, PA 15213
reitter@cmu.edu, katia@cs.cmu.edu, cl@cmu.edu,
jerryv@andrew.cmu.edu, ajuarez@andrew.cmu.edu

Michael Lewis
University of Pittsburgh
Pittsburgh, 15260
ml@sis.pitt.edu

Keywords: Communication Networks, Social Networks, Belief Propagation, Cognitive Modeling

ABSTRACT *In an experiment involving teams of humans playing a cooperative game, we study the effect of local communication policies on the efficiency and the performance of teams and of individuals in different positions within a network. This design provides an experimental model of human communities, where information may spread from peer to peer by word of mouth. With this model, we explain the realistic tradeoff between liberal information dumping and targeted information sharing by human peers. Human subjects exchanged natural-language messages with relevance to a task, thereby sharing knowledge across a community. Communication took place along the edges of a small-world graph. Cooperation and individual efforts were incentivized. In one condition, participants were asked to request specific information and only supply information that they knew was needed. In another condition, they were asked to supply and forward as much information as possible. We found that a targeted communication policy was successfully implemented by the participants, increased task success, shortened the time it took to get answers to questions, increased efficiency (task success per communication bandwidth), and may have done so selectively for nodes with fewer connections.*

1. Introduction

The basic tenet of teamwork is that joint problem-solving increases efficiency: the whole of all collaborators is stronger than the sum of their contributions. The alternative view of collaboration is that "none of us is as dumb as all of us." Such inefficiencies are commonly blamed on communication overheads. In many modern professional settings, access to information is increased (e.g., through real-time retrieval on the internet), or has been proposed to be increased, for instance with soldiers on battlefields having access to a range of real-time data (Owens, 1995). Two scenarios represent opposite approaches: information is either disseminated widely to all people regardless of its utility, or it is selected by the information consumer based on filtering by, e.g., keywords. Neither of these extremes is satisfactory: wide dissemination burdens the recipient with filtering, while filtering by the sender does not allow downstream nodes to aggregate information of which they are unaware.

To our knowledge there has not been any systematic exploration of system architectures that deliver information to the right recipient at the right time through a communication network, or what the potential benefits and pitfalls may be, taking into consideration cognitive abilities and limitations. Once established, computational and quantitative models could predict the effect of communication policies on information propagation and task performance. These appear particularly useful in the context of large networked organizations, possibly with connections to adversarial networks. In this study, we therefore look beyond the one-on-one situation of dialogue to iterated communication in social word-of-mouth networks. To quantify the success of communication, we use a combined individual and team task as a benchmark. It measures individual and team performance in dynamically changing and time stressed environments.

A computational-cognitive model needs to take into account the simple fact that humans are limited by attention span and by memory. To illustrate, consider a thought experiment. Suppose a network of communicating, cognitive agents that pass questions and answers about distributed information from node to node. We propose two extreme cases: in one, agents have unlimited attention span, that is, their capacity to process information in the short term is not limited. However, agents have no memory to retain information that may be needed later on. In the second case, agents have accurate and infinite memory, but lack the ability to process more than a single piece of information at a time. Which communication policies would be appropriate for the two cases?

Perhaps, in the first case, agents would have to re-send questions regularly, while in the second, they would send out questions only once. Human attention span and memory are both limited; hence, they are subject to a strategic tradeoff. The costs of producing and comprehending linguistic communication, of attention and decaying memory are relevant to a model, as is the structure of the network. Lab-based experiments, in which humans are networked artificially, playing an artificial game, represent a class of non-computational models that explain and predict emergent phenomena observed in the real world. In our case, we strive to elicit the tradeoff between bandwidth optimization through careful information filtering and the dissemination of useful knowledge. This paper introduces a distributed task and describes results from two conditions on either side of the tradeoff. In our communities of experiment participants, we found that increased availability of information in an information-sharing task does not necessarily lead to improved team performance. This reflects a phenomenon we know (or believe we know) from real-life teamwork situations.

