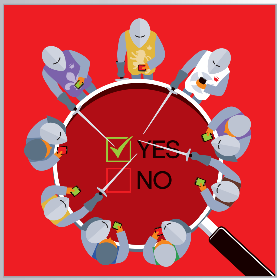

Reasoning about Sentiment and Knowledge Diffusion in Social Networks

Social media platforms, taken in conjunction, can be seen as complex networks; in this context, understanding how agents react to sentiments expressed by their connections is of great interest. Here, the authors show how Network Knowledge Bases help represent the integration of multiple social networks, and explore how information flow can be handled via belief revision operators for local (agent-specific) knowledge bases. They report on preliminary experiments on Twitter data showing that different agent types react differently to the same information — this is a first step toward developing tools to predict how agents behave as information flows in their social environment.

Fabio R. Gallo, Gerardo I. Simari, Maria Vanina Martinez, and Marcelo A. Falappa
Departamento de Ciencias e Ingeniería de la Computación, Universidad Nacional del Sur (UNS) & ICIC (UNS–CONICET), Bahía Blanca, Argentina

Natalia Abad Santos
Departamento de Matemática, Universidad Nacional del Sur (UNS), Bahía Blanca, Argentina

Social media has without a doubt changed the way people and institutions communicate with each other; in particular, the flow of information has been immensely affected by this revolution. Whereas just 15 years ago marketing or political campaigns solely relied on TV, radio, billboards, or flyers, today we can easily see a transformed landscape, where political discussions and attention-hungry ads are staples of daily activity on Facebook, Twitter, YouTube, Instagram, Google+, LinkedIn, Pinterest, Tumblr, and others. A recent example of this is the 2016 elections held in the United States, where Facebook and Twitter played an important role both in candidates' efforts to communicate with their constituency[1,2] and in the spread of so-called *fake news* stories.[3,4] To understand the extent to which social media played a role in the elections, it's helpful to consider the fact that Donald Trump's team was testing thousands of variants of their ads – 40,000 to 50,000 on usual days, and reaching 175,000 on the day the candidates engaged in their third debate.[5]

One useful way of looking at this landscape is through the lens of *multiagent systems* (MAS),[6] which can be described as a software engineering metaphor for a set of autonomous entities called *agents* with a shared environment. This metaphor is quite useful both for developing the basic machinery underlying the problems that arise (logic-based reasoning, synchronization,

 1089-7801/17/$33.00

Related Work

Networks have been used to model different kinds of diffusion processes in real-world domains, such as epidemics spreading through a population, cascading electrical power failures, marketing, and the spread of mutant genes. Many different disciplines therefore have studied variants of this model, such as biology,[1] economics,[2] physics,[3] sociology,[4] and of course computer science.[5] As we mentioned in the introduction, more ad hoc models also quite recently have become central to world events, as both the US presidential election and the Brexit vote were influenced by the use of social media.[6,7]

What distinguishes these models from complex networks is their expressive power — they lack the capability to represent multiple attributes of nodes and edges, model competing diffusion processes, and others such as representing time. Our current line of research continues the work of Paulo Shakarian and colleagues,[8] where the authors propose a general formalism to model complex networks. Whereas their work focuses on cascading processes, they don't contemplate individual knowledge bases for each agent, and thus our work can be seen as a generalization. Our end goal is to build on these first steps — modeling how local revisions can be performed by agents — to eventually model the combination of these processes throughout the network as they give rise to cascades. The dynamics of how data and more general knowledge is communicated, processed, and adopted or rejected makes this problem quite difficult, as it generalizes the early models both in terms of expressive power and, consequently, the kinds of problems that must be solved. In prior work,[9] we also analyzed how traditional belief dynamics operators could be applied toward a solution — the conclusion was that any straightforward application of such operators has serious flaws, because basic desirable properties are violated, so new machinery is needed to tackle the general problem of performing social revisions.

Finally, another related line of work is presented by Luciano Tamargo and colleagues,[10] where belief revision is also studied in the multiagent systems setting considering the credibility or trust associated with each informant (agent) represented as a strict partial order among them. They define different kinds of change operators (expansion, contraction, and both prioritized and nonprioritized revision), and each operator is also able to modify the informant credibility according to new perceptions.

