

Belief Dynamics in Complex Social Networks

Fabio R. Gallo¹, Natalia Abad Santos², Gerardo I. Simari¹,
 Maria Vanina Martinez¹, and Marcelo A. Falappa¹

¹ Institute for Computer Science and Engineering (CONICET–UNS)
 Department of Computer Science and Engineering, Universidad Nacional del Sur
 Av. Alem 1253 – B8000CPB Bahia Blanca – Argentina

² Department of Mathematics, Universidad Nacional del Sur
 Av. Alem 1253 – B8000CPB Bahia Blanca – Argentina
 {fabio.gallo,gis,mvm,mfalappa}@cs.uns.edu.ar, nasantos@uns.edu.ar

Abstract. People are becoming increasingly more connected to each other in social media networks. These networks are complex because in general there can be many different types of relations, as well as different degrees of strength for each one; moreover, these relations are dynamic because they can change over time. In this context, users’ knowledge flows over the network, and modeling how this occurs—or can possibly occur—is therefore of great interest from a knowledge representation and reasoning perspective. In this paper, we focus on the problem of how a single user’s knowledge base changes when exposed to a stream of news items coming from other members in the network. As a first step towards solving this problem, we identify possible solutions leveraging preexisting belief merging operators, and conclude that there is a gap that needs to be bridged between the application of such operators and a principled solution to the proposed problem.

Keywords: Social Web, Complex Networks, Belief merging operators.

1 Introduction and Related Work

Nowadays, most people have an account in at least one social network; such networks can be seen as dynamic environments containing knowledge about the relationships that exist among their members. From a knowledge representation and reasoning perspective, there are certain aspects that make these environments *complex*; for instance, it is not evident how news being communicated through the network should be treated, especially when the news contradicts a user’s own knowledge, or when contradicting news is received from different sources. To make things worse, a user could share an item—making it visible to their friends—even when this user does not adopt the knowledge it contains.

In this work we focus on the problem that arises when a user must integrate many sources of knowledge coming from their connections in a (complex) social network. We first formalize the notion of complex social networks and the aforementioned problem in a simplified setting, and then explore how belief merging

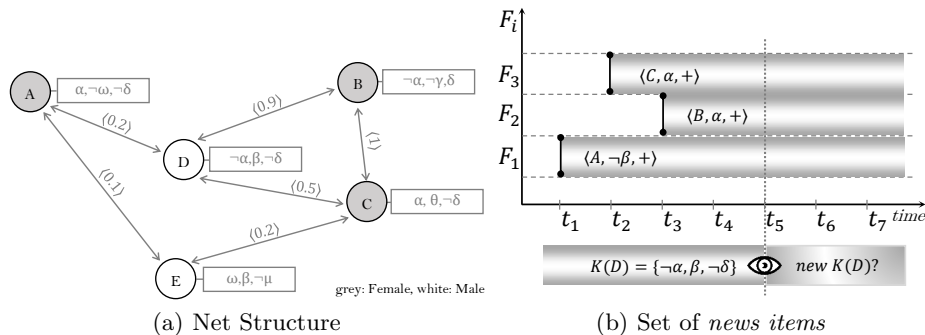


Fig. 1. *Left:* Example complex social network modeling weighted friendship and local KBs; *Right:* Depiction of news items as seen from user D 's perspective.

operators from the literature could be applied towards solving our problem. We propose and analyze three different possibilities for doing this, and in each case conclude that they are not adequate as a general solution.

Example 1 (Social Media Network). We now describe a setting that we use as a running example to motivate and illustrate the presentation. Consider a social media network (like Facebook, Google+, Instagram, or Twitter) where people are connected to each other, and share thoughts, ideas, or multimedia content, post comments, etc. We wish to also model different levels of strength over relations (*weight*). Figure 1(a) shows a network consisting of five users; their (bidirectional, in this case) relationships are represented with labeled edges—though in general multiple relationships can exist between two people, for simplicity in this paper we will assume a single “friendship” relation, where labels indicate how close they are as a real value in $[0, 1]$. Weights could, for instance, be calculated as a function of how long they have known each other, their common interests, etc. Furthermore, each vertex has associated a local *set of beliefs*. Here, vertex A has belief base $\{\alpha, \neg\omega, \neg\delta\}$. ■

In this setting, posts made by users are seen by *all* of their connections; when this occurs, we assume that each user adopts a position regarding the information that they receive from their friends. For instance, it could be seen but ignored, or perhaps an evaluation is made to see if the knowledge has enough epistemic value to be integrated into their own knowledge base. In this context, users generally have many possible sources of information (friends) that could make mutually inconsistent posts, and/or posts that contradict their local belief base. In such cases, it would be natural for users to resort to the closeness of the friend who made the post, how many friends share the opinion, etc.