Yet, human cognition has evolved in a mix of individual and social environments, which supports the intriguing hypothesis that cognitive limitations are advantageous in a collaborative, social setting where other individuals cannot process an endless stream of information. In networks, humans may act as filters that prevent information overload in the community. A recent simulation has demonstrated how limitations on cognitive capacity can be beneficial as a forcing function for agents to focus on specific sub-problems (Bhattacharyya and Ohlsson, 2010).

We examine networks of humans in small-world configurations, whose structure resembles that of naturally evolved social networks. Our experimental task is a foraging game, called the Geo Game, where participants are motivated by a combination of individual foraging success (finding hidden items in a spatial environment) and the success of their peers. Participants play the game as a group. Their success in the Geo Game depends on the exchange of information. Players are organized by a small-world network; each player is a node in the network and may communicate with its network neighbors. Because it is fully connected, information can be passed throughout the network.

2. Recent Work

Much recent work on social networks has focused on their structural and resulting computational properties, but does so independently of two major aspects of real-world networks: *Humans*, and their *joint objectives, or task*. The study design presented here uses human participants and employs a task that individuals connected over the network have to execute. We provide a measurable objective and a task that depends both on individual performance and collaboration. We connect to work by Bavelas (1950) and Leavitt (1951), who detailed the effects of network structure with human nodes, arguing that networks with centrality show more stable performance, but increased dependency on those central nodes and decreased flexibility with respect to the integration of information. The influence of structural properties in social tasks is evident even when payoffs are determined by individual performance: Judd et al. (2010) show how a social differentiation task (coloring) was harder for subjects when long-distance connectivity in a small world network was increased, while a social agreement task (consensus) appeared easier. Kennedy (1999) used a particle swarm performance model to investigate the effects of network structure. Recently, Bhattacharyya and Ohlsson (2010) show, by way of simulation, how network structure influences the creativity of a community of agents, who exchange partial results in order to achieve creative goals. There, both cognitive agent properties and network structure interact in predicting task performance.

With our task we intend to also complement readily available datasets with an experimental design that gives us