References

1. E. Lieberman, C. Hauert, and M.A. Nowak, "Evolutionary Dynamics on Graphs," *Nature*, vol. 433, no. 7023, 2005, pp. 312–316.
2. T.C. Schelling, *Micromotives and Macrobehavior*, W.W. Norton and Co., 1978.
3. V. Sood, T. Antal, and S. Redner, "Voter Models on Heterogeneous Networks," *Physical Review E*, vol. 77, no. 4, 2008, p. 041121.
4. M. Granovetter, "Threshold Models of Collective Behavior," *Am. J. Sociology*, vol. 83, no. 6, 1978, pp. 1420–1443.
5. D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the Spread of Influence through a Social Network," *Proc. Knowledge and Data Discovery*, 2003, pp. 137–146.
6. C. Shao et al., "Hoaxy: A Platform for Tracking Online Misinformation," *Proc. World Wide Web Companion,* 2016, pp. 745–750.
7. H. Allcott and M. Gentzkow, "Social Media and Fake News in the 2016 Election," *Nat'l Bureau of Economic Research*, tech. report 23089, 2017; www.nber.org/papers/w23089.
8. P. Shakarian, G.I. Simari, and D. Callahan, "Reasoning about Complex Networks: A Logic Programming Approach," *Theory and Practice of Logic Programming*, vol. 13, nos. 4–5, online supplement, 2013.
9. F.R. Gallo et al., "Belief Dynamics in Complex Social Networks," *Proc. Simposio Argenteno de Inteligencia Artificial*, vol. 45, 2016, pp. 70–77.
10. L.H. Tamargo et al., "On the Revision of Informant Credibility Orders," *Artificial Intelligence*, vol. 212, 2014, pp. 36–58.

communication, uncertainty, learning, and so forth) as well as for implementation purposes.

In this article, we're interested in adopting the MAS viewpoint to tackle the problem of reasoning about the diffusion of knowledge and sentiment in social media. The variety of options available make it necessary to model the underlying communication structure as a *complex network*, in which nodes are agents and we have the capability to model multiple relations between them, as well as multiple attributes that label both the nodes themselves and the relations in which they participate. For instance, nodes could be labeled with gender (male or female, or richer non-binary categorizations), date of birth, phone number, and political views (again, either a simple binary flag or a richer structure), whereas examples of their relationships are friends, coworkers, couples, family, pets, or between people and organizations (jobs, support, and services such as gym membership). Clearly, the potential to integrate many different data sources is enormous.

Social Knowledge Bases

In recent work,[7] we proposed a model called *Social Knowledge Bases* (SKBs, for short) that was conceived from a series of desirable properties for systems that work with the kind of data that are produced by agents in social media environments. The unique combination of challenges in this setting involves, among others

- multiple attributes with different domain data types,
- multiple relations between agents,
- uncertainty stemming either from data integration (inconsistency/overspecification or incompleteness/underspecification) or from inherently uncertain information,
- reasoning with agents' preferences,
- dealing with groups of agents as agents in their own right,
- computational tractability constraints,
- revising beliefs arising from interactions among the agents, and
- cascading processes (an action in the network can have a "domino effect" and reach many other agents).

The main focus of this article is on the last two points, which refer to issues arising from the way in which information flows in the network, as well as how agents dynamically react.

Data and Belief Dynamics

The problem of modeling how knowledge bases (KBs, for short) change in response to different kinds of events is commonly known as *belief dynamics*; in particular, deciding how to integrate an epistemic input into a KB is called *belief revision*. The latter has been studied from the point of view of KBs comprised of formulas closed under consequence (called *belief sets*)[8,9] as well as not closed (called *belief bases*)[10,11] — for a survey on different aspects of belief revision, see Pavlos Peppas's work.[12] When the epistemic input is a set instead of a single sentence (useful, for instance, to combine different sources of information), we use the term *belief merging*[13–16] or *multiple change*[16] instead.

As we mentioned in the introduction, one of the key aspects of interactions that arise among agents in a social media environment is the dynamic aspect of information — in particular, we're interested in agents' beliefs, which can be analyzed by observing the content they share, as well as the sentiments they express while doing so. In prior work,[17] we describe a formalization of a special kind of SKB that we call *Network KBs* (NKBs, for short) dealing specifically with this problem. Informally, NKBs are directed graphs $G = (V, E)$, where vertices represent agents and edges represent their relationships, augmented with the following:

- Labeling functions l_{vert} and l_{edge} for vertices and edges, respectively. Vertices are labeled with attributes and their values, and edges are labeled both with attributes and values, as well as a number in the [0,1] interval expressing the relationship's weight.
- A *local knowledge* base for each agent, representing its current set of beliefs. The language used to represent such knowledge can vary in expressivity; in this work, we'll assume they're expressed in propositional logic.
- A set of *constraints* characterizing situations that shouldn't be possible; they're comprised of formulas that can make reference both to the network structure and the agents' KBs. These can be simple data constraints such as vertices having more than one spouse (if this is a valid constraint in the domain being modeled), or they can be more complex, such as the fact that two very close friends shouldn't disagree on certain aspects.

Figure 1 shows a sketch of how an NKB models the integration of social network data. The details of this creation process are outside the scope of this article, because it's mainly a knowledge engineering task — for instance, determining the degree to which one user follows another, or populating the local KBs. Here, we're primarily interested in the challenges that arise when NKBs are already available to use as reasoning tools.

