Understanding Human Social Communication: A Computational Model of Gossip

Jeungmin Lee (jeungminlee@kaist.ac.kr)
Department of Bio and Brain Engineering, Program of Brain and Cognitive Engineering
College of Engineering, KAIST, Daejeon, 34141, South Korea

Jerald D. Kralik (jerald.kralik@gmail.com)
Department of Bio and Brain Engineering
College of Engineering, KAIST, Daejeon, 34141, South Korea

Jaeseung Jeong (jsjeong@kaist.ac.kr)
Department of Bio and Brain Engineering, Program of Brain and Cognitive Engineering
College of Engineering, KAIST, Daejeon, 34141, South Korea

Abstract

Updating people about the actions of others—social communication—is a powerful means by which humans learn about the world and maintain stable societies. However, how the mind/brain achieves this ability computationally remains unclear. Our goal is to model when, how, and why people choose to communicate information about others to others. Here we present current progress. We first describe our social communication framework, the test paradigm for model development and assessment, and an empirical experiment we conducted to obtain novel data to test model predictions. We then present our model, and compare it with two others. Our model outperformed the others, capturing the main patterns of the empirical data and matching the specific results most closely (i.e., percent of cases deciding to communicate about a target individual). Thus, our model successfully simulates human social decision-making, helping to understand how it is achieved by the human mind/brain.

Keywords: evolution of social cognition; theory of mind; communication; decision-making; computational model

Introduction

Observing the actions of others is a principal means by which humans learn and update knowledge—both about the world as well as the person performing the act—greatly extending our reach beyond our own individual experiences. Moreover, learning from and about others ratchets up even further with communication, not only from the performer to the observer, but in turn from the observer to someone else, and so on. In this way, information quickly disseminates across the social network (in turn enabling social networks to scale). Additionally, the ability to influence a person's future actions increases dramatically (e.g., via social influence or appeals to authority when someone's actions are in question). Communication about the action of others, then, provides extraordinary value to social groups, and likely played a leading role in the evolution of the human brain (Dunbar, Marriott & Duncan, 1997; Dunbar, 2004). To date, however, how we achieve this capacity at a computational level remains unclear.

This is because the ability is deceptively elaborate and complex. For example, focusing on the observer as the *central agent*, there are multiple critical factors that determine whether to transmit information about the actions of a *target* individual to someone else, i.e., a *receiver*. In general, it requires assessing the significance of the target's action, and whether a receiver would be interested in learning of it (and/or whether the information could likely feedback to the target person and influence future behavior). On first order, the significance of the action can be measured in terms of potential *benefits* and *costs* to self and others. Making this evaluation requires the central agent to have an internal *value scale* that assigns the degree of significance to particular target actions. In other words, the central agent must possess a minimal *affective* apparatus (Gazzaniga, Ivry & Mangun, 2013). In assigning value, a target's actions can again be categorized in two general categories: whether significant as world knowledge independent of the target, or valuable knowledge about the target him/herself. For the latter, given that social interactions comprise such a significant portion of our daily lives (whether at home, workplace, or almost anywhere else), information about others (i.e., their locally stable traits and behaviors) is critical. Indeed, Dunbar and colleagues (1997) found that over 60% of conversations involve discussing others. Because of this significance, and the general lack of development to date, our model currently focuses on this *social knowledge*.

Intriguingly, this type of social communication—telling others about someone else—has normally been defined as *gossip*. Although gossip may seem superficial, it is in fact an important mechanism underpinning society (Dunbar, Marriott & Duncan, 1997; Dunbar, 2004; Foster, 2004). For shorthand, we thus also use the term "gossip" to refer to this form of social communication: in which one informs another individual about events involving someone else not present.