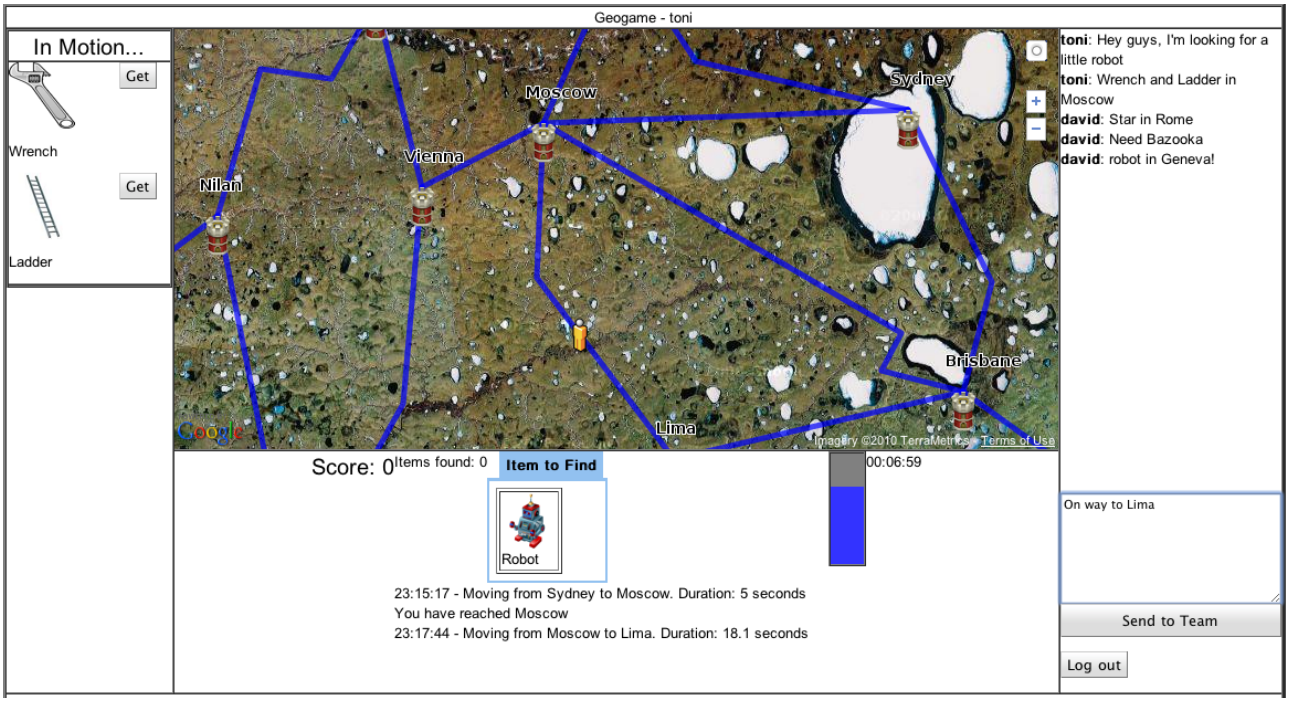

Figure 1: The Geo Game participant interface.

communication data ready to be analyzed in terms of its semantics and its timing. Datasets of language-based communication show the spread of memes or opinions (e.g., Twitter datasets), or they represent socializing or debate that is difficult to analyze and operationalize for the purposes of problem-solving research (Klimt and Yang, 2004). Such datasets, however, are not the result of explicit human collaboration in the context of a well defined task. Tasks, in such datasets, are coincidental, while the task in Geo Game is central to driving communication and provides an objective basis for evaluating the effect of communication on performance.

3. The Geo Game

The Geo Game is a spatial search game, where all players simultaneously engage in a foraging task. Players are shown a map of several named cities, connected by a road network. At any given time, each player is located in one city or is moving between two connected cities; players are shown their own location, but not that of the others (see Fig. 1). The key features of this game are as follows.

Collaborative problem-solving: Players are tasked to find *items* by moving via roads to a city; they find one item at a time (their *goal item*). Participants can visit a neighboring city (directly connected by a road to his current location) by clicking on its symbol on the map. Each city has a small number of items available; this item set differs for every city. Items located in a city are revealed to the participant only while the participant is "visiting" the city. After finding the goal item, a subsequent item is shown to the participant. Moving from city to city takes time, so players are pressured to rely on their knowledge and that of others to find the city efficiently rather than to merely scavenge for items. The time cost for each road is randomized initially, but constant throughout the game. Participants are asked to find as many items as possible within the duration of the session; a timer displays the remaining time.

Dependency on communication: The key feature of the game is that players can improve their performance through communication. They may exchange information through natural language, such as requests for an item they need or responses about the whereabouts of items. A chat interface allows each player to broadcast written messages to a fixed set of other players ("neighbors"). A player, receiving a message, may chose to re-broadcast the information since his set of neighbors is likely to differ from the original message sender. Players use this facility to ask about the whereabouts of their goal item, to reply to others' queries by telling them where the item is available, or to disseminate recently acquired knowledge about the location of items. The task was designed so that the crucial decision a participant had to make was whether to send a piece of information or not, possibly based on its relevance.

Access to the underlying network structure: Players are organized in a graph structure, with participants as nodes, and edges indicating communication channels (Figure 2). This graph is independent of the road network and the locations of the participants. A node y is a network *neighbor* of node x if y and x are connected by one link. In this case, the *distance* between x and y, $d(x, y)$, is 1. The distance between any non-neighbor node z and x is the smallest $1 + d(z, y')$ for any y' that is a neighbor of x. In our experiments, the underlying network topology is a graph in which the mean distance between two random nodes increases logarithmically with the number of vertices. Such networks are considered *small worlds*, because it is possible to connect any two vertices in the network through just a few links. The graphs used in the present experiment are "re-wired ring lattices" (Watts and Strogatz, 1998): starting with a ring in which each node is connected to exactly two neighbors, long-range links are added to a few pairs of randomly chosen nodes (Figure 2). The communication network embodies principles well-known to participants from online networks such as Facebook.

