Mnemonic Convergence: From Empirical Data to Large-Scale Dynamics

Alin Coman[1,⋆], Andreas Kolling[1,⋆], Michael Lewis[1], and William Hirst[2]

[1] University of Pittsburgh, School of Information Sciences, Pittsburgh, PA 15260
[2] The New School for Social Research, New York, New York 10011

Abstract. This study builds on the assumption that large-scale social phenomena emerge out of the interaction between individual cognitive mechanisms and social dynamics. Within this framework, we empirically investigated the propagation of memory effects (retrieval induced forgetting and practice effects) through sequences of social interactions. We found that the influence a public figure has on an individual's memories propagates in conversations between attitudinally similar, but not attitudinally dissimilar interactants, further affecting their subsequent memories [3]. The implementation of this transitivity principle in agent based simulations revealed the impact of community size, number of conversations and network structure on the dynamics of collective memory.

1 Introduction

Community identity and both intergroup hostility and cooperation often rest, at least in part, on the collective memories communities form of their past [9, 14]. We are interested here in using psychologically informed agent-based simulations to understand the dynamics underlying the formation of collective memory. We treat collective memory, following Hirst and Manier [7], as a representation of the past shared across a community that bears meaningfully on the identity of it's members. We want to understand the factors that affect the spread of a memory across a network of individuals and the convergence of the community onto a shared representation of the past. In this paper, we integrate the empirical data emerging out of the distributed cognition literature into an agent-based modeling framework [5, 8]. The assumption underlying this work is that macro-level social phenomena could emerge in predictable ways out of micro-level local dynamics.

In using the recent psychological work on social aspects of memory as our starting point for modeling, we are not only avoiding what Sun [13] has referred to as a shallow conceptualization of cognition in agent-based modeling, we are also extending the psychological work. This work mainly focuses on interactions in small groups, often consisting of just two or three people. The use of agent-based simulations that involve a large number of psychologically plausible agents allows us to overcome this limitation. In this paper, we focus on two psychological processes: practice effects and induced forgetting effects.

⋆ The first two authors contributed equally to the paper.

S.J. Yang, A.M. Greenberg, and M. Endsley (Eds.): SBP 2012, LNCS 7227, pp. 256–265, 2012.

We envision a community as a network of connected nodes. Each node possesses a memory, which can be transmitted to connected nodes while affecting memories of both the recipient and sender. We chose the two processes because they have received a great deal of attention to date [1, 4]. Future work can incorporate other psychological processes governing social influences on memory.

Although the model could be used to address a large number of substantive issues, we address the following questions here:

(1) Is the mnemonic convergence on a shared representation dependent on community size and number of conversations among individual agents? Inasmuch as groups differ radically in size, it is important to understand the consequences of these different group sizes on the ability of a group to form a collective memory. In recent years, some researchers have argued that, from an evolutionary perspective, core group configurations afford the evolution of specialized functions [2] and have highlighted group size as a pertinent variable [12]. They hypothesized, for instance, that bands or demes, usually with a group size of around 30, are adapted for the shared construction of reality. Following this lead, we expect that smaller networks would provide an appropriate context by which psychological principles of memory might promote swift formation of collective memory, while larger networks would not.

(2) How does the rate of convergence vary across different network parameters (e.g. network density)? There are good reasons to believe that network structure plays a role in the formation of collective memories. In particular, researchers have highlighted how network structure affects information transmission, and through it, collective problem-solving [11].

1.1 The Empirical Work Underlying the Two Processes

We mean by practice effects (PE) the finding that, when one is exposed to an event, and recalls it subsequently, information that is recalled is much more likely to be remembered later on than information that was not recalled [10]. Because this effect is well-established, we focus in this introduction on the other psychological process we plan to incorporate into our model: retrieval-induced forgetting (RIF). RIF has been extensively studied in situations in which a person initially studies material, then selectively practices some of the studied material, and then finally is asked to recall the initially studied material again. RIF occurs when people find it more difficult to remember memories unpracticed in the selective practice phase that were related to the practiced ones than those that were unrelated to the practiced memories [1]. Cuc et al.[4] refers to this as within-individual retrieval-induced forgetting (WI-RIF). Recent work established that the same pattern of induced forgetting occurs for those listening to other selectively practice. For instance, it occurs for listeners who attend to the selective recollections of speakers in a conversation. Such socially shared retrieval-induced forgetting (SS-RIF) is thought to occur because, on many now well-specified occasions, listeners concurrently, albeit covertly, recall selectively along with the speaker. This concurrent retrieval triggers the same processes in listeners as in speakers [4, 6].

