Social Cognition: Memory Decay and Adaptive Information Filtering for Robust Information Maintenance

David Reitter and Christian Lebiere

Department of Psychology, Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
`reitter@cmu.edu cl@cmu.edu`
Content areas: memory and knowledge, natural language processing, dialogue and interaction

Abstract

Two information decay methods are examined that help multi-agent systems cope with dynamic environments. The agents in this simulation have human-like memory and a mechanism to moderate their communications: they forget internally stored information via temporal decay, and they forget distributed information by filtering it as it passes through a communication network. The agents play a foraging game, in which performance depends on communicating facts and requests and on storing facts in internal memory. Parameters of the game and agent models are tuned to human data. Agent groups with moderated communication in small-world networks achieve optimal performance for typical human memory decay values, while non-adaptive agents benefit from stronger memory decay. The decay and filtering strategies interact with the properties of the network graph in ways suggestive of an evolutionary co-optimization between the human cognitive system and an external social structure.

Introduction

The human cognitive system has evolved as a general mechanism to process information and store knowledge provided by our environment. Humans are social animals, and as such, their cognition may serve as a model of information processing in general multi-agent systems. A seemingly disadvantageous property of human cognition is forgetting: memory retrievals are contextualized according to recent exposure to knowledge. Similarly, attentional resources are limited, and humans will reduce their interactions with others when necessary. Can such filtering heuristics help multi-agent networks cope with dynamic environments? This study examines memory decay and communication filtering in a multi-agent simulation that uses cognitive models in place of simple agents.

In this paper, we use an interactive foraging task to analyze the interaction between human memory and a network of peers in a multi-agent simulation. The empirical version of the experiment (Figure 1, Reitter et al. 2011) serves to parametrize and motivate a cognitive model, which is an algorithm bounded by human processing limitations. The

model allow us to manipulate network structure, individual memory decay, and communication strategies in a system consting of multiple independent, networked copies of the model. We observe the interaction of these variables with task success and accuracy in passed messages.

Rational Cognitive Architectures

We formulate our model within a cognitive architecture, that is, within a principled set of human-like bounds. Such architectures combine a number of goals. They represent the invariant properties of the human mind, unifying and standardizing the basis on which cognitive models can operate. A model describes sensory processes, thought, memorization, and action-taking within a defined task. It predicts and explains reaction times and an experimental participant's decisions and the associated learning effects and bounds on performance, but also neurophysiological measures such as activation found in different brain regions. But beyond the scientific explanation, cognitive modeling can leverage its faithful representation of human abilities and limitations to implement human-like systems, or complex systems with emerging intelligence. The design of cognitive architectures can be motivated by many factors. One intriguing idea has served as a remarkably useful default assumption: that the environment contains clues to explain invariant architectural properties of the mind. The mind's design is *rational*, that is, adaptively addresses the needs imposed by the structure and statistics of our environment (Anderson 1991; Oaksford and Chater 1999).

Described this way, the rational assumption glosses over the fact that most of the available information is created by ourselves. Some of the fundamental principles underlying the structures of society and information sharing must be encoded in the mind. For instance, the predominant type of network structure found in social networks (such as graphs of friends, collaborators, lovers, business relationships) has "small-world" properties, which have been proposed to be created using the preferential attachment principle: each individual agent prefers to create connections with those who already possess a large number of connections (Simon 1955; Barabasi and Albert 1999).

An example of a cognitive property is the decay rate for information stored in declarative memory in the cognitive architecture ACT-R (Anderson 2007). Activation, a proxy
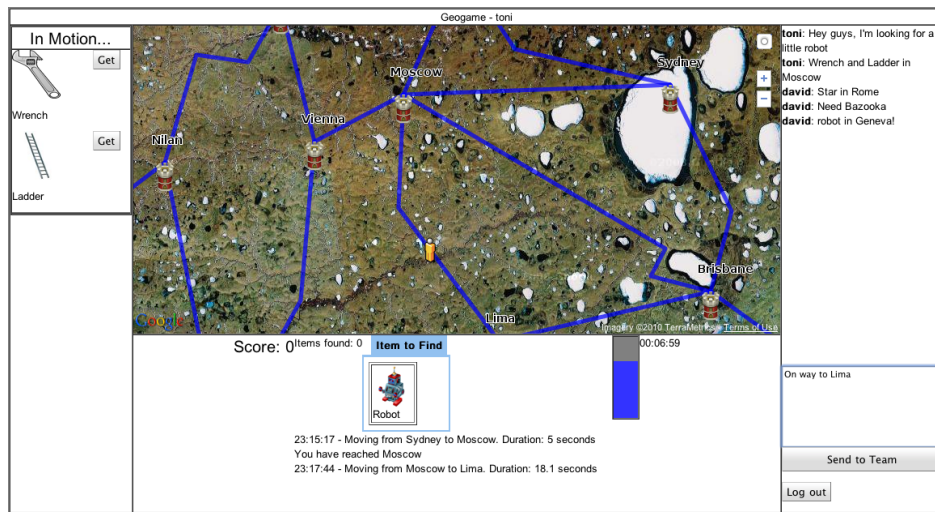
Figure 1: The Geo Game interface for human players.