In fact, disseminating knowledge about an individual's actions across a social network (boot-strapping culture and societies) belies a formidable affective-sociocognitive engine

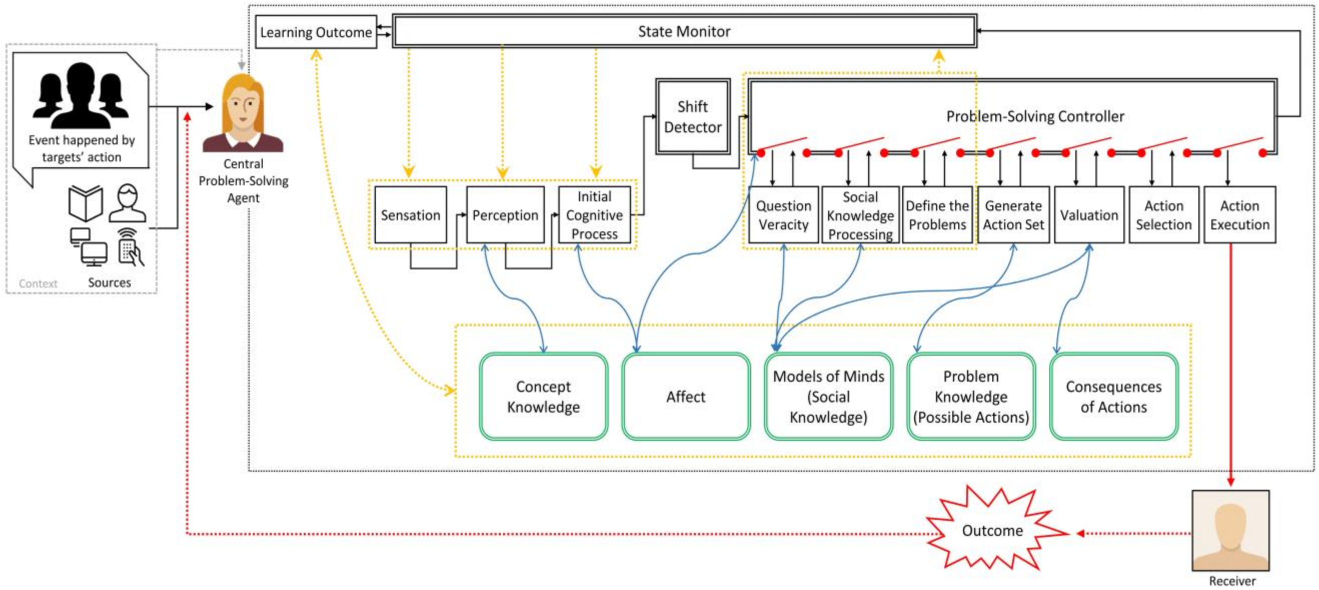

Figure 1. The complete framework of internal process regarding social information and gossip decision. The central agent goes through a series of internal processes (black boxes) managed by higher metacognitive processes (double-lined boxes) by accessing the knowledge (green rounded boxes). See text for details (Lee, Kralik, & Jeong, 2018; 2019).

under the hood: one that includes not only the affective valuation assessment, but also mind-reading (i.e., interpreting intentions underlying someone's actions and whether others know or would care about it), and social accounting (i.e., a social currency, based on the value the information provides to others) (Cosmides & Tooby, 1992; Gazzaniga, Ivry & Mangun, 2013; Lee, Kralik, & Jeong, 2018; 2019). Our goal, then, is to model when, how, and why people choose to communicate social information to others.

The current paper presents our progress, by first describing our overall social communication framework, as well as the test paradigm for model development, and the empirical experiment we conducted in our cognitive neuroscience laboratory to obtain novel data to test model predictions (Anonymous, *submitted*). We then present the computational development of our main model, together with two alternative models. We then compare the models on how well they fit the empirical findings.

Framework, Methods, and Models

In this section we describe the framework and test paradigm used for the empirical experiment and model development. We then describe our main computational model, and then two others based on simpler versions of the main model.