Content posted by agents to their social media sites are modeled by what we call *news items*, which are simply triples of the form $\langle$*agent*, *content*, *action*$\rangle$, where *content* represents pictures, videos, status updates, links, comments, and likes, and *action* is either *add* or *remove*. When an agent checks its feed, it's subject to a set of zero or more such news items, and must perform a *local revision* to see if any changes must be made to its individual KB.

Example 1. Consider a typical scenario in which presidential elections will be held soon, and there are two prominent candidates, A and B. Alice thinks that candidate A is the better option for her country than B — we can represent this fact with α; so we have α in Alice's KB.

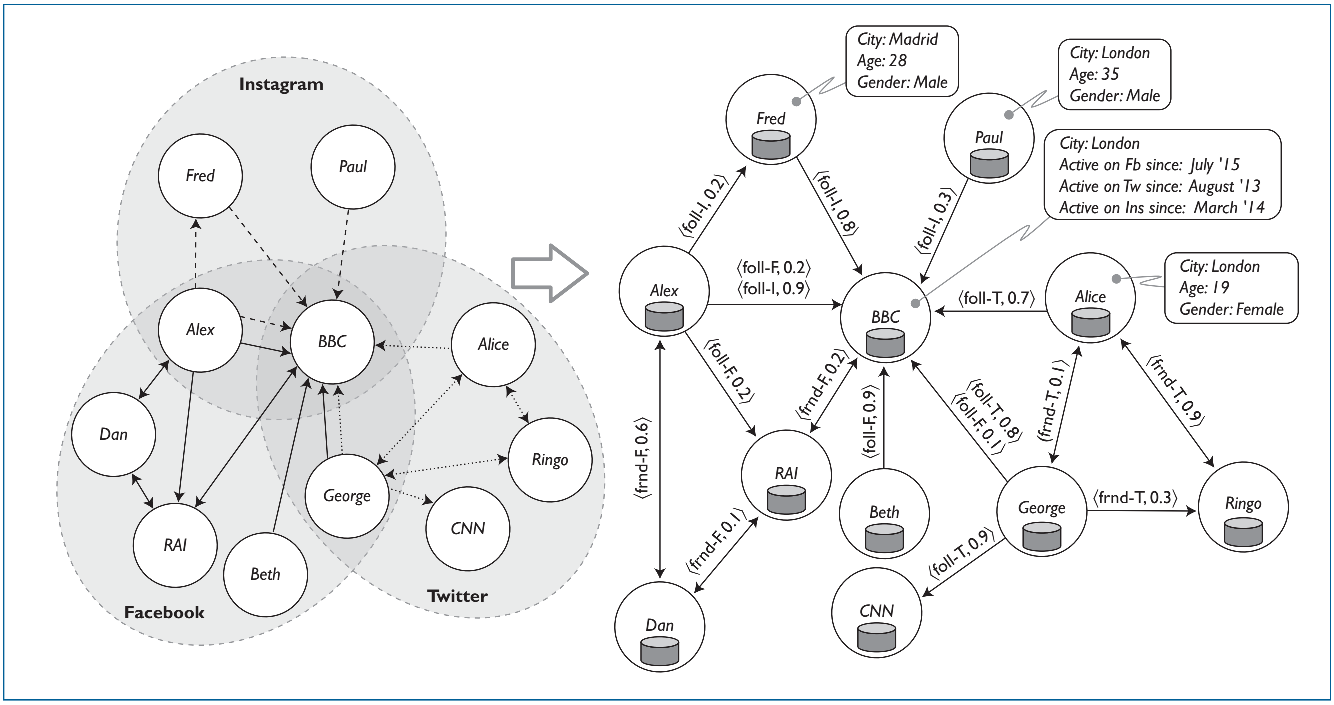

Figure 1. The Network Knowledge Base (NKB) model provides a way to represent the result of integrating information from multiple sources. Here, we have three social networks that share users; dashed, dotted, and full arrows denote relations in Instagram, Twitter, and Facebook, respectively. On the right, we show the resulting NKB, comprised of a complex network augmented with local KBs for each node. Edge and node labels store additional information, *such as the fact that George follows the BBC to a larger degree on Twitter than Facebook, or that Paul is a 35-year-old man who lives in London. We only show examples of node labels for a small subset to aid readability.*

George tweets that candidate A plans to reduce the budget for education if they win (and so he is against A, denoted with $\neg\alpha$); on the other hand, BBC posts both on Facebook and Twitter that candidate A is concerned about low academic achievement and they plan to increase the budget by 10 percent (a support for A). These news items can be represented with $\langle George, \neg\alpha, add\rangle$ and $\langle BBC, \alpha, add\rangle$.