for the log-odds of a piece of information being needed, is determined as a function of the time elapsed since past presentations. That is, it decays at a certain rate. The parameter dictating this rate has been remarkably steady in the literature. From a rational perspective, the decay can be seen as an adaptation to an average rate of change in the environment (Anderson and Schooler 1991). The change rate is, however, by no means the only environmental property that may have influenced human memory. Humans are surrounded by mnemonic devices: visual ones for example, and also networks of (human) agents that retain their own memories and communicate with each other to relay them. In short, these networks serve as buffers and filters in relaying information.

Humans overcome their own limitations as well as environmental shortcomings, not just through the evolutionarily-determined, invariant cognitive properties, but also dynamically based on experience. *Metacognition* refers to an ongoing monitoring loop that leads to success-driven choice of strategies. The study presented in this paper will show how decay of memories as well as metacognitive adaptation of communicative behavior can optimize the spread of information through a communication network.

We see decay, metacognition and adaptive communication as the result of an adaptation to the environment. Where the environment is man-made (as is the case for naturalistic communication networks), we expect co-adaptation on network structure and memory decay. We will pay special attention to the performance of our models with naturalistic (small-world) network topologies and with decay parameters near the values documented in the cognitive modeling literature.

Background

Information foraging has been studied from a cognitive perspective in a number of environments, such as information retrieval systems (Pirolli and Card 1999). Pirolli and Fu (2003) developed an ACT-R model of information foraging

that reflects the constraints and mechanisms of our cognitive architecture, especially the subsymbolic processes such as spreading activation. Pirolli (2005) developed a rational analysis of information foraging on the web to study the optimality of our strategies given known cognitive limitations. While that approach to information foraging focused on a single cognitive agent, Bhattacharyya and Ohlsson (2010) simulated the results of decomposing a complex task and assigning it to a network of cognitive agents. Performance of the network reflects differing individual characteristics of the agents such as working memory. Recent efforts have also attempted to include human-like characteristics in networks of artificial agents (e.g., Paruchuri et al. 2010).

The specific biases and communication modulation we are concerned with relate to a depreciation of information over time; this property has been observed in human-produced data and leveraged by a number of algorithms. Effects at many levels of representation show such forgetting, e.g., see Klinkenberg and Renz (1998) for concept drift in categorization, and Jelinek et al. (1991) for cache models in automatic speed recognition, and Reitter, Moore, and Keller (2006) for the use of recency-based adaptation in predicting task success. Coman et al. (2012) combine aspects of networked communication and forgetting in a cognitive simulation, showing the effects of group size and homogeneity of initial attitudes on the adoption of ground truth by an agent community.

The Geo Game

The Geo Game is an interactive simulation that involves a network of participants who communicate with each other. This design provides an experimental model of human communities (Figure 1), where information may spread from peer-to-peer by word-of-mouth, and an equivalent agent-based simulation. It is intended as a model of real-world cooperative foraging tasks. Communication paths are defined by the edges of a social network graph; each participant may

broadcast to their network neighbors by typing short messages and read the messages sent by those neighbors. Participants are asked to engage in a foraging task: find named items that are hidden in virtual cities shown on a map that connects these cities with roads. The task requires them to scout the area; it is more efficient if they have information about where their items are located. Thus, the task operationalizes facts as city-item associations; messages typically contain either requests for the location of an item, or facts. When an item is taken at a location, it is re-created at a different, random location; this implements a dynamic ground truth and invalidates existing facts. The number of items taken is our primary measure of task success.