Framework for Social Communication

We developed the general framework to capture the fundamental components of social intelligence and communication (Figure 1) (Aronson, Wilson, Akert & Sommers, 2016; Cosmides & Tooby, 1992; Dunbar, Marriott & Duncan, 1997; Dunbar, 2004; Foster, 2004; Gazzaniga, Ivry & Mangun, 2013; Haidt, 2007; Kralik, 2017; Kralik et al., 2018; Lee, Kralik, & Jeong, 2018; 2019). It is based on a *target* individual being involved in some event, such as hitting a coworker, caught cheating on an exam, helping people escape a burning building, or going to the movies. The *Central Agent* then learns of the event, and subsequently processes it via a series of subprocesses (black boxes in Figure 1) that also utilize specific knowledge stores (green boxes).

The overall process begins with the initial sensory input passed from *sensation* to *perception* to *initial cognitive processing*, in which for example, auditory input is transformed to sentences, and initial meaning is ascertained, including the event content evoking an initial affective (i.e., emotional) response, such as cheating on an exam being bad. Although we have been developing these early modules, the current model simply begins with a three-element vector representing the event involving a target individual, which then is mapped to an *affect score* (described further below). The affect score reflects the level of initial interest or concern the target's action evokes, with the overall affect assignment process being the affective/emotion center of our framework and model (Gazzaniga, Ivry & Mangun, 2013). Next, if the affect value is sufficiently high, the *shift detector* activates the *problem-solving controller*, which in turn activates individual subprocesses by closing the corresponding affect gates.

The subprocesses include determining the reliability of the information source, updating the model of the target's mind, determining the specific problem(s) at hand, generating the relevant action set (such as whether or not to communicate to another person as *receiver*), valuation of the possible actions (i.e., assessing the benefits versus costs), action selection based on these action valuations, action execution, and then monitoring of the action outcome for potential learning (though learning is not yet explicitly modeled) (see Figure 1). Computational development has thus far focused on *valuation*, the central process that determines whether to take action based on the event, described next.

Test Paradigm and Behavioral Experiment

Input to the system is an event or *scenario* involving someone, such as someone caught cheating on an exam or verbally abusing a coworker. Each event is then represented by three fundamental factors: the *target* individual, the event *content*, and the *valence* of the content, that is, whether positive or negative. Each independent variable (i.e., our three event factors) is then further divided into a number of fundamental categories. As stated, content *valence* is subdivided into positive and negative events. For *target*, we examined *ingroup* versus *outgroup* versus *celebrity*, in order to test the important social factors of contact, caring, and status (described further below). Finally, for *content*, our intention was to produce a comprehensive set of social events that occur in daily life (either rarely or frequently). Based on theoretical considerations and literature review, this resulted in eight content domains that can be roughly aligned according to how much we care about them, that is, how much affect or emotion they evoke, represented by an affect score, listed in Table 1 (Aronson et al., 2016; Cosmides & Tooby, 1992; Dunbar, Marriott & Duncan, 1997; Dunbar 2004; Foster, 2004; Gazzaniga, Ivry & Mangun, 2013; Haidt, 2007; Kralik, 2017; Lee, Kralik & Jeong, 2018; 2019). Based on all combinations of these factors, we developed 48 different scenarios (3 target · 2 valence · 8 content). Using this comprehensive set of scenarios, we developed our main computational model and generated a set of predictions of how the three independent variables (target, content, and valence) influence gossip spreading.

Table 1. Eight content domains and their affect score (where *m* signifies morality domain: Haidt, 2007).

Content	***Affect Score***
Prosociality (care/harm)^{m}	7
Fairness (fair/cheating)^{m}	6
Competition (positive/negative)	5
Social-oriented (altruism/selfishness)	5
Community (loyalty/betrayal)^{m}	4
Respect (authority/subversion)^{m}	4
Purity (sanctity/degradation)^{m}	3
General social affairs (positive/negative)	1

To collect the empirical data, we recruited 102 participants (59 females and 43 males, mean age 23.8 years, range 20-32), and each participant was shown a gossip scenario and asked if they wanted to spread it to other people (i.e., receivers) (Lee, Kralik & Jeong, *submitted*). The 48 scenario types were replicated three times. As a result, 144 different gossip scenarios were used per participant in the experiment. Gossip rates of the scenarios were calculated by taking the mean across replications and participants.