The Revision Process

Any belief revision process is characterized by what are called *revision operators;* in our setting, we identify two such processes, one that's performed at the *local level*, where agents receive news items from their social media feeds (which causes them to revise their local KBs), and one that's performed at the *global level* (which can potentially have effects throughout the entire NKB, including the network structure). Figure 2 shows an overview of the overall process: the system evolves in a loop — in the current state, agents receive news items from their social media feeds, and revise their local KBs. Once this process has finished, a global operator is applied to the entire NKB to address any violations of constraints. Note that we're assuming synchronicity among independent local operators — this isn't an essential assumption, and we could apply the global operator either at regular intervals or whenever certain conditions are met. In this setting, there will inevitably be intermediate steps in which the NKB is inconsistent (that is, some of the constraints in the NKB might be violated).

In the remainder of this article, we'll focus solely on the local revision process. Therefore, here, *social revision operators* take as input an NKB, denoted Δ; a local KB belonging to a node v in Δ, denoted $K(v)$; and a set of news items that we call the *epistemic input*, denoted P. The output of such operators is a new (modified) NKB: we'll denote the output with Δ' and therefore the modified local KB corresponding to v with $K'(v)$.

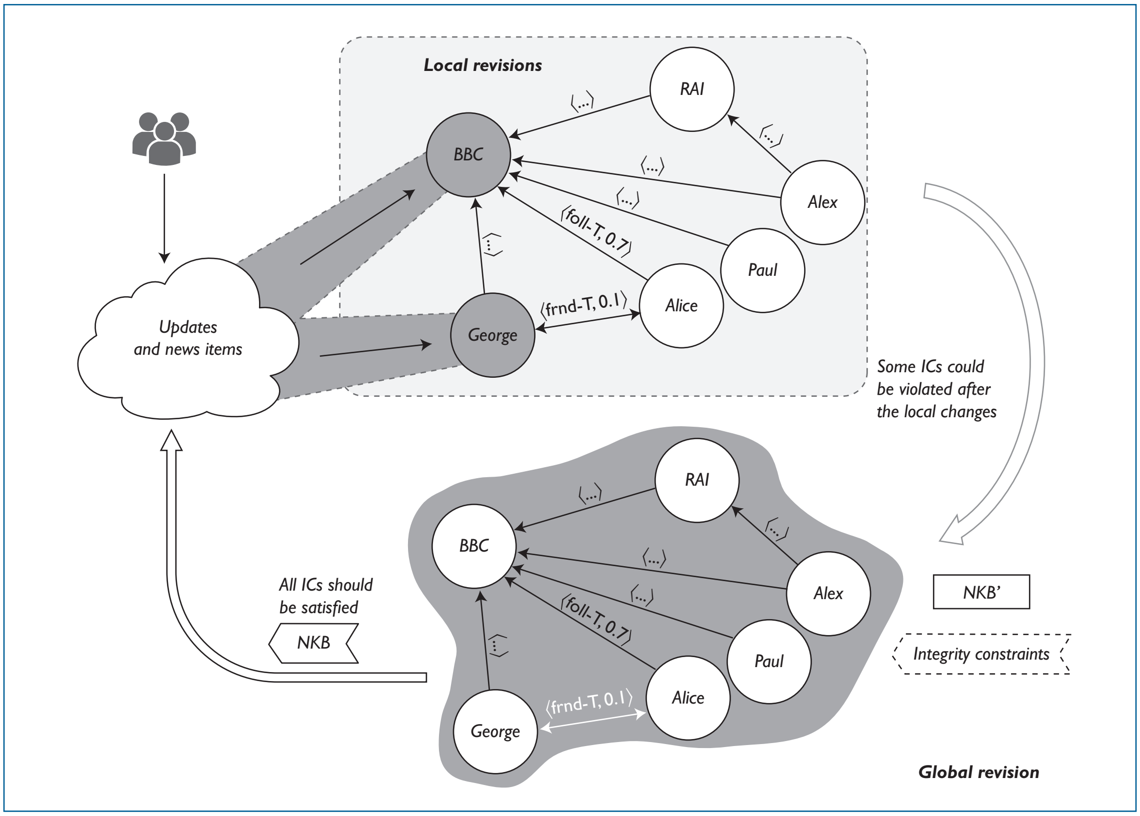

Figure 2. Overview of the local and global revision processes. Updates and news items can come from external sources or as the result of a global revision; for instance, in this case George and Alice are no longer friends on Twitter (perhaps because of a disagreement over the elections). This update can show up in their friends' feeds and thus be part of the next set of local revisions.