Previous work has examined the effects of communication strategies in a study with human participants. Studies of local communication policies in the Geo Game have shown the efficiency and the performance of teams and of individuals in different positions within a network (Reitter et al. 2011). 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. These experiments forced a realistic tradeoff between directing attention to communication and to exploring the world: both activities lead to information gain, but are associated with attentional costs. In one condition (*target*), participants were asked to request specific information and only supply information that they knew was needed. In another condition (*dump*), they were asked to supply and forward as much information as possible. Participants did not benefit from strategies that increased communication volume, compared to a targeted strategy where they request information as needed, pass on requests, and judiciously reply and forward based on needs. These empirical results form the basis for the multi-agent simulations we discuss here, and in some analyses we will adopt the dump/target contrast.

Model

We model the Geo Game as a multi-agent simulation, where each agent is implemented as a cognitive model in the ACT-R framework (Anderson 2007). Each agent moves from city to city in the Geo Game world, looking for its goal item. Agents regularly read incoming messages and decide whether to forward them. When looking for a new goal item, they request the help of their network neighbors via the communication system, just like humans do ("Where is the cake?"). Agents acquire facts from experience by moving between cities. They broadcast facts as they are learned, provided they see a need for the facts to be known by their network neighbors (i.e., a request was received earlier). They may also re-broadcast facts and requests to pass them on across the network. They always do so if they know of a need for the fact, but in addition, a stochastic choice is made to pass on these facts even if there is no known need. This dampening may be meta-cognitively regulated. We generally avoid broadcasting the same fact or request multiple times within a short time frame to preserve attentional resources or bandwidth. These strategies serve to implement

basic human communication behavior: being relevant, designing for the audience as far as memory allows, and being economical in one's communications.

The model stores facts about item locations having learned about them from others, or having observing them locally. Retrieval is governed by a learning function:

$$B_j = \ln \sum_i^n t^{-\alpha} + \beta$$

B_j is the base-level activation if memory item j, which determines (via a threshold function, and modulated by noise) whether j can be retrieved from memory. t_i indicates the time since fact j was presented (at the ith occasion of all n prior presentations). Presentations occur when the agent becomes aware of the new fact, but also shortly afterwards, while the information is still visible on the screen and if the agent finds spare time for rehearsal. A minimum activation is needed to allow retrieval. Retrieval is a stochastic process due to the noise parameter β.

This model, with standard parameter values ($\alpha = .5$), and given the same time frame as used in the empirical literature, predicts performance at or somewhat below the empirical levels for groups of 20 subjects and small-world graphs. In these experiments, humans primarily used simple natural-language messages to efficiently convey the two types of semantics; we tune the model to display comparable messaging rates. The model is implemented in ACT-UP, a scalable and rapid-prototyping implementation of a subset of ACT-R (Reitter and Lebiere 2010).

Simulations

We were interested in optimal values for α, and whether those optima vary with the network topology. We contrast various types of network topology, all of which are defined initially and remain unchanged throughout each trial of simulated 30 minutes:

- small-world networks with a mean degree of d,
- Erdös-Renyi random graphs with a mean degree of d,
- tree hierarchies with a branching factor of $d - 1$, and,
- as a control, unconnected networks (representing no communication between agents).

All of the multi-agent experiments are Monte-Carlo simulations that randomly choose a base-level learning exponent ($0.05 \leq \alpha < 0.85$), a mean degree of the network ($|N| * 0.08 \leq d < |N| * 0.28$, where $|N| \in 25, 250$ is the number of nodes in the network). Plots in this paper show data from the networks large enough to give stable results ($|N| = 250, \bar{d} = 27.3$). We randomize the initial ground truth and agent states (locations of agents) for each trial. The first set of simulations does not enable any metacognitive decision-making; messages are passed between agents only when they see a need (target strategy).