Decodable message content: In practice, the language used by participants is simple and easy to decode automatically. Messages could be analyzed as either *Information Requests (Item)* ("I need a towel!") or *Fact (Item, Location)* ("The cat is in Pittsburgh"). There was little opportunity to aggregate or interpret data ("It seems that there are many household items in the Western region."). A *time-to-answer* analysis can match requests to facts.

Individual and collective payoffs: Players move around the road network until they find a city that can provide the item. Players are rewarded for finding a goal item with $r_0 = 1000$ game points. As a further incentive to not only ask for information, but also provide useful information to their contacts, we reward each participant (node x) with game points whenever a node y to which they are connected obtains a goal item. The reward is highest for immediate neighbors. It is $r_d = \frac{r_0}{2^{d(x,y)}}$ for all y for which $d(x, y) \leq 3$. This reward distribution system follows common pyramid schemes and also Pickard et al. (2010). The final game score of a player is the sum of all rewards obtained during one session.

Task success metrics: The game was designed to give measurable task success metrics, where the communication within the player network would be critical to task success. These include the accumulated payoff (the communicated objective), the number of goal items found, and the average time it takes to receive an answer to one's information request.

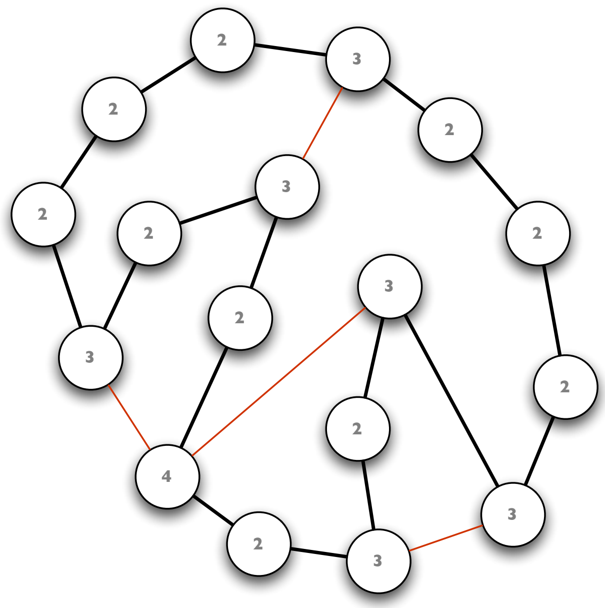

Figure 2: The communication network assigned to one of the experimental groups (17 subjects). The initial links from the network creation process are shown bold and black; the other ones (red, thin) can be considered long-range connections. Nodes are labeled with their degrees.

4. Method

34 participants from our university communities each played the Geo Game in computer labs for reward. They were split into two groups with 17 participants each. Each group formed a team. Participants were briefed about the game through introductory screens. They were instructed to score as many points as possible. It was emphasized that they benefited both from finding items themselves, as well as from helping those that they were connected to via the chat interface. Experimenters proctored the session to ensure that they did not communicate with other participants outside of the chat interface or view their screens. Each group played a 10-minute familiarization phase, and then two 30-minute experiments in the conditions outlined below. The order of conditions was changed between the two groups to control for learning effects.

The experiment had two conditions. Participants were instructed to adapt their communication strategy: In the first condition (*dump*), participants were asked to indiscriminately broadcast a maximum of information available. In the other condition (*target*), they were asked to request and target information so that only such knowledge was disseminated (and passed on) that was known to be relevant to others in the network through prior requests.