Test Paradigm and Behavioral Experiment

We now describe the computational models, developed and tested using Matlab (The Mathworks, Natick, MA, USA). For all three models, on every trial, that is, when one of the 48 scenario events occurs, the given event's *content* is converted into (a) an *affect score* (Table 1) and (b) a valence flag (i.e., 0 for positive, 1 for negative), which occurs in the *Initial Cognitive Process* module in Figure 1. The *target* for the given event is then identified as *ingroup*, *outgroup*, or *celebrity* in the *Social Information Processing* module and then converted to *contact*, *caring*, and *status* values shown in Table 2.

Table 2. Parameter values of the best fits of the three components of target (contact, caring, and status) for ingroup, outgroup, and celebrity for the three models.

	Model 1		
	$t_{contact}$	t_{care}	t_{stat}
Ingroup	1	0.9	0.2
Outgroup	0.1	0.1	0.1
Celebrity	0.1	0.3	1
	Model 2		
	$t_{contact}$	t_{care}	t_{stat}
Ingroup	X	0.8	X
Outgroup	X	0.05	X
Celebrity	X	0.5	X
	Model 3		
	$t_{contact}$	t_{care}	t_{stat}
Ingroup	X	0.8	X
Outgroup	X	0.05	X
Celebrity	X	0.5	X

This content, valence, and target information are all then sent to the valuation module, the critical subprocess in which the central agent decides whether or not to take a particular action (in our case, gossip) among a set of possible choices. To maximize expected outcome, the central agent needs to carefully consider all the pros and cons of the possible actions. We next describe this valuation process (and then action selection) further.

Our Main Model: Model 1 To decide which action to take, the central agent needs to estimate the outcome of every possible action based on *benefits* and *costs*. Because our current focus is the conditions under which one would or would not communicate to someone (a receiver) about someone else (the target), our model considers two actions: (1) gossiping or (2) not gossiping. In this case, *the costs of* *gossiping become the benefits of not gossiping*. The equations, then, are the following:

$$Value_{Gossip} = A \cdot B_{Total} \quad (1)$$
$$Value_{Not\ Gossip} = A \cdot C_{Total} \quad (2)$$

where A is the affect score of the given information (listed in Table 1), such that the given action value increases with the affective response.

Based on literature review and our own theoretical development, Table 2 shows the list of potential benefits and costs involved in gossip (Aronson et al., 2016; Cosmides & Tooby, 1992; Dunbar, Marriott & Duncan, 1996; Dunbar 2004; Foster, 2004; Gazzaniga, Ivry & Mangun, 2013; Haidt, 2007; Kralik, 2017; Kralik et al., 2018; Lee, Kralik, & Jeong, 2018; 2019). The most obvious benefit of gossip is that one can avoid facing the target directly (B1), especially if an unfavorable outcome is expected (e.g., the target becoming upset). Additionally, the central agent may obtain further information from the receiver about the incident or target (B2). Third, and especially critical, the information provided to receivers can update their 'broken' models about the target and the world (B3). Fourth, gossip can also promote fairness balance and societal stability by rewarding positive actions and punishing negative ones. This can potentially be accomplished by influencing the social status of the target, via affecting their reputation (B4). Additionally, fifth, receivers may be in better position to directly contact the target to reward or punish the behavior (B5). Finally, the target and receivers may enjoy entertaining target activities and learn from them (B6).