Decay in Individual's Memories

If temporal decay of retrieval probability is a rational adaptation to the environment, we would expect the model to perform particularly well at tasks requiring memorization for
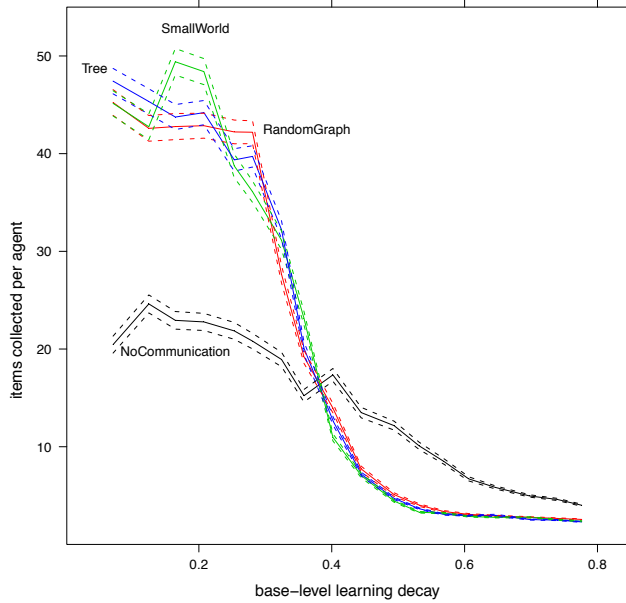
Figure 2: Task performance (items collected per agent) for different network types over a range of base-level learning decay values. The version of the game used here gives models ample time to communicate while changing location. 95% confidence intervals (uncorrected for multiple comparisons) are shown, which were obtained by bootstrapping.

decay variable settings that are realistic with respect to decay commonly used in cognitive models. The dynamic environment provides a changing ground truth, but, as we aim to show here, it also consists of a network of individuals. As a consequence, we would expect that realistic temporal decay of the accessibility of memorized facts adds a performance benefit for typical social network structures.

Figure 2 shows the effects of decay on performance, and its interaction with communication structure. Decay seems to negatively impact performance across the board. At a realistic decay of 0.5, performance is comparable to the performance that was obtained empirically.

However, examining task performance in more detail shows a bimodal distribution, with peaks at about 4 and 92 items taken. Some agents achieve extraordinary performance, whereas we found no observable correlation with the node's degree in the communication network (size of neighborhood). As any means-based analysis is strongly affected by outliers, we exclude results with performance ≥ 30 (the top 17%) in Figure 3. We do so in this case only. Here, we see a substantially changed effect of base-level decay: decay improves performance by up to a factor of 2 for communicating and non-communicating agents. This improvement slows down for decay values of around $\alpha > 0.5$, provided agents communicate. For low and high decay values, the benefits of communication seem diminished.

A modicum of decay seems to purge outdated information from memory. The amount of decay in humans is considered an innate constant and not controllable. These results lend
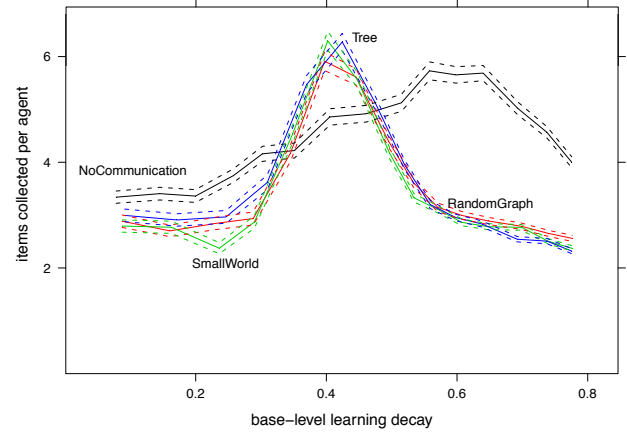

Figure 3: As in Figure 2, but excluding outliers (agents with performance of 30 or better).

no support yet for the idea that networks with more common structural properties (high clustering factor, or scale-invariance, or low betweenness, as in small worlds) show some form of co-adaptation with (human) decay. However, this effect is only seen for networked communication compared to the case of agents searching for themselves.

Among the questions arising from these data is the one referring to external validity. To what extent do these results depend on the set of parameters used in this simulation? Coman et al.'s (2012) recent results suggest co-adaptation of forgetting and the size of the social group. Running the same simulation in smaller teams (25 agents), we obtain similar effects. Small teams fare considerably better: their task performance measured in items obtained was approximately twice as high. The effect of memory decay and its interaction with network topology is consistent with that in the large communities.