Desirable Properties for Social Revision Operators

Following the way in which traditional belief revision operators are constrained by establishing desirable properties (called *postulates*), we propose a set of new ones especially suited for our setting. Some of them are based on the same ideas that can be found in classical postulates; however, given the richer format of the epistemic input, their formalization requires extensive adaptations. On the other hand, the last four postulates are completely novel and specific to NKBs. It should be noted that these properties need not all be satisfied by every operator; as we'll see, different operators arise depending on the subset of properties they enjoy. Let $K(v)$ be the KB of agent v, P be the epistemic input, and $K'(v)$ be the revised KB of agent v:

- *Inclusion*. $K'(v)$ is a subset of the union of $K(v)$ together with all literals in P. Thus, no unwarranted information should be added as part of a revision.
- *Success*. If this property holds, the epistemic input is guaranteed to be accepted; so, the changes contained in P are materialized in $K'(v)$.
- *Weak Success*. Similar to *Success*, except that it only applies when the information in P is consistent with $K(v)$.
- *Consistency*. $K'(v)$ must be consistent.
- *Vacuity 1*. If e isn't inferred from $K(v)$, and all the news items in P that refer to e are of the form $\langle u, e, \mathit{remove}\rangle$, then e shouldn't be inferred from $K'(v)$.
- *Vacuity 2*. As a kind of dual of Vacuity 1, this property states that if an element e is inferred from $K(v)$, and all the news items in P that refer to e or $\neg e$ are of the form $\langle u, \neg e, \mathit{remove}\rangle$ or $\langle u, e, \mathit{add}\rangle$, then e should be inferred from $K'(v)$.
- *Weak Vacuity 1*. If e isn't inferred from $K(v)$, and *none* of the news items in P refer to e, then e shouldn't be inferred from $K'(v)$.

- *Weak Vacuity 2.* As a kind of dual of Weak Vacuity 1, this property states that if e is inferred from $K(v)$, and *none* of the news items in P refer to $\neg e$, then e should be inferred from $K'(v)$.
- *Strong Congruence.* If two sets of news items P and P' are equivalent, the revision of $K(v)$ by P and by P' should be identical; that is, the result doesn't depend on the syntax used to express the epistemic input.
- *Weak Congruence.* As a weaker version of the previous property, if two sets of news items P and P' are identical with respect to the content they *add*, and P removes a superset of what P' removes, then the revision of $K(v)$ by P is a subset of the revision of $K(v)$ by P'.
- *Uniformity.* If two sets of news items P and P' are equivalent, then the revision of K by P and the revision of K by P' should be identical.
- *Majority.* This property refers to the number of news items that are *for* and *against* a certain content e. If the positive items outweigh the negative ones, then e shouldn't be inferred from $K'(v)$; otherwise, e can't be inferred.
- *Weighted Majority.* A generalization of the previous property, which allows votes to be weighted according to the relationship between the agent performing the revision and the origin of each news item.
- *Local Effect.* Applying an operator must not have any effect on other nodes' KBs; so, $K'(u) = K(u)$, for every node u different from v.
- *Structural Preservation.* No vertex, edge, or label is modified by a *local* revision operation.

Working with this list of properties, we focus next on proposing different classes of operators, depending on which ones they satisfy.

Classes of Operators

We begin with a bare minimum, considering that all operators should satisfy Structural Preservation, Local Effect, Consistency, Uniformity, and Inclusion; such operators are called *minimal*.

- A minimal operator is called *Restrained* if it satisfies Strong Congruence, Vacuity 1, and Vacuity 2. This is the most constrained kind of operator — it only has a few opportunities to take any kind of liberty in how the revision is performed.
- A minimal operator is called *Weakly Restrained* if it satisfies Weak Congruence, Weak Vacuity 1, and Weak Vacuity 2. As a weaker form of the previous class, these operators can make different revision decisions, such as removing an element from the local KB when all news items received actually delete its negation.
- A minimal operator is called *Social* if it satisfies Weak Success and either Majority or Weighted Majority. This class focuses mostly on the opinions of others instead of constraining how revisions are made by logical properties.

Next, we'll look at how agent types can be defined.

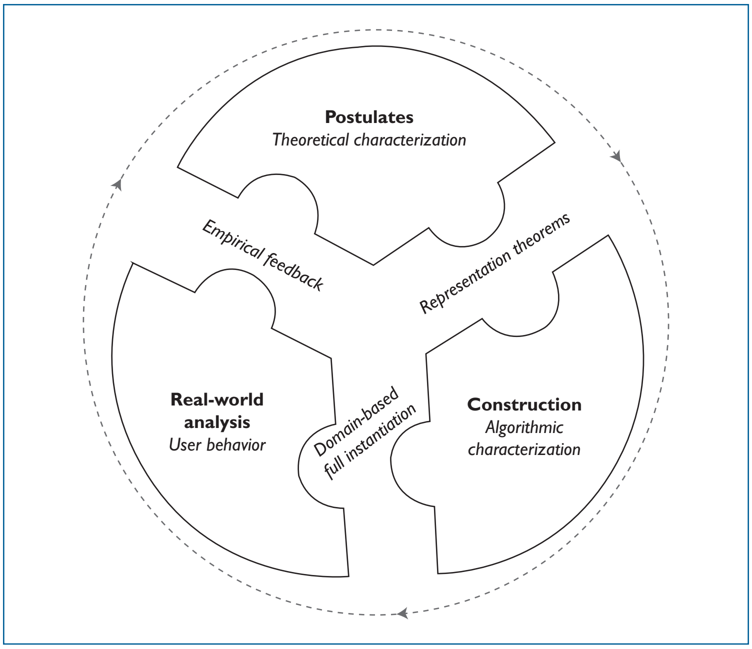

Figure 3. A schematic view of the relation among postulates, operator construction, and empirical evaluation/real-world applications.