Data from one participant in the *dump* condition was ex-

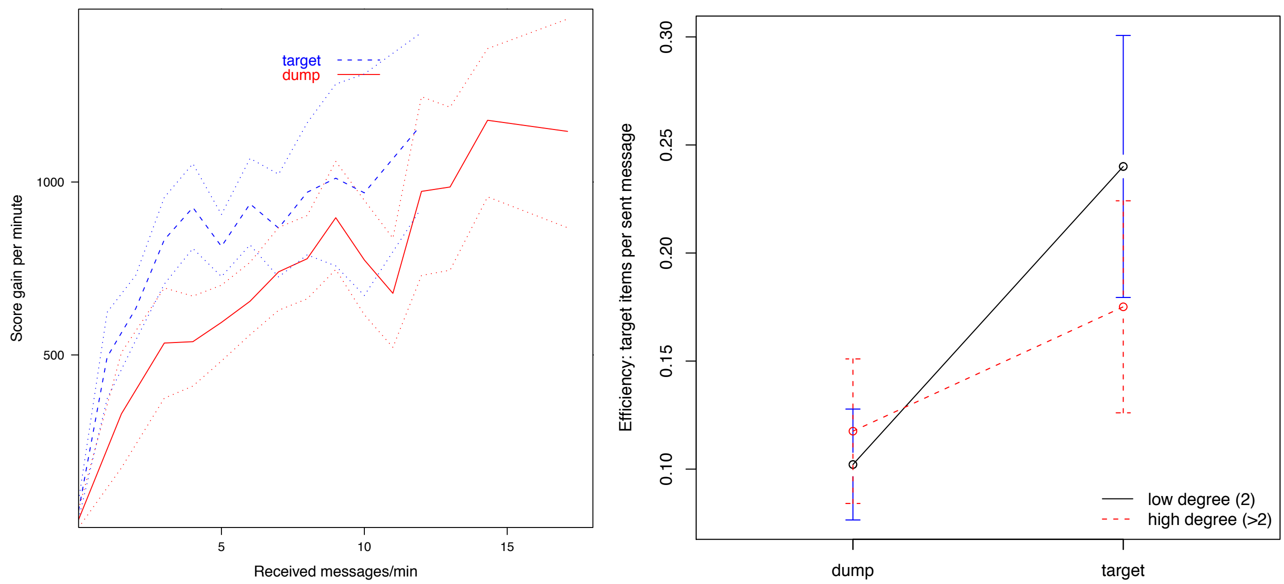

(a) Score gain per minute

(b) Efficiency gain through targeting

Figure 3: (*a*) Score gain per minute vs. number of messages received: At the same numbers of received messages, communities gained a higher score in the *target* condition than in the *dump* condition. (Dotted lines show 95% confidence intervals, assuming bins per user for each minute elapsed and gaussian distributions). (*b*) Sender-side efficiency gain through targeting: low-degree nodes (those with only two connections) vs. others.

cluded, as they were assigned no network neighbors by the random network generation. Note that this participant reached only 4,000 points compared to a mean of 23,660 points (std.dev 6,001) for all other participants, providing anecdotal evidence for the claim that communication is essential to task success in this game.

5. Results

In the analysis of the results, we focus on three measures. First, we compare the overall task success as measured by *game score* and *goal items found*, analyzed by participant. The game score is the objective that was communicated to the subjects. The game score depends on the number of network neighbors and the overall connectivity of the node, as each node was assigned additional points when network neighbors found items. The total number of goal items found by a participant, on the other hand, is not confounded with neighborhood size (degree).

Second, we analyze *communication efficiency*, which is the number of goal items found relative to the number of messages received.

Third, we describe the *time* it took for each request to be answered. Requests (for example, “Need cake”) and answers (“Cake in Munich!”) were annotated using keyword and punctuation spotting, with keywords comprising goal items and locations. The coding was manually checked on samples. We matched requests and responses by identifying requests for items sent by a node, and subsequent answers with the item and a location received by the same node.

5.1. Targeting increases task performance

The first question we ask is whether the information propagation policy influences task performance in the community. If either *dump* or *target* condition prove substantially advantageous, then the manipulation would appear to interact with the specific task. Under the assumption that humans can effortlessly integrate information, we would expect that the team directly benefits from unlimited communication in the *dump* condition. This is not the case: Subjects score 2,653 points higher when targeting (mean score in *dump* : 22,367), which amounts to a trend ($p = 0.073$, t-test of linear regression coefficients). For the basic summary statistics of scores, goal items taken, and messages exchanged, refer to Table 1.