A fundamental task for each individual playing the Geo Game is to decide how much to communicate, and how much to concentrate on visiting different locations. We see this as a trade-off between benefitting from a distributed sensor network (the other agents) and the cost of communication. On a functional level, this amounts to a decision between internalizing facts (storage in individual memory) and externalizing them (storage in the network). For this reason, the following experiments parametrize the Geo Game such that agents move ten times as fast between locations. This speed-up further limits the time available for communication, but also for individual rehearsal of knowledge. Figure 5 demonstrates the baseline performance of the agents in this round of simulations.

Metacognitive Communication Filtering

Humans can adapt their choice of action strategy upon explicit reflection, but also according to ongoing monitoring. This ability is referred to as *metacognition*. In Geo Game, communication bandwidth is limited. Agents have no a-priori information about their position in the network, or
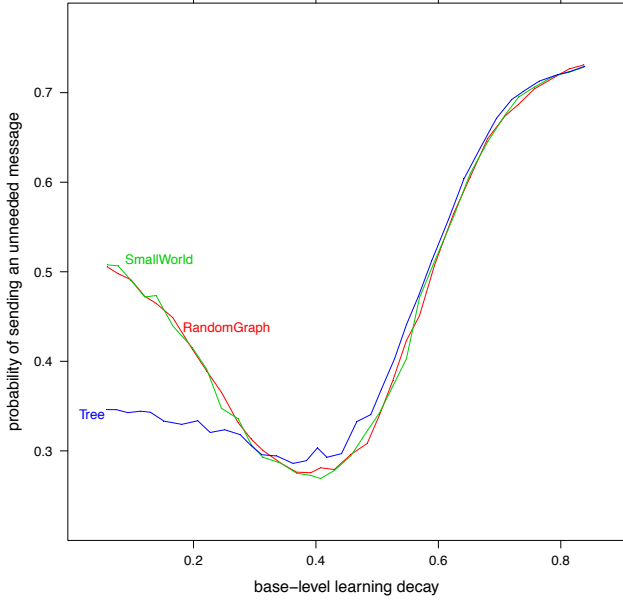
Figure 4: Agents send more unnecessary messages when internal memories decay too slowly or too quickly.

Figure 5: Agent performance after speed-up of the game (which limits communication bandwidth).

variables influencing communicative choices, such as trust for network neighbors. Communication is not universally helpful. For instance, Figure 4 shows the model's derived utility of communication across a range of decay values. When memory is particularly short-lived, the models benefit much more from communication.

In our agent-based simulation, we introduce a mechanism that adjusts the communication strategy dynamically based on perceived success. Its goal is to maintain an expectation s' of how the utility of communication. The metacognitive model will then send a message whenever it knows of a need; in addition, it may send it according to a stochastic decision function, i.e., with probability s'.

Instance-based learning (IBL, Gonzalez, Lerch, and Lebiere 2003) is a design pattern that stores episodic information of different categories in order to retrieve an *ideal* version of each category. For instance, an IBL model may experience task success in different runs, associated with different strategies. Later, it may determine an evaluation of each strategy according to its average task success. In our model, IBL stores an episode $e_i = < s_i, t_i >$ whenever an agent becomes aware of a useful fact after time period t_i, having used a certain communication strategy s_i (1.0 for communication, and 0.0 for no communication). That is, when the item's location was originally requested in a message sent by the agent and received in a message from another one, $s_i = 1.0$ and t_i indicates the time passed since the first request. When the item's location was discovered by exploration, $s_i = 0.0$ and t_i indicates time spent between item assignment and discovery. To decide how much to communicate, IBL produces a blended episode e', where s' is a weighted mean of the stored episodes. We aim to minimize the time it takes to obtain information about an item's loca-
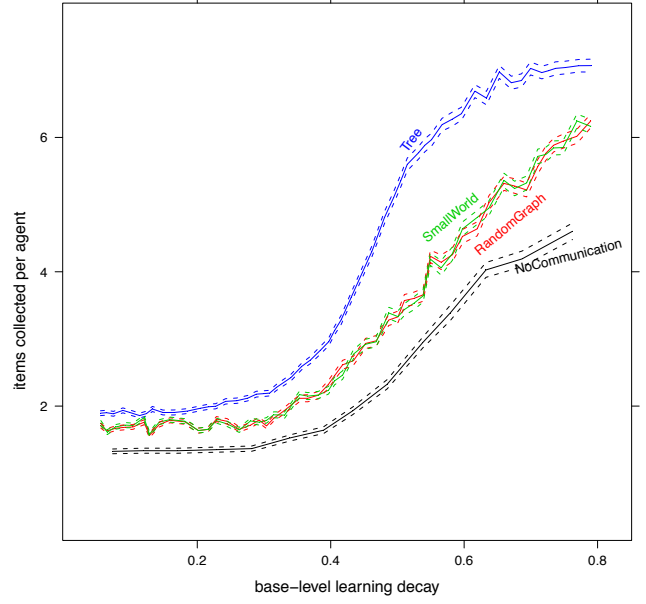
tion by requesting $t' = 0.0$ as part of the blending process. Thus, recent episodes are weighted more heavily as well as ones with a low t_i. Based on the calculated s', agents throttle communication.