Agent Types

Having formulated a general set of properties for local NKB revision operators, we now focus on the next steps in the process of materializing them in actual applications. Figure 3 depicts this process, which is essentially an iterative refinement that moves from the formulation of postulates to algorithms designed to satisfy certain subsets of the postulates, and then to concrete instantiations of such algorithms; the circle closes when real-world feedback informs the formulation and choice of theoretical properties. This process is akin to the spiral-based models of software development in general.

After tackling the theoretical characterization in the previous section, we now move to an initial study of how local NKB operators can be applied in a real-world domain. We perform this step out of order in this first iteration — the goal of this effort is to inform the initial construction of general algorithmic characterizations by selecting appropriate subsets of postulates according to types of users that we can recognize in an actual dataset. Toward this end, and as an orthogonal classification of how revisions are performed, we propose different classes of agents:

1. *Credulous* agents adopt all new knowledge that they see, even if it's against their current beliefs.
2. *Incredulous* agents are reluctant to incorporate new knowledge appearing in their feeds, regardless of whether it has any relation with their previously held beliefs.
3. An agent with *herd behavior* accepts all new information, as long as there are enough agents adopting it.
4. A *blind follower* agent accepts all new knowledge, as long as it's shared by others that are close enough to it.
5. *Cautious* agents don't immediately adopt new knowledge, but rather wait until enough reasons to do so are presented.
6. *Self-confident* agents give more value to their previously held beliefs, making it difficult for them to incorporate new information that contradicts them.

These types of agents allow us to define a set of orthogonal personality traits that each agent can be characterized by, such as *Credulousness*, *Cautiousness*, and *Self-confidence*. Each agent will then have a value for each trait, and together these values define their type.

Example 2. Consider again the scenario from Example 1. Because George and Alice are friends on Twitter and she follows BBC on Facebook, Alice receives news items with both α and $\neg\alpha$, and revises her KB with these inputs. Suppose Alice decides to use the weights of her relations to assign importance to posts, and thus decides to favor BBC's post instead of George's tweet. As a result, α still belongs to $K(Alice)$; if the revision were global, Alice could also decide to defriend George on Twitter (see Figure 2). Here, we could perceive Alice as a self-confident user, because she prefers to keep her prior beliefs unchanged. Her behavior is also in agreement with an operator that satisfies the Weighted Majority postulate.

Experimental Evaluation Using Twitter Data

We carried out a preliminary empirical evaluation to study (and discover) the presence of the aforementioned types of agents in real-world data, with the future objective of grounding and refining the revision operators in order for them to be applicable as models of how agents immersed in social networks behave (see Figure 3).

Our dataset is comprised of 18,292,721 tweets posted between 15 July 2013 and 25 March 2015; of these, 16,780,489 are in English — in this first analysis, we focus only on these. Hashtags are present in 5,107,986 tweets, and there are a total of 136,809 distinct hashtags. Finally, the dataset also includes information regarding who each user follows. This data was collected with the purpose of analyzing various election periods in India; note that the objective of our experiments is to analyze user behavior with respect to incoming information — the actual content of the tweets, or the domain itself therefore isn't relevant, and the same kind of analysis could be performed on other datasets.

The main objective of the present empirical study was to show that by analyzing how information flows in a social network such as Twitter, we can build a map of users that indicate their type (as we introduced in the previous section). This is valuable information to have, because knowing an agent's type is the first step toward understanding their behavior, and ultimately making predictions about it. Toward this end, we used hashtags as proxies for the knowledge of interest conveyed in the tweets — in future work, we'll generalize this by analyzing content more deeply via, for instance, entity extraction and keywords. To understand the way in which each hashtag is referred to, we carried out sentiment analysis using PHPInsight (see https://github.com/JWHennessey/phpInsight). The

*Table 1. Summary of agent activity when receiving news items from their connections.**