In a recent study, Coman & Hirst [3] explored the propagation of practice effects and RIF through short sequences of social interactions, with the goal of understanding the formation of mnemonic convergence in small groups. They looked at how listening to a lecture on the legalization of euthanasia reshapes memories of learned material and whether the influence of the lecture propagates into a conversation and then through the conversation to a final recall test (see Figure 1). In the first phase, participants were exposed to arguments in favor and against the legalization of euthanasia, grouped into categories; each category contained two arguments in favor and two arguments against euthanasia. In a practice phase, similar to a lecture, participants were exposed to half of the arguments (only pro-euthanasia) from half of the categories. We call this phase Person-Pro exposure. After their individual memories were assessed for practice and induced forgetting effects triggered by Person-Pro, two participants were paired and were asked to jointly remember the arguments that they were exposed to initially. Subsequent to the conversation, in a final recall test, participants were asked to remember individually the arguments.

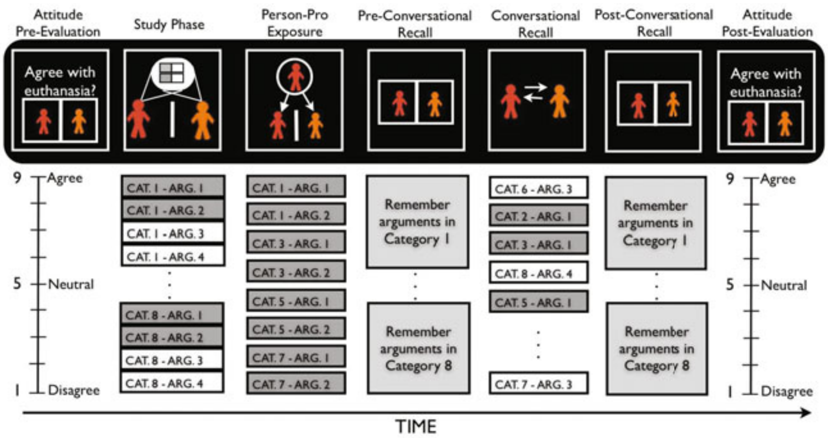

Fig. 1. The different phases of the experimental procedure in [3]. In the Study, Person-Pro Exposure and the Conversational Recall phases, darker shades indicate arguments in favor of the legalization of euthanasia.

Exposure to Person-Pro created three types of items: items mentioned by Person-Pro, which are always in favor of legalization (Rp+), unmentioned items related to the ones mentioned by Person-Pro, which are always against the legalization (Rp-), and items unrelated with the ones mentioned by Person-Pro, which are either in favor or against euthanasia (Nrp-pro and Nrp-anti). A practice effect due to Person-Pro emerges if the recall proportion of Rp+ items is greater than the recall proportion for Nrp-pro items. Similarly, an induced forgetting effect due to Person-Pro emerges if the recall proportion of Rp- items is less than the recall proportion of Nrp-anti items. Coman & Hirst [3] examined whether the practice and induced forgetting effects triggered by the lecture propagated into the conversations between participants, and through the conversation, into the final recall. They found that, when conversations were between like-minded individuals (both participants in the pair in favor of euthanasia or both against euthanasia), Person-Pro shaped what was remembered in the conversation. This influence further propagated into a final recall test, suggesting a lasting influence. Moreover, as practice effects and RIF effects propagate through the conversation to the final recall, they increased in size, thus strengthening with transmission.

The results indicate that mnemonic influences exhibit a principle of transitivity as they propagate through a sequence of social interactions. Importantly, mnemonic convergence around the information provided in the lecture emerged following conversations between like-minded individuals, even in situations where like-minded people disagreed with Person-Pro (see [3] for detailed analyses). Politicians can have a profound influence on what people remember, even when their listeners turn to each other to discuss the issue in an effort to remember the original material as best as they can. Coman & Hirst [3] explored how micro-level processing in a single social interaction shapes the emergent memories in a small sequence of social interactions. Their work suggests that macro-level phenomena specifically, the formation of a collective memory – could emerge in predictable ways out of micro-level dynamics. In what follows we want to extend their work to large networks, using agent-based simulations.