We found that the metacognitive filter, overall, hurts performance (Figure 6). Yet, for small-world and random graphs, performance is enhanced at realistic (human) decay values (around 0.5). We also compare performance of metacognition against more extreme communication strategies (Figure 7): either targeting (only communicating when information is needed), or dumping (always communicating). We find that task success is enhanced for a range of realistic decay values (approx. 0.3–0.5). (This holds also true for smaller networks ($|N| = 25$) for decay values 0.3–0.6.)
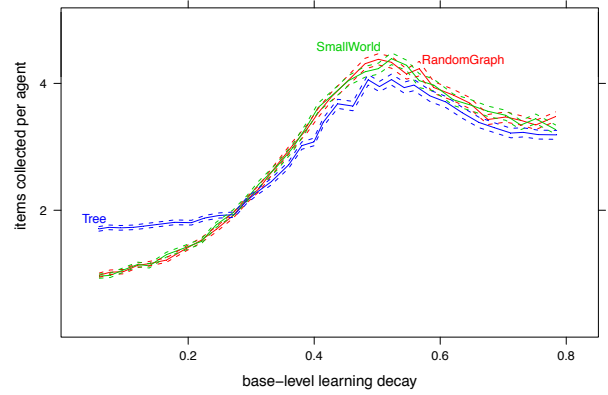

Figure 6: As Figure 5, but with throttling based on metacognitive monitoring of the utility of communication.
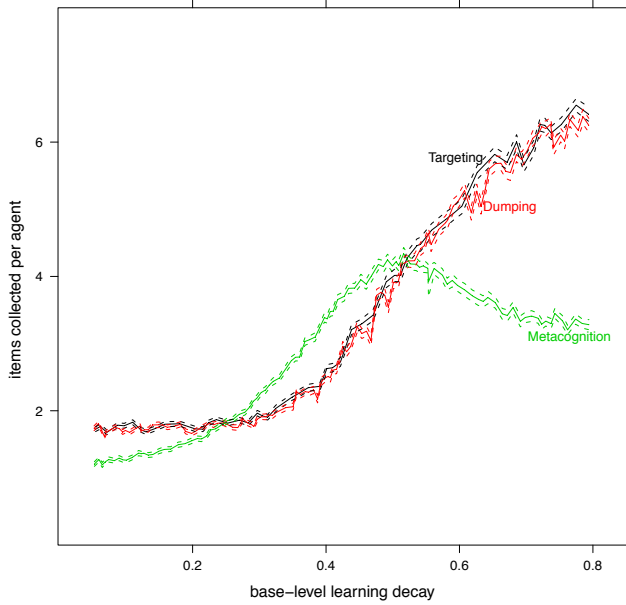
Figure 7: Metacognition beats maximally and minimally talkative agents at cognitive plausible values of individual decay, but it is outperformed for high decay values.
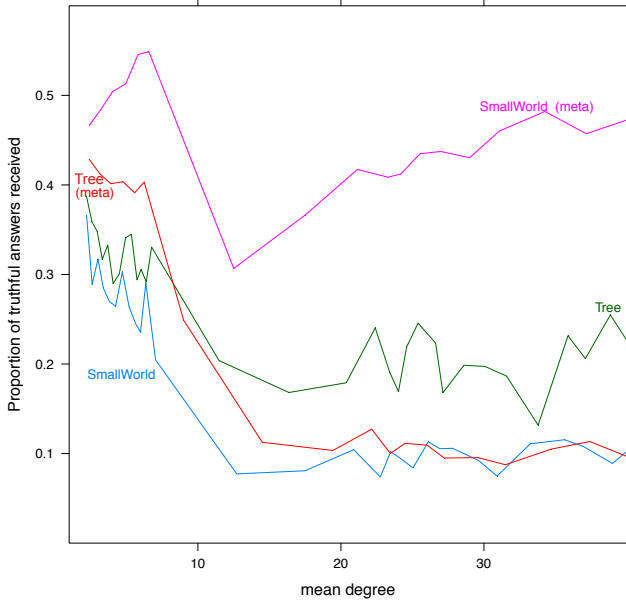