Agent	Sentiment distribution (received)	Average behavior (tweeted)
a_1	32% / 19% / 49%	21% (76% / 3%) 2% (95% / 3%) 70% (24% / 6%)
a_2	33% / 17% / 50%	7% (93% / 0%) 14% (86% / 0%) 90% (10% / 0%)
a_3	33% / 18% / 49%	6% (79% / 15%) 0% (82% / 18%) 78% (6% / 16%)
a_4	31% / 23% / 46%	67% (7% / 26%) 0% (70% / 30%) 6% (65% / 29%)
a_5	26% / 19% / 55%	0% (70% / 30%) 0% (54% / 46%) 59% (0% / 41%)
a_6	44% / 2% / 44%	0% (75% / 25%) 0% (50% / 50%) 75% (0% / 25%)
a_7	34% / 16% / 50%	0% (72% / 28%) 0% (60% / 40%) 73% (0% / 27%)
a_8	11% / 2% / 87%	0% (90% / 10%) 0% (65% / 35%) 97% (0% / 3%)
a_9	33% / 22% / 45%	2% (59% / 39%) 3% (68% / 29%) 56% (6% / 38%)
a_{10}	33% / 18% / 49%	7% (91% / 2%) 17% (80% / 3%) 79% (19% / 2%)

* The second column contains the distribution of the sentiment in the tweets seen by each agent (positive, negative, and neutral, respectively). The third column describes how the agent reacted: the first line contains the percentage of times that it was also positive, while in parentheses we include first the percentage of times that it changed the sentiment and second the percentage of times it didn't use the hashtag at all; the other two lines are analogous for negative and neutral.

sentiment analysis over a given tweet yields one out of three possible results: positive, negative, or neutral.

To begin, we identified the 50 most prevalent hashtags, which appear in 1,466,169 distinct tweets by 337,492 users. Out of these users, we selected the 10 that had the most activity to analyze their behavior with respect to the hashtags they see in their feeds; it's easy to expand this number, but we kept it low for presentation purposes. For each of these users, we analyzed the sentiment of the tweet via which each hashtag reached their feed (positive, negative, or neutral), and then analyzed how the user reacted — the possible reactions were to tweet using the hashtag (again, with positive, negative, or neutral sentiment) or to not use that hashtag in any of their subsequent tweets.

Table 1 shows the results; for each agent, the table summarizes how they behaved on average over all hashtags, and provides clues as to their types. For instance, agent a_8 is likely to be of the type *self-confident*: it never reuses hashtags with the same sentiment when it receives them with positive or negative — instead, it uses a different sentiment 90 percent and 65 percent of the time, respectively. Other agents (a_3, a_5, a_6, a_7, and a_9) exhibit similar behavior; this makes sense, because we chose the most active users, who are presumably highly opinionated individuals. Expanding the number of users in further experiments will likely expose other kinds of active users, such as those that exhibit *herd behavior* by reusing hashtags using the same sentiment with which they receive them.

Agent a_4 is an interesting case: when receiving tweets with positive sentiment, it's likely to pass it on (67 percent of the time), while only changing sentiment 7 percent of the time. On the other hand, when receiving negative sentiment, it's likely to change it (70 percent of the time).

Another observation is that we can easily detect agents that are highly prone to respond to their connections' comments.

Both observations suggest that the self-confident type of agents could be further refined: the former group could be seen as adversarial, while a_4 is conciliatory. Further experiments will shed light on this refinement (and possibly others). Another observation is that we can easily detect agents that are highly prone to respond to their connections' comments (such as a_1, a_2, and a_{10}) or — on the other end of the spectrum — "dead end" users who don't reuse the hashtags that they see in their feeds.

The main focus of this work was to lay the foundation for modeling the dynamics of data and knowledge flowing through social media; our work is based on the NKB model. Based on a multiagent systems approach, we abstract the combination of multiple social media platforms into a single NKB; this gives rise to complex networks in which agents have local KBs where their beliefs are stored. Each agent has multiple attributes describing their features, as well as different kinds of relations with other agents — these relations can also have attributes that characterize them. We showed how agents' activities and other network updates can be represented by basic elements called news items; these are the starting points of local revisions performed by each agent, and thereafter global revisions are also carried out (at the network KB administrator level) to ensure that integrity constraints aren't violated. We specified a set of rationality postulates as desirable properties for social revision operators at the local level, and proposed several agent types that can occur in real-world domains. Finally, we also presented the results of a preliminary experimental evaluation performed on Twitter data collected during several electoral periods in India, analyzing the behavior and identifying clues pointing to the types of each agent from real-life interactions when receiving and sharing hashtags with an associated sentiment.

Future work involves formalizing the local and global revision operators based on observed (learned) types of agents — that is, working on the "Construction" piece of the puzzle in Figure 3 using as a basis both the postulates and what we learned so far from empirical inquiry. We'll also continue the empirical evaluation by performing experiments that validate these operators' effectiveness. For this, the next two important steps are to aggregate data from multiple social platforms, and to develop a systematic approach to generating synthetic data. This will allow us to adequately evaluate different aspects of the implemented systems by varying key parameters as needed.