2 Agent Model

We designed an agent-based model that incorporates the sequence of interaction in [3]. In the study phase, we set the level of initial activation for the agents' memory. Agents were then exposed to Person-Pro. Subsequent to the exposure, agents were allowed to interact with each other. As a consequence of an interaction, activation was increased for items recalled during the interaction, and decreased for other items, with a greater decrease for those more closely related to recalled items. This differential decrease captured the induced forgetting effect. In the simulations presented here, activation updates were based on values obtained from empirical data collected by [3]. We first sought to determine whether our model produced results similar to those found by these researchers without introducing additional artifacts. Once satisfactory fitting was achieved, we ran simulations to explore whether mnemonic consensus is influenced by: 1) community size, 2) number of conversations among agents, and 3) the conversational network structure.

2.1 Model Details and Formalism

Let $\mathcal{A} = \{a_1, \ldots, a_n\}$ be a set of n agents. Agents are connected in a communication graph $G^c = (\mathcal{A}, E^c)$ with agents receiving edges whenever they are capable of exchanging messages. An agent a_i has a unique identifier and a memory model, i.e. $a_i = \{i, M_i\}$. The memory model $M_i = (X, f_i, G^m)$ consists of m_{size} items written $X = \{1, \ldots, m_{size}\}$, an associated activation function $f_i : X \to [0,1]$, and a weighted memory graph $G^m = (X, E^m, w^m)$ that describes the strength of connections between items in X with a weight function $w^m : E^m \to \mathbb{R}$. Notice that only f_i depends on the agent and X, G^m, w^m are identical across all agents. Patterns of activation given by f_i are the distinctive feature of an agent's memory. This simplifies the memory model and forces to relate all changes in memory across agents to changes in f_i. Two effects on activation levels are modeled whenever an agent is exposed to an item $x \in X$, namely a practice effect and an induced forgetting effect. To describe these we interpret $f_i(x)$ as a function that changes with regard to the number of practice and induced forgetting instances. As a shorthand, we write $f_x(n)$ for $f_i(x)(n)$ where n denotes the number of practice or forgetting instances. Upon exposure to a item $x \in X$ its activation f_x changed according to: $f'_x = \sigma(f_x)$, with f'_x denoting first order derivative of f_x and $\sigma : [0,1] \to [0,1]$ modulating the increase in activation. We require that $\sigma(a) \leq 1 - a$ and $\sigma(a) > 0, \forall a \in [0,1]$ to ensure that f_x remains within $[0,1]$. The choice of σ determines whether low activation items and high activation items are affected differently during exposure. To model induced forgetting, all neighbors of x in G^m, written $\bar{x} \in N(x)$ with $[x, x'] = e' \in E^m$, are subject to: $f'_{\bar{x}} = -w^m(e') \cdot \bar{\sigma}(f_{\bar{x}})$, where $\bar{\delta} : [0,1] \to [0,1]$, analogous to σ, modulates the decrease in activation. Notice that the magnitude of the induced forgetting effect also depends on the $w^m(e')$.

A conversation is now defined as the repeated recall and exposure of m_{conv} items from X for two agents a_1 and a_2. At first we have $X' = X$ and the probability of an item x being recalled by agent $a_i, i = 1, 2$, is:

$$P(x, a_i) = \frac{f_i(x)}{\sum_{j=1,2} \sum_{x \in X'} (f_j(i))}. \tag{1}$$

After an item is recalled both agents are exposed to it and it is removed from X'. Notice that an agent with a higher overall activation level will get to recall more items during a conversation. In addition to the simple conversational dynamics above we also consider biased conversations in which agents have an attitude bias given by $b_i \in \mathbb{R}^+$. This bias models the differences in attitude towards the recall of items in memory X. More precisely, two agents a_1 and a_2 in a conversation are said to have a different bias if $b_1 \geq \delta$ and $b_2 < \delta$ or vice-versa for some threshold δ. In this case, the biased activation level $f_b(x)$ for agent a_1 for a subset of items $x \in X_1 \subset X$ becomes: $f_b(x) = f_(x) + f_{bias}$. Similarly, for agent a_2 and a subset of items $x \in X_2$ the biased activation is given by: $f_b(x) = f(x) - f_{bias}, \forall x \in X_2$. Agent a_1 therefore mentions X_1 more frequently and a_2 mentions X_2 less frequently. This biased activation models changes in recall probabilities observed in experimental data as we shall see later.