Figure 8: Metacognition allows agents to make use of larger network neighborhoods. (Random graphs behave like small worlds and are omitted for clarity.

Metacognition can be seen as an additional filter for the agents. In this game, agents have to choose between externalizing and internalizing information. For internally maintained knowledge, activation decay allows each agent to discard old and possibly outdated information. For externally maintained knowledge, metacognitive filtering can provide a

similar decay. At the network level, information is discarded when it is no longer passed on. When agents choose to not pass on messages based on a stochastic function, information decays rather than persists in the network.

The Geo Game uses the number of items taken as a measure of success. This can be seen as a proxy for the availability of information and integrates both the time spent waiting for information, and the accuracy of information obtained. Accuracy is particularly relevant in systems with changing ground truth. We define accuracy as the proportion of correct facts received by a node over the number of facts received overall. We find that metacognition substantially improves the accuracy of messages in the system for small-world and random graph networks; it worsens accuracy for trees.

The mean size of a node's network neighborhood (*degree*) contributes to the availability of information; however, can *popular* agents make better use of their large neighborhoods? Figure 8 illustrates that the smart choice of communication strategy lets agents make better use of their large neighborhoods. This does not apply in tree hierarchies. We find that a small-world network topology does perform best, but it requires selective filtering of messages. These simulations used randomly sampled BLL parameters (0.3–0.6). Random Graphs and small-world topologies showed very similar behavior. We believe that saturation of the network with messages in the non-filtered configuration may account for the similarity in random graphs and small worlds. Similarly, the task at hand does not necessarily benefit from the higher clustering coefficient in small worlds. The average degree that can affect attentional load in the agents was controlled across these two types of networks.

Discussion

The ideas we explore here are basic: can *forgetting* help multi-agent systems maintain a representation of true facts when ground truth changes? Our cognitive models carry out a foraging task; they are developed according to empirical observations of humans working on this task. We find that average task performance improves with increasing individual memory decay. This suggests that time-based decay is a useful heuristic for agents that allows them to discard outdated information. It is important to examine the interactions of forgetting with network topology, and to see how much forgetting is useful.

We show two simple ways to improve information transmission and flexibility in multi-agent communication networks with dynamic ground truths: temporal decay in each agent's belief state, and metacognitive reasoning to maximize communication efficiency. Both techniques are inspired by human cognition. They may be useful not just for simple agents, but also in complex computer-based or mixed human-computer systems.

Under the rational assumption that forgetting mechanisms have co-evolved with social structure, we expect that cognitively plausible forgetting is most useful for commonly found network topology with small-world properties. Stepping outside of cognitively bounded simulation, we can explore more extreme assumptions in order to determine useful strategies for multi-agent systems.

Time-based decay plays a role in the context of distributed memory. When agents communicate with each other, they externalize memory storage. As agents maintain more reliable knowledge, we find that decay particularly helps communities that have less clustering and relatively long distances, and that take many steps to propagate information: tree hierarchies (Figure 5).

Network topology does not always affect task performance. In our simulations, we find little differences between networks with small world properties and random graphs; however, tree structures reach at times higher performance. We observe, however, that when network-level forgetting is added, that is, when agents are judicious and adaptive in how much information they relay to others, small world networks outperform tree structures in one particular measure: the proportion of truth in the answers communicated. Real-life organizational networks re-configure their topology in time of crisis (via the Enron e-mail corpus, Diesner, Frantz, and Carley 2005); the higher accuracy of messages may explain some of the shortcuts taken by humans.