Although this indirect form of communication has many benefits, there are obvious costs. Although gossip allows avoiding direct contact with the target (B1), the target might yet ascertain the source of the gossip (i.e., the central agent) and retaliate (C1). Moreover, there is a risk that event information (and thus, about the target) is incorrect (C2), resulting in deleterious effects such as altering the target's social status. Because maintaining accurate models of others' minds is critical within a multi-agent society, sharing false information sows confusion. Third, the central agent may also earn a bad reputation as a gossiper (C3). Fourth, related to B2, additional information from receivers may be wrong or misleading (C4). Fifth, because of its indirectness, influencing the target via gossip may not match what the central agent intended (C5). Finally, choosing to gossip, like any action, requires both cognitive and behavioral effort, both to take the action and monitor its effects (C6) (Lee, Kralik & Jeong, 2019).

Total benefits and costs (B_{Total} and C_{Total} in the equations) are the summation of all potential benefits and costs. That is:

$$B_{Total} = B_1 + B_2 + \cdots \quad (3)$$
$$C_{Total} = C_1 + C_2 + \cdots \quad (4)$$

where each benefit Bi and cost C_i are then calculated as a function of the target factors T (i.e., contact, care, and status), valence V, and benefit-cost weighting factor w thus:

$$B_i = T_{contact,Bi} \cdot (T_{care,Bi} + T_{status,Bi}) \cdot V_{Bi} \cdot w_{Bi} \quad (5)$$
$$C_i = T_{contact,Ci} \cdot (T_{care,Ci} + T_{status,Ci}) \cdot V_{Ci} \cdot w_{Ci} \quad (6)$$

The benefit-cost weighting factors w are listed in Table 3, and are again based on theoretical considerations, literature review, and adjusting for model fitting (with yet maintaining the general relative positions among them).

To calculate target and valence factors (i.e., $T_{contact}$, T_{care}, T_{status}, and V), we combine the target and valence values input to valuation — i.e., those in Table 2 and the valence 1 (positive) or 0 (negative) flag — with gating values, g, as 0 (irrelevant) or 1 (relevant), based on whether each factor is relevant to the given benefit or cost. The gate values are listed in Table 3.

Table 3. Table 3. Benefits and costs of gossiping. Cells contain weighting (w) and gating values. The gates are social or valence filters and used in the valuation equation.

Valuation Categories	***w***	***Target (T)***			***Valence (V)***
		$g_{contact}$	g_{core}	g_{status}	v
B_1: Avoid direct contact with the target	1	1	0	0	1
B_2: Feedback to the gossiper from receiver	0.4	0	1	0	0
B_3: Update receiver's knowledge	1	0	1	0	0
B_4: Influence target's social status	0.9	0	0	1	1
B_5: Receiver influences target's behavior	0.8	0	1	0	1
B_6: Entertainment and social learning	0.4	0	1	0	1
C_1: Potential direct contact from the target	0.6	1	1	0	1
C_2: Risk of spreading wrong information	0.4	0	0	1	1
C_3: Earn bad reputation as a gossiper	0.7	0	0	0	1
C_4: Get wrong/misleading information from receivers	0.9	0	0	0	0
C_5: Influence target improperly	0.8	0	1	0	1
C_6: Cost effort both cognitively and behaviorally	0.6	0	1	0	1

Contact is based on whether the target can actually reach the central agent; *care* indicates how much the central agent cares about the target, such that the parameter is high if the central agent is invested in the outcome; and s*tatus* represents the target's position within the social hierarchy, such that celebrities are high, ingroup low, and outgroup the lowest. The input values and gates are then combined thus:

$$T_{contact} = 1+g_{contact} \cdot (t_{contact}-1) \quad (7)$$
$$T_{care} = g_{care} \cdot t_{care} + (1- g_{care}) \cdot \ ave\text{-}t_{care} \quad (8)$$
$$T_{status} = g_{status} \cdot t_{status} \quad (9)$$
$$V = [(1- v_{default}) \cdot g_{valence} \cdot v] + v_{default} \quad (10)$$

In general, for contact, if the central agent and target cannot directly contact each other (i.e., contact gate is 0), then the particular benefit or cost has no effect for the given scenario; for care, if caring is relevant for the given benefit or cost (gate=1), then it becomes the value for the target group in Table 2, otherwise (gate=0) it is the average of the three target group values. Finally, for valence, some benefits and costs would be expected to have a greater impact for negative events (such as directly contacting the target when they've done something egregious), and we represented this by having a default value for positive valence which increases for negative valence if the gate is 1.