3 Application of the Model

In this section we determine the parameters for our model to fit the experiment data from [3]. The memory graph G^m is set to represent the relationships between the items in each category. A total of eight categories with four items each leads to $m_{size} = 32$. Every four consecutive items are in the same category, i.e. $\{4i-3, 4i-2, 4i-1, 4i\} = C_i \subset X, i = 1, \ldots, 8$. Furthermore, every two consecutive items are Pro and Anti items, $X_{pro} = \{1, 2, 5, 6, \ldots, 29, 30\}$ and $X_{anti} = \{3, 4, 7, 8, \ldots, 31, 32\}$. With slight abuse of notation we define G^m as a weighted adjacency matrix:

$$G^m(i,j) = \begin{cases} w_1 + w_2 & \text{if } j \in C_{\lfloor i/4 \rfloor} \text{ and } i,j \in X_{pro} \text{ or } i,j \in X_{anti} \\ w_1 & \text{otherwise if } i,j \in C_{\lfloor i/4 \rfloor} \\ w_2 & \text{otherwise if } i,j \in X_{pro} \text{ or } i,j \in X_{anti} \\ 0 & \text{otherwise} \end{cases}$$

In essence, we interpret X_{pro} and X_{anti} as broader categories in addition to the connections within the smaller categories $C_i, i = 1, \ldots, 8$. To specify σ and $\bar{\sigma}$ we make two additional assumptions, namely that the practice effect decays exponentially and induced forgetting is linear, leading to: $\sigma(x) := \delta_l \cdot (1-x)$ and $\bar{\sigma}(x) := \delta_f$, which, enforcing the boundary condition $f_x(0) = 0$, is solved by $f_x(n) = -e^{-\delta_l \cdot x} + 1$, with e denoting Euler's constant. Finally, we determine values for δ_l, δ_f, w_1, and w_2 from experimental data, specifically the recall of RP+,RP-,NRP-pro, and NRP-anti items. After a study phase, all participants were exposed to Person-Pro discourse. Person-Pro recalls items from a specified list $\{4 \cdot i - 3, 4 \cdot i - 2\}, \forall i = 1, 3, 5, 7$. After this conversation, items of type RP+, RP-, NRP-pro, and NRP-anti, are recalled with probabilities $0.4537, 0.3199, 0.3067$, and 0.4366, respectively. In our model the recall for NRP-anti is unaffected by Person-Pro and hence the initial activation for all items, assuming uniform study effects, is $p_{init} = 0.4366$. The difference to p_{init} for RP+, RP-, and NRP-pro, allows us to determine the values of our parameters. After setting $w_1 = 1$ the others are completely specified by fitting the above recall probabilities with $w_2 = 0.2781$, $\delta_l = 0.3357$ and $\delta_f = 0.0584$. Using these parameters, an initial activation level of 0.4366 for all items and Person-Pro's influence, we reproduce the following recall probabilities: $0.4826, 0.3199, 0.3067$, and 0.4366 for RP+,RP-,NRP-pro, and NRP-anti respectively. When comparing these values with the ones obtained in [3], only RP+ items have a small error of 0.0299 suggesting that the initial recall probabilities are in fact close to uniform and that this assumption is not too strict. After exposure to Person-Pro the conversation phase starts. Two agents are paired into a heterogenous pair with an attitude bias or a homogenous pair without bias. Agents each have an attitude sampled from $[1, 9] \subset \mathbb{N}$ with $\delta := 5$ and are paired with another random agent. Items affected by the bias are RP- items, as X_1, and RP+ items, as X_2. This models the observed increase in recall of RP- items by participants with a pro bias and the reduced recall of RP+ by participants with an anti bias. The conversation length m_{conv} is eleven, the mean conversation length in [3].