The basic premise of the Geo Game task is that it relies on communication. An individual’s success should depend on their communications. Figure 3(a) shows task success (score gain per minute) as a function of the number of

Cond.	Degree	Score	Received Msgs	Msgs Sent	Items Taken	n
dump	> 2	27455	299.6	64.4	6.36	11
target	>2	28000	166.9	40.8	6.5	14
dump	2	19512	198.6	68.7	6.45	20
target	2	22576	111.9	37.3	7.1	18

Table 1: Performance and communication data: Means over participants in the two conditions, split into low-degree and higher-degree sets.

messages received, per minute. A correlation here is predicted, given the experimental design: more messages lead to higher task performance, at the level of individual nodes. However, the network structure moderates the opportunity to communicate. To differentiate the effects of communication and network location, a linear regression model was fitted to the participant data with the score as response, and as predictors the condition, numbers of received and sent messages as well as the node's degree[1]. Another model was estimated to predict the number of items obtained, with the same predictors.

The more messages a participant received, the higher their final game score ($p < 0.0005$) and the more items they seem to find ($p < 0.075$). Targeting increased participant's scores substantially. Once the main effects of message *quantity* and network position are regressed out, we observe strong benefits from a communication policy on residual task performance, i.e., we see an effect of message *quality*. Participants scored an average of 6,930 points more in the *target* condition than in the *dump* condition ($p < 0.0005$), and appeared to find 1.28 more items ($p < 0.08$). This means that while an increase in communication is generally good for performance in this task, indiscriminate dumping of information is detrimental relative to communication targeting. Participants fail to make efficient use of the dumped information, and individual nodes lack the cognitive capacity to utilize all of it. Thus, targeting at the network level contributes substantially to the knowledge of the community. A failure to target at the network level is not effectively compensated by individual participants. This difference is not an effect of the task, which could simply reward participants for avoiding communication. As the case of a non-communicating participant shows, communication is necessary, and the absence of it also severely hurts performance. Thus, team performance is highest when participants tailor their information to their communication partners.

[1]All models: Linear Mixed-effects regression. Predictors were centered. A random intercept was fitted grouped by session and subject. P-values for the linear models calculated with Markov-Chain Monte-Carlo sampling. Where the number of items obtained is predicted, generalized regression with a poisson transform was used. A preliminary regression model did not find significantly different scores or items-taken between the two experimental groups, and across both conditions.

5.2. Targeting increases efficiency

Communication in this experiment comes at a small cost: it takes time to send and read messages. In real-world tasks, communication may be more expensive, both in terms of time and attention. The tradeoffs are specific to the tasks. This study can, however, speak to the effects of human networks and network communication policies on *communicative efficiency* in more general terms. We operationalize efficiency as the number of items obtained per message received (overall mean $\mu = 0.044$).

A linear regression model was fitted to participant data, predicting the total number of items taken divided by the total number of received messages (response, henceforth *efficiency*), with condition, node degree and number of received messages as predictors, and random intercepts grouped by session and participant to account for repeated measures (see Table 2).

Efficiency was higher during targeting than dumping ($\beta = 0.014, p < 0.05$). Efficiency decreased slightly with the number of received messages ($\beta = -0.0002, p < 0.0005$). I.e., independently of whether participants were asked to target their messages or not, participants did not benefit proportionally from a large number of messages. Controlling for this small effect by including it in the regression model, we find that the efficiency increase in the *target* condition is not explained solely by an overall lower number of messages. Thus, the *target* manipulation affected message quality. Participants increased the relevance of their messages and/or paid more attention to each message. Figure 3(a) shows how the effect of communication policy holds over the full range of message productivity, taking the score (instead of items obtained) as response measure.

These results are further corroborated by our time-to-answer analysis. Here, a regression model was fitted to message data, predicting time-to-answer (seconds) as a function of condition and requester's degree.[2] In the *target* condition, participants received answers, on average, 49.5 seconds earlier than in the *dump* condition ($p < 0.05$).

[2]Linear mixed-effects regression. Random intercept for session and requester ID. Predictors centered. P-values by Markov-Chain Monte-Carlo Sampling.