Network-level forgetting can be implemented in many ways. One seems particularly useful in the context of cognitive systems. When agents dynamically adapt the rate of communication according to their performance with and without it, then this can outperform non-dynamic strategies for certain levels of individual memory decay (Figure 7). This range of decay parameters (0.3–0.5) includes the decay level observed in empirical studies with humans (0.5). In this context, the decay parameter influences not just forgetting of facts. It also determines how much agents prioritize recent over longer-term experience in making their communication decision. Thus, more work is needed to determine the exact mechanism behind this effect.

A preliminary regression model predicting task success as a function of network parameters, decay and metacognition is not shown for reasons of space; however, it suggests that our findings are reliable. The cognitive models employed as agents exemplify a hybrid multi-agent system: they are constrained by human memory and communication bandwidth. In recent work, a model, even one of the same Geo Game paradigm (Wang, Sycara, and Scerri 2011), which was not constrained by cognitive architecture, has demonstrated disadvantages of collaboration. This was the case as communication in fully connected networks was moderated according to several non-adaptive strategies. According to our results, a number of properly parametrized mechanisms to discard old beliefs and to moderate agent communications have to come together to derive benefits from social cognition.

References

Anderson, J. R., and Schooler, L. J. 1991. Reflections of the environment in memory. *Psychological Science* 2:396–408.

Anderson, J. R. 1991. Cognitive architectures in a rational analysis. In VanLehn, K., ed., *Architectures for Intelligence*. Lawrence Erlbaum Associates.

Anderson, J. R. 2007. *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.

Barabasi, A. L., and Albert, R. 1999. Emergence of scaling in random networks. *Science (New York, N.Y.)* 286(5439):509–512.

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

Coman, A.; Kolling, A.; Lewis, M.; and Hirst, W. 2012. Mnemonic convergence: from empirical data to large-scale dynamics. In *Int'l Conf. on Social Computing, Behavioral-Cultural Modeling and Prediction*.

Diesner, J.; Frantz, T. L.; and Carley, K. M. 2005. Communication networks from the enron email corpus "it's always about the people. enron is no different". *Comput. Math. Organ. Theory* 11(3):201–228.

Gonzalez, C.; Lerch, F.; and Lebiere, C. 2003. Instance-based learning in dynamic decision making. *Cognitive Science* 27(4):591–635.

Jelinek, F.; Merialdo, B.; Roukos, S.; and Strauss, M. 1991. A dynamic language model for speech recognition. In *HLT'91 Workshop on Speech and Natural Language*.

Klinkenberg, R., and Renz, I. 1998. Adaptive information filtering: Learning drifting concepts. In *AAAI98/ICML-98 Workshop Learning for Text Categorization*, 33–40. Press.

Oaksford, M., and Chater, N., eds. 1999. *Rational Models of Cognition*. Oxford University Press.

Paruchuri, P.; Varakantham, P.; Sycara, K.; and Scerri, P. 2010. Effects of human bias on human-agent teams. In *IEEE/WIC/ACM Int'l Conf. on Intelligent Agent Technologies*.

Pirolli, P., and Card, S. K. 1999. Information foraging. *Psychological Review* 106:643–675.

Pirolli, P., and Fu, W.-T. 2003. Snif-act: A model of information foraging on the world wide web. In *Ninth Int'l Conf. on User Modeling*.

Pirolli, P. 2005. Rational analyses of information foraging on the web. *Cognitive Science* 29:343–373.

Reitter, D., and Lebiere, C. 2010. Accountable modeling in ACT-UP, a scalable, rapid-prototyping ACT-R implementation. In *10th Int'l Conf. on Cognitive Modeling*, 199–204.

Reitter, D.; Sycara, K.; Lebiere, C.; Vinokurov, Y.; Juarez, A.; and Lewis, M. 2011. How teams benefit from communication policies: information flow in human peer-to-peer networks. In *20th Behavior Representation in Modeling & Simulation*.

Reitter, D.; Moore, J. D.; and Keller, F. 2006. Priming of syntactic rules in task-oriented dialogue and spontaneous conversation. In *28th Annual Conference of the Cognitive Science Society*, 685–690.

Simon, H. A. 1955. On a class of skew distribution functions. *Biometrika* 42(3–4):425–440.

Wang, Y.; Sycara, K.; and Scerri, P. 2011. Towards an understanding of the value of cooperation in uncertain world. In *IEEE/WIC/ACM Int'l Conf. on Web Intelligence*.