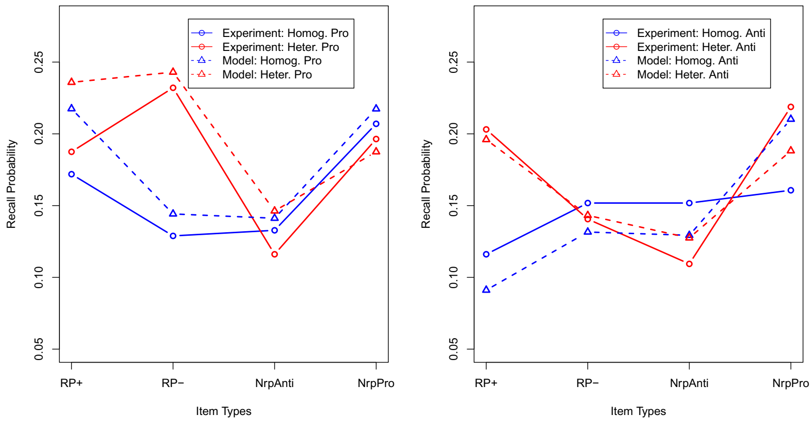

Fig. 2. Model and experimental data for pro (left) and anti (right) participants and agents in the conversation

To determine f_{bias} we tested values between 0.0 and 0.5 in intervals of 0.1 and after obtaining 0.15 as the best fit in terms of the mean squared error (MSE) for values compared in 2 we tested intervals of 0.01 from 0.1 to 0.3, leading to $f_{bias} = 0.23$ with a MSE of 0.0083. The recall probabilities, measured by the frequencies with which pro and anti agents recall items of a certain type, are seen in Fig. 2 and compared to the experimental data on which the MSE was computed. The main effects from [3] are modeled with the exception of reduced RP+ recall for pro participants (left on Fig. 2).

4 Results

To measure the level of convergence of memory in a network we define mnemonic distance between any two agents by $\|f_i - f_j\| = \sum_{x \in X} |f_i(x) - f_j(x)| \in \mathbb{R}^+$. Across the network we consider the mean mnemonic distance of all agents to all other agents. In the following simulations we keep the parameters as determined for the experiments with the exception of the G^c which is now a *small-world* network with rewiring probability (the proportion of links that cut across the network) $p = 0.20$, and $k = 3$ neighbors for every agent. This network configuration corresponds to the configuration of a plausible social network [15]. Every agent had C conversations in the network, choosing a random neighbor in G^c for each, and thus making homogeneous and heterogeneous conversations equally likely. Each model was run for ten trials.

Is mnemonic consensus influenced by network size and by the number of conversations among agents? In order to answer this question, we varied $C = 0, 1, 2, \ldots, 9, 10, 20, \ldots, 60$ and $n = 10, 50, 100$. Results in Fig. 3 indicate that smaller communities arrive at a mnemonic consensus with fewer conversations than larger communities and also converge to a different asymptote. Convergence occurs after an initial increase in mnemonic distance due to the conversational dynamics following exposure to Person-Pro. A possible explanation for this effect is that most memory items are unlikely to be mentioned in consecutive conversations in smaller communities. In larger communities the likelihood for consecutive mentions, at least for some local neighborhoods, is larger and once mentioned sufficiently often these memory items can persist locally and hence lead to a larger mnemonic distance in the network.

Does Person-Pro have a differential impact on mnemonic consensus as a function of network density? To explore this dynamic, we manipulated network density by varying $k = 1, 2, 3, 4, 5$, and we ran simulations for networks with and without the influence of Person-Pro where $n = 100$, $p = 0.10$, and $C = 10$. As seen in Fig. 3, the number of neighbors k influences consensus with denser networks reaching consensus faster. In addition, Person-Pro's discourse decreases the average mnemonic consensus, most likely because it differentially affects homogeneous and heterogeneous pairs. We expect that introducing homophily in the network would result in a reversed pattern [15]. Finally, these results also suggest that mnemonic convergence is possible even for large networks in which agents have relatively few conversations ($C = 10$), but only if the agents are richly connected.