Covariate	β	SE	p_{MCMC}
Linear Model predicting Items Taken / Rec'd Msgs.			
Intercept	0.0440	0.0060	< 0.0001
Condition	0.0137	0.0064	< 0.05
Degree	-0.0098	0.0137	< 0.005
Received Msgs	-0.0002	0.0033	< 0.0001
Cond.:Degree	-0.0190	0.0100	< 0.07

Table 2: Condition was coded as -0.5 (dump), +0.5 (target). The model contains effects for residualized Condition (against Received Messages and Degree), conservatively attributing effects to the number of received messages and the degree; it also contains the residualized Received Messages (against Degree). Centering and residualization were used to control for collinearity; correlation coefficients between relevant predictors were, as in all models in this paper, lower than 0.2.

5.3. Leaf nodes benefit more from targeting

While all participants carry out the same task, some of them may benefit from an improved network position. Naively, one would expect that a *target* -like communication policy would benefit especially the network hubs, through which many messages need to pass.

For a time-to-answer analysis, we fitted a model similar to the one in Section 5.2, but added an interaction between the condition (*target* vs. *dump*) and node degree (2 to 5). There was no significant interaction between these variables in the model, indicating no evidence for a difference in time-to-answer benefit from targeting for low or high-degree nodes.

For the analysis of efficiency, we turn again to the model described in Section 5.2, which included the degree and its interaction with the condition. To recall, we find increased efficiency in the *target* condition. Comparing different network positions, we see that the leaf nodes that have few connections benefit more than the well connected ones: The degree (2-5) shows a trend to interact with the main effect of condition ($\beta = -0.019, p < 0.085$), indicating a possibly decreased efficiency advantage from targeting with each additional neighbor a node has. Task success measured as items taken is consistent with this analysis (Table 1). When targeting, low-degree nodes collected more items than high-degree nodes, while when dumping, degree did not matter as much.

The difference in efficiency with the two communication policies also changes markedly for the message senders. Figure 3(b) shows the sender efficiency (items obtained by messages sent) of degree-2 nodes compared to the aggregate efficiency of higher-degree nodes, again suggesting a diminished efficiency advantage for high-degree nodes.

6. Discussion

Human networks can filter information and direct it to where it is needed, provided that humans are encouraged to make decisions about local information distribution. Communication processes are not necessarily mechanistic: people can pay attention to the informational needs of dialogue partners if asked to do so. In networks and in a situation where individuals communicate one-to-one, this appears to be beneficial to their task success.

The increased task performance suggest that "targeting" of communications accommodates attentional limitations at the cost of losing information. Given memory decay and interference, information left on just a few nodes risks being forgotten and has to be rediscovered at cost. Conversely, decreased targeting of information may improve the life and utility of information by spreading its availability and hedging against its decay, at the cost of attentional overload and interference with more important knowledge. Finding the optimum requires a computational account of human memory and attention.

We focus on the case where humans act as sensors and communication relays. The experiment provides no condition where nodes pull information directly from across the network: the game would be trivial otherwise. This assumption is sensible, as realistic knowledge is stochastic, noisy and also biased (based on interpretation): knowledge is more complex than "the cheese is in Paris". Future experiments will ask nodes to integrate noisy data.

The experiment is designed such that participants seek out information that originates from places far beyond their network neighborhood. Thus, the results show that careful communication practices impact the information state and attentional demands not just among interaction partners, but also further away in the network. The clearest effects of the manipulation were observed in the efficiency measures. Greater efficiency at the local level translated to greater efficiency at the game level; this was, to some extent, by design. However, we also observed greater communicative success when participants increased local efficiency by targeting. This pattern of results would not have been predicted by a model in which human participants have sufficient capacity to process all incoming information while playing the game (they would have benefitted from additional messages, even if their average usefulness was lower). Thus, participants employed targeted communication, or *audience design,* successfully to overcome attentional limitations of their interaction partners.