Model 2 Although we believe the factors and values for our main model are well justified, it nonetheless is important to test their significance in the model. We thus developed two competing models that simplified prominent factors. For Model 2, gating values for the *contact* and *status* components of the *target* were all set to 0 while the gates for *care* were all set to 1. That is, here we collapsed the target components of contact, care, and status into one general factor that represented the difference among the three target categories of ingroup, outgroup, and celebrity. Table 2 shows the best fit values for this vector.

For *valence*, Model 2 does not consider cases where negative events may be more impactful than positive ones, and thus sets all valence gates to 0, using only the default value for all events. The best fit $v_{default}$ was 0.7.

As $g_{contact}$, g_{status}, and $g_{valence}$ are all set to 0, the valuation equations for Model 2 become:

$$B_i = T_{care,Bi} \cdot v_{default} \cdot w_{Bi} \quad (11)$$
$$C_i = T_{care,Ci} \cdot v_{default} \cdot w_{Ci} \quad (12)$$

This equation in turn means that the individual benefits and costs are considered as one total value for each. To best fit these benefit and cost weights, then, one parameter was used for benefits (α_{gossip}) and one for costs ($\alpha_{not\text{-}gossip}$). For Model 2, the best fit values were α_{gossip}=1.05; and $\alpha_{not\text{-}gossip}$=1.

Model 3 The second competing model, Model 3, was the same as Model 2 except for *valence*. For this model, for all benefits and costs, all negative events were considered more impactful than positive ones; and thus, all valence gates were set to 1. The Model 3 valuation equations then are expressed as:

$$B_i = T_{care,Bi} \cdot [(1- v_{default}) \cdot v + v_{default}] \cdot w_{Bi} \quad (13)$$
$$C_i = T_{care,Ci} \cdot [(1- v_{default}) \cdot v + v_{default}] \cdot w_{Ci} \quad (14)$$

Table 3 shows the best fit weights for the singular target effect. The best fit values for the relative benefit to cost effects were again α_{gossip}=1.05; and $\alpha_{not\text{-}gossip}$=1.

Results

We first examine the main effects for content *valence* and *target*. We then examine the *target* by *valence* interactions.

Figure 2 shows how the three competing models compare to the empirical findings for percent gossip for the *valence* and *target* main effects. For *valence*, percent gossip in the empirical data was significantly higher for negative than for positive content. All three models obtained the same pattern. Additionally, Model 1 had a gossip percentage closer to the empirical results than Models 2 and 3.

For *target*, percent gossip in the empirical data was significantly higher for celebrities, followed by ingroup, then outgroup. The same pattern was obtained by Model 1, whereas Models 2 and 3 showed a different pattern, with ingroup the highest. Thus, our main model provided better fits, showing that the target effects indeed appear to be a function of contact, care, and status, and both target and valence have distinct relevance to different benefits and costs.

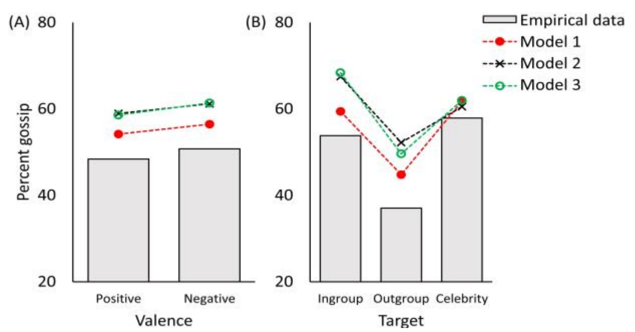

Figure 2. Results of the three models compared to the empirical data for gossiping with respect to (A) valence and (B) target main effects.