Acknowledgments

This work was supported by funds provided by CONICET, Agencia Nacional de Promoción Científica y Tecnológica, and Universidad Nacional del Sur, Argentina. Some of the authors of this work were also supported by the US Department of the Navy, Office of Naval Research, grant N00014-15-1-2742. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research. We thank V.S. Subrahmanian for providing us with the India elections dataset.

References

1. B.L. Ott, "The Age of Twitter: Donald J. Trump and the Politics of Debasement," *Critical Studies in Media Comm.*, vol. 34, no. 1, 2017, pp. 59–68.
2. S. Oates and W.W. Moe, "Donald Trump and the 'Oxygen of Publicity': Branding, Social Media, and Mass Media in the 2016 Presidential Primary Elections," *Am. Political Science Assoc. Ann. Meeting*, 2016; dx.doi.org/10.2139/ssrn.2830195.
3. C. Shao et al., "Hoaxy: A Platform for Tracking Online Misinformation," *Proc. World Wide Web Companion*, 2016, pp. 745–750.
4. H. Allcott and M. Gentzkow, "Social Media and Fake News in the 2016 Election," *Nat'l Bureau of Economic Research*, tech. report 23089, 2017; www.nber.org/papers/w23089.
5. I. Lapowswky, "Here's How Facebook Actually Won Trump the Presidency," *Wired*, Nov. 2016; www.wired.com/2016/11/facebook-won-trump-election-not-just-fake-news.
6. M.J. Wooldridge, *An Introduction to MultiAgent Systems*, second ed., John Wiley and Sons, 2009.
7. F.R. Gallo et al., "A Desiderata for Modeling and Reasoning with Social Knowledge," *Proc. CACIC*, 2015.
8. C.E. Alchourrón, P. Gärdenfors, and D. Makinson, "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions," *The J. Symbolic Logic*, vol. 50, no. 2, 1985, pp. 510–530.
9. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, 1988.
10. S.O. Hansson, "Kernel Contraction," *J. Symbolic Logic*, vol. 59, no. 3, 1994, pp. 845–859.
11. S. Hansson, "Semi-Revision," *J. Applied Non-Classical Logics*, vol. 7, nos. 1-2, 1997, pp. 151–175.
12. P. Peppas, "Belief Revision," *Handbook of Knowledge Representation*, 2008, pp. 317–359.
13. S. Konieczny and R.P. Pérez, "On the Logic of Merging," *Proc. Knowledge Representation and Reasoning*, 1998, pp. 488–498.
14. S. Konieczny and R. Pino Pérez, "Merging Information under Constraints: A Logical Framework," *J. Logic and Computation*, vol. 12, no. 5, 2002, pp. 773–808.
15. P. Liberatore and M. Schaerf, "Arbitration (or How to Merge Knowledge Bases)," *IEEE Trans. Knowledge and Data Eng.*, vol. 10, no. 1, 1998, pp. 76–90.
16. M.A. Falappa et al., "Prioritized and Non-Prioritized Multiple Change on Belief Bases," *J. Philosophical Logic*, vol. 41, no. 1, 2012, pp. 77–113.
17. F.R. Gallo et al., "Belief Dynamics in Complex Social Networks," *Proc. Simposio Argenteno de Inteligencia Artificial*, vol. 45, 2016, pp. 70–77.

Fabio R. Gallo is a doctoral student in computer science at Universidad Nacional del Sur. His research interests are centered in knowledge representation and reasoning, with a special focus on social knowledge. Contact him at fabio.gallo@cs.uns.edu.ar.

Gerardo I. Simari is an assistant professor in the Department of Computer Science and Engineering, Universidad Nacional del Sur, and a researcher in the Institute for Computer Science and Engineering (Universidad Nacional del Sur–CONICET). His research interests include reasoning under uncertainty, preferences, social knowledge, and databases. Simari has a PhD from the University of Maryland, College Park. Contact him at gis@cs.uns.edu.ar.

Maria Vanina Martinez is a researcher at the Institute for Computer Science and Engineering, Universidad Nacional del Sur–CONICET. Her research interests are at the intersection of knowledge representation and reasoning and database theory, with a special focus on reasoning with inconsistent and uncertain knowledge. Vanina Martinez has a PhD in computer science from the University of Maryland, College Park. Contact her at mvm@cs.uns.edu.ar.

Marcelo A. Falappa is an associate professor and dean of the Department of Computer Science and Engineering, Universidad Nacional del Sur, and a researcher at the Institute for Computer Science and Engineering (Universidad Nacional del Sur–CONICET). His research interests are centered in belief revision and argumentation-based reasoning. Falappa has a PhD in computer science from Universidad Nacional del Sur. Contact him at mfalappa@cs.uns.edu.ar.

Natalia Abad Santos is an assistant professor in the Department of Mathematics, Universidad Nacional del Sur. Her research interests are focused on reasoning under uncertainty, particularly knowledge representation tools for probabilistic reasoning. Contact her at nasantos@uns.edu.ar.