An important consideration is the interaction of communication policy and network position. Our results suggest that well-connected people (hubs) derive less benefit from a selective message-passing policy. Conversely, openness in communication may be less useful for leaf nodes. This may have to do directly with the attentional or memory limitations of hubs, who fail to pass on relevant information when their bandwidth is exceeded. Yet, they act selfishly by making use of all received information.

7. Conclusion

With the Geo Game, we have presented a new experimental paradigm to investigate group performance in communities of human agents. The game is neither limited to a single group objective (as in joint problem-solving), nor do humans act as individual, adversarial agents. Instead, individuals benefit both directly and indirectly form the efforts of others. The experimental platform allows us to observe performance and communication efficiency.

Teamwork benefits from communication policies that make use of active information filtering by humans. What such policies should be, however, is subject to a system of interacting constraints: individual cognitive limitations and objectives, payoff distribution, network structure and the individual's position within the network. We point out some of the quantitative relationships between those variables. Human content selection strategies strike a balance between the need to pass on information quickly while remaining relevant. Human Geo Game players do very well with the *target* strategy, which perhaps is a natural or acquired communication maxim.

Acknowledgements

This work was funded by the Air Force Office of Scientific Research (MURI grant FA95500810356).

References

Alexander Bavelas. 1950. Communication patterns in task-oriented groups. *J. Acoustical Soc. America*, 22:725–730.

Siddhartha Bhattacharyya and Stellan Ohlsson. 2010. Social creativity as a function of agent cognition and network properties: A computer model. *Social Networks*, 32(4):263–278.

J. A. Cannon-Bowers and E. Salas. 1997. Teamwork competencies: The interaction of team member knowledge, skills, and attitudes. In Jr. H. F. O'Neil, editor, *Workforce readiness: Competencies and assessment*, pages 151–174. Lawrence Erlbaum, Mahwah, NJ.

H. Guetzkow and H. A. Simon. 1955. The impact of certain communication nets upon organization and performance in task-oriented groups. *Management Science*, 1(3-4):233–250.

S. Judd, M. Kearns, and Y. Vorobeychik. 2010. Behavioral dynamics and influence in networked coloring and consensus. *PNAS*, 107(34).

J. Kennedy. 1999. Small worlds and mega-minds: effects of neighborhood topology on particle swarm performance. In *Proc. Evolutionary Computation*.

Bryan Klimt and Yiming Yang. 2004. The Enron corpus: A new dataset for email classification research. In Jean-François Boulicaut, Floriana Esposito, Fosca Giannotti, and Dino Pedreschi, editors, *Machine Learning: ECML 2004*, volume 3201 of *Lecture Notes in Computer Science*, pages 217–226. Springer Berlin / Heidelberg.

H.J. Leavitt. 1951. Some effects of certain communication patterns on group performance. *Abnormal and Social Psychology*, 46:38–50.

W. A. Owens. 1995. The emerging system of systems. *Military Review*.

Galen Pickard, Iyad Rahwan, Wei Pan, Manuel Cebrián, Riley Crane, Anmol Madan, and Alex Pentland. 2010. Time critical social mobilization: The DARPA network challenge winning strategy. *CoRR*, abs/1008.3172.

Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.

Author Biographies

DAVID REITTER is a research psychologist at Carnegie Mellon University. His research interests concern computational psycholinguistics as well as communication, collaboration and emergent behavior of human and mixed human-computer groups.

KATIA SYCARA is a professor of computer science and is interested in large man-machine systems.

CHRISTIAN LEBIERE is a research faculty in the Psychology Department at Carnegie Mellon University. His main research interests are computational cognitive architectures and their applications to psychology, artificial intelligence, human-computer interaction, decision-making, intelligent agents, robotics and neuromorphic engineering.

YURY VINOKUROV is a research programmer in the Psychology Department at Carnegie Mellon University and works on cognitive architecture.

ANTONIO JUAREZ is a PhD student in computer science at Carnegie Mellon University.

MICHAEL LEWIS is a professor of Information Sciences and Intelligent Systems at the University of Pittsburgh. His research focuses on human interaction with intelligent automation.