We next examined the *target* by *valence* interactions. Figure 3 shows the results of the three competing models compared to the empirical data for percent gossip for ingroup, outgroup, and celebrity targets, broken down by positive and negative valence. For ingroup targets, as opposed to the general main effect finding of more gossiping about negative events, we have the opposite: percent gossip for positively valenced events was significantly higher than for negatively valenced ones in the empirical data. Indeed, Model 1 obtained the same pattern, whereas Model 2 and 3 showed the opposite pattern, similar to the main effect result.

For outgroup and celebrity targets, the empirical data showed higher percent gossiping for the negatively valenced events than for the positive ones (as for the main effect). In these two cases, all three models obtained the same pattern. Thus, Models 2 and 3 were unable to obtain different patterns across the three target groups, while Model 1 was able to explain the flipped relationship for ingroup targets.

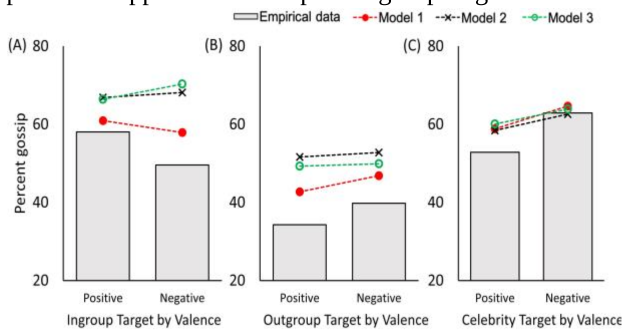

Figure 3. Results of the three models compared to the empirical data for gossiping about (A) ingroup, (B) outgroup, and (C) celebrity targets for different valence.

The main effect in the results, captured by Model 1, is that both the relative benefits of gossiping positively and costs of gossiping negatively are heightened for ingroup members. At the same time the relative lower cost and higher interest in justifying status led to increased negative gossiping about celebrities (Foster, 2004). For outgroup, a relative lack of interest dominated the findings, although the relative lower cost in negative gossiping was also observed.

We next examine the results for the *target* by *valence* interaction by examining positive and negative events for the three different target groups. For positively valenced events, Figure 4A shows that the percent gossiping in the empirical data was highest for ingroup targets, followed by celebrity, and finally outgroup. All three models obtained the same pattern, however, Model 1 obtained values closer to the empirical data. In contrast, for negative events, Figure 4B shows that percent gossiping in the empirical data was highest for celebrity targets, followed by ingroup, then outgroup. Model 1 again obtained the same pattern, whereas percent gossip about ingroup targets was higher than for celebrities for both Models 2 and 3. Thus, once again Models 2 and 3 showed the same patterns for positive and negative events, while Model 1 found different patterns for the two valences, matching the empirical data.

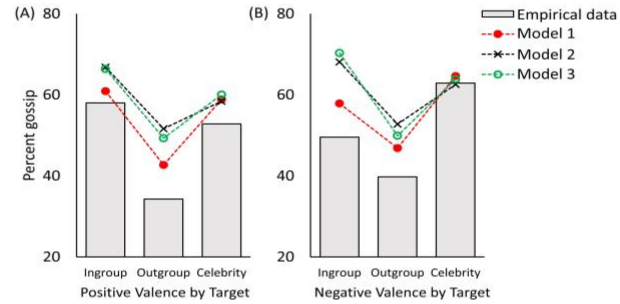

Figure 4. Results of the three models for (A) positive and (B) negative valence with respect to the different targets.

In sum, our main model, Model 1, was superior in not only the level of gossiping predicted, but in importantly capturing the general patterns in the findings, especially flipping the prediction for ingroup, with an increase in positive gossiping and decrease in negative for ingroup members.