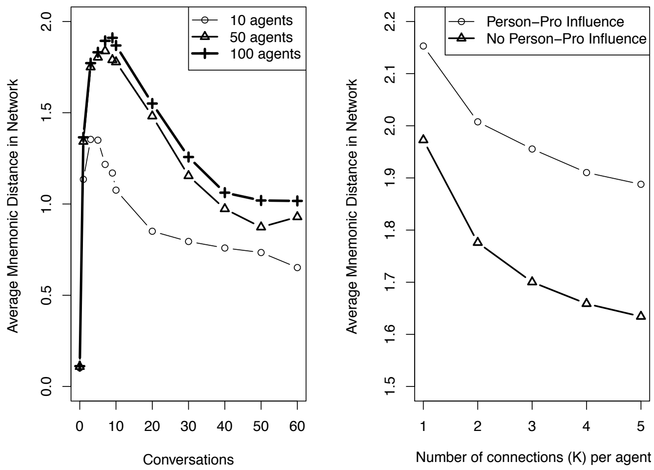

Fig. 3. Consensus dynamics (lower score=larger consensus) as a function (left) of community size and number of conversations per agent; and (right) of network structure (k neighbors) and presence/absence of Person-Pro's influence. Each data point represents the average of 10 runs.

5 General Discussion

In this paper we constructed an agent-based model with psychologically realistic assumptions about memory and conversational dynamics. It builds upon solid empirical work on the propagation of memories in conversations, focusing on two ways memories can be altered: through practice effects and through retrieval-induced forgetting. By extending this empirical work into the domain of agent-based modeling, the current work has underscored how: (1) psychological principles can guide the construction of simulations, and (2) how agent-based modeling can allow for the study of memory propagation and convergence across large network structures. Even at this early stage of model development, the results bear on critical issues concerning the formation of collective memory that should be of interest to both social scientists and policy makers. Specifically, the results supply some insight into why researchers might insist that bands or demes are well-suited for the construction of social reality [15]. It also underscores the fact that certain network structures can facilitate the construction of shared reality even in larger communities. In particular, the results indicate that mnemonic convergence depends on group size and number of conversational exchanges, but also hint at the fact that connectedness of agents in a network, as well as the nature of these connections are critical factors to be considered. With technological advances in human subjects experimentation (e.g. Mechanical Turk), we're planning on validating our results by using what Mason and Watts [11] call virtual experiments; that is, experiments in which groups of up to 20 participants interact in experimenter controlled tasks. Finally, although more detailed models charting the interaction between network structure, network size, memory, and conversational exchange still need to be formulated, the research presented here provides a solid framework for undertaking this endeavor.

References

[1] Anderson, M.C., Bjork, R.A., Bjork, E.L.: Retrieval induced forgetting: Evidence for recall specific mechanism. Psychonomic Bulletin & Review 7, 522–530 (2000)
[2] Caporael, L.R.: The evolution of truly social cognition: The core configuration model. Personality and Social Psychology Review 1(4), 276–298 (1997)
[3] Coman, A., Hirst, W.: Cognition through social networks: the propagation of practice and induced forgetting effects. Journal of Experimental Psychology: General (in press, 2012)
[4] Cuc, A., Koppel, J., Hirst, W.: Silence is not golden: A case for socially-shared retrieval-induced forgetting. Psychological Science 18, 727–733 (2007)
[5] Epstein, J.: Generative social science: Studies in agent-based computational modeling. Princeton University Press, Princeton (2006)
[6] Hirst, W., Echterhoff, G.: Remembering in conversations: the social sharing and reshaping of memories. Annual Review of Psychology (in Press)
[7] Hirst, W., Manier, D.: Towards a psychology of collective memory. Memory 16, 183–200 (2008)
[8] Hutchins, E.: Cognition in the Wild. MIT Press, Cambridge (1995)
[9] Kaplan, R.D.: Balkan Ghosts: A Journey Through History. Picador (1993)

[10] Karpicke, J.D., Roediger, H.L.: Repeated retrieval during learning is the key to long-term retention. Journal of Memory and Language 57, 151–162 (2007)
[11] Mason, W., Watts, D.: Collective problem solving in networks. Proceedings of the National Academy of Sciences (in press)
[12] Power, C., Dunbar, R., Knight, C. (eds.): The Evolution of Culture. Edinburgh University Press (1999)
[13] Sun, R.: Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation. Cambridge University Press, New York (2006)
[14] Zerubavel, Y.: Recovered Roots: Collective memory and the Making of Israeli National Tradition. University of Chicago Press (1997)
[15] Watts, D.: Six degrees: The science of a connected age. Norton, New York (2004)