Discussion

To understand how the human mind/brain has harnessed the ratcheting power of social information exchange, it is important to model how people process social events, and when, how, and why they choose to communicate the information to others. To this end, we have developed a general framework for human social intelligence and communication based on literature across the social sciences, and have begun developing a computational model detailing the processes. Here we presented a significantly elaborated computational version of our model, based on this literature and our own evolutionary and affective-sociopsychological theoretical considerations of why people should choose to communicate this information (Aronson et al., 2016; Cosmides & Tooby, 1992; Dunbar, Marriott & Duncan, 1997; Dunbar, 2004; Foster, 2004; Gazzaniga, Ivry & Mangun, 2013; Haidt, 2007; Kralik, 2017; Kralik et al., 2018; Lee, Kralik, & Jeong, 2018; 2019). We have also conducted an experiment in the laboratory in which we collected data on whether people would choose to communicate to others in order to test the predictions of our model (Lee, Kralik, & Jeong, *submitted*). Indeed, our model predictions were supported, successfully capturing the main patterns of results.

There are, of course, multiple avenues for future development. These include a more detailed consideration by the central agent of why the target acted as he/she did. The answers to "why" will require a richer set of social factors (beyond contact, care, and status), which will in turn be compared against the central agent's own model of the target's mind. For possible responses, we also plan to include communicating directly to the target (rather than to others about them). Eventually, more action detail also needs to be included, to carry out, for example, an extended conversation with the target and/or receiver. For event content, more activities are needed. We also intend to add learning to our model, including the need for active monitoring of outcomes to assess actual action effectiveness, especially with indirect communication. Learning can also potentially capture cultural influences on moral dimension weightings.

Evidence shows that social intelligence and communication are comprised of relatively hard-wired components (and thus to some extent expert-like), together with more malleable ones, with their combination enabling *general* social intelligence across multiple content domains. Moreover, a comprehensive integration of the relevant literature across multiple fields of study shows that human social ability is indeed elaborate and complex. We thus believe that approaches such as ours that face this challenge head-on are necessary to ultimately understand human social processing, and endow artificial systems with the affective-sociocognitive processing machinery that truly leverages the power of sociality.

References

Aronson, E., Wilson, T. D., Akert, R. M., and Sommers, S. R. (2016). *Social Psychology*. Pearson: NYC.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In: J. H. Barkow, L. Cosmides, J. Tooby (Eds.). *The Adapted Mind*.

Dunbar, R. I. M. (2004). Gossip in evolutionary perspective. *Review of General Psychology*, 8(2), 100–110.

Dunbar, R. I. M., Marriott, A., and Duncan, N. D. C. (1997). Human conversational behavior. *Human Nature*, 8(3), 231–246.

Foster, E. K. (2004). Research on gossip: taxonomy, methods, and future directions. *Review of General Psych*. 8, 78–99.

Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2013). *Cognitive neuroscience: the biology of the mind.* WW Norton & Company.

Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316(5827), 998–1002.

Kralik, J. D. (2017). Architectural design of mind & brain from an evolutionary perspective. *Proc. AAAI 2017 Fall Symposium: A Standard Model of the Mind*.

Kralik, J. D., Lee, J. H., Rosenbloom, P. S., Jackson, Jr., P. C., Epstein, S. L., Romero, O. J., Sanz, R., Larue, O., Schmidtke, H. R., Lee, S. W., McGreggor, K. (2018). Metacognition for a Common Model of Cognition. *Procedia Computer Science*: 145, 730–739.

Lee, J., Kralik, J. D., & Jeong, J. (2018). A Sociocognitive-Neuroeconomic Model of Social Information Communication: To Speak Directly or To Gossip. *Proceedings of the Annual Meeting of the Cognitive Science Society* (*CogSci).*

Lee, J., Kralik, J. D., & Jeong, J. (2018). A General Architecture for Social Intelligence in the Human Mind and Brain. *Procedia Computer Science: 145, 747–756*.

Lee, J., Kralik, J. D., & Jeong, J., Submitted.