To approach the belief merging problem presented in Example 1, we consider some of the most important merging operators that have been developed in the literature. *Belief merging* [5, 6] refers to the problem of coherently combining several sources of information. When merging, new evidence can be partially or

completely ignored if old information has more epistemic value. The operation integrates all the information to a consistent whole, whose result depends on informational value and not necessarily novelty. A considerable amount of work has been developed in this area, in particular regarding axiomatic characterizations. In [5], the authors propose an axiomatic characterization of merging operators, where operators satisfying those axioms are called *pure merging operators*. Then, in [6], the problem of merging several belief bases in the presence of integrity constraints is considered. Furthermore, the authors show that Liberatore and Schaerf commutative revision operators [7] are a particular class of merging operators. Another line of work is related to *majority* and *arbitration* operators. The former is the class of operators in which a piece of information will persist in the result of merging if it came from many sources. The latter focus on consensual behavior in extracting information from several sources. Later, this line of work is continued in [1] by dealing with merging operators in a finite logical framework (countably many propositional variables). Their main result is a representation theorem, where certain postulates had to be redefined to be appropriate for the infinite case. Then, in [2] the authors investigate the use of a unanimity condition, i.e., to accept pieces of information shared by all agents. Also related to merging are *multiple change belief revision operators*; according to [3], there are two types of multiple change operator where the epistemic input is a set of sentences: *prioritized change* (the input set is fully accepted) and *symmetric change* (input sentences could be accepted or rejected).

In the following, we first propose a formalization of *Network Knowledge Bases*, and then discuss possible ways in which the problem of knowledge integration in this context can be addressed using existing belief merging operators.

2 Network Knowledge Bases

We assume a language \mathcal{L} built from a finite set of propositional symbols \mathcal{P} and constants, in which the only connective is \neg . Therefore, all elements in our language are ground literals. We also assume the existence of two arbitrary but fixed disjoint sets VP and EP of vertex and edge predicate symbols, respectively, such that $VP \cap \mathcal{P} = \emptyset$ and $EP \cap \mathcal{P} = \emptyset$. Vertex predicates have arity 1, and edge predicates have arity 2. We recall the definition of a Social Network from [8].

Definition 1 ([8]). *A Social network is a 4-tuple $(V, E, l_{vert}, l_{edge})$ where:*

1. V is a finite set whose elements are called vertices.
2. $E \subseteq V \times V$ is a finite set whose elements are called edges.
3. $l_{vert} : V \rightarrow 2^{VP}$ is a function called a vertex labelling function.
4. $l_{edge} : E \rightarrow 2^{\mathcal{T}}$ is a function called an edge labelling function, where $\mathcal{T} = \{\langle b, w \rangle \mid b \in EP, w \in [0, 1]\}$.

For our purposes, we associate a set of propositional literals with each vertex.

Definition 2. *A Network KB (NKB for short) is a 5-tuple $(V, E, l_{vert}, l_{edge}, K)$, where the first four elements comprise a social network, and $K : V \rightarrow 2^{\mathcal{L}}$ is a*

mapping assigning a knowledge base to each vertex; $K(v_i)$ is called the knowledge base (*KB*, for short) associated with vertex v_i .

Additionally, we further enrich the notion of network with a set of constraints that conditions (and relates) both the structural part of the network (*i.e.*, its relationships) and the knowledge bases of their users.

Definition 3. A constraint C over a NKB $(V, E, l_{vert}, l_{edge}, K)$ is a pair (S, B) where, given $\{v_1, \dots, v_n\} \in V$, and $e_1, \dots, e_m \in E \cap \{v_1, \dots, v_n\} \times \{v_1, \dots, v_n\}$,

1. S , called the structural part, contains a conjunction of conditions each of which can be of the following forms:
 - $l_{vert}(v) = a$, $a \in VP$, $v \in \{v_1, \dots, v_n\}$;
 - $\langle b, w \rangle \in l_{edge}(e)$, $\langle b, w \rangle \in \mathcal{T}$, $w \in [\alpha, \beta]$, for some $0 \leq \alpha \leq \beta \leq 1$.
2. B is called the belief part and contains a conjunction of conditions involving elements in $K(v_1), \dots, K(v_n)$.

Informally, given a constraint $C = (S, B)$, we say a set of vertices $\{v_1, \dots, v_n\}$ and edges $\{e_1, \dots, e_m\}$ satisfy S if all the conditions of the conjunction hold when we replace the vertices and the edges properly. We say that a set of vertices $\{v_1, \dots, v_n\}$ satisfies B if there exists a way to replace the vertices such that the conjunction holds. Finally, an NKB satisfies C if for every set of vertices and edges it holds that whenever S is satisfied, B is also satisfied. Given this semantics, it may also be convenient to write constraints as $S \rightarrow B$.

Example 2. Let $v_1, v_2 \in V$, and $e(v_1, v_2), e(v_2, v_1) \in E$; we have a constraint $C_1 = (S_1, B_1)$ defined by:

1. S_1 : $l_{vert}(v_1) = l_{vert}(v_2) = \text{female}$ and $\langle \text{friend}, 1 \rangle \in l_{edge}(e)$;
2. B_1 : If $p \in \mathcal{P}$ and $K(v_1) \models p$ then $K(v_2) \not\models \neg p$.

This constraint says that if two people are *female* and *friends* with a *weight* of 1 (S_1) then their knowledge bases should not have contradictory elements (B_1). In Figure 1(a) we can see that C_1 is not satisfied since there is an edge $\langle \text{friends}, 1 \rangle$ between vertex B and C such that δ belongs to $K(B)$ and $\neg\delta$ belongs to $K(C)$.

Consider another constraint $C_2 = (S_2, B_2)$ where: S_2 : If $\langle \text{friend}, w \rangle \in l_{edge}(e)$ then $w \in (0, 1)$ and B_2 : \emptyset . Constraint C_2 expresses that when two people are friends, the associated *weight* should be neither 0 nor 1—note that the belief part in this constraint is empty. In this case we can see in Figure 1(a) that C_2 is not satisfied since there is an edge labeled with *friend* between B and C with *weight* 1. If that weight were, say, 0.9, constraint C_2 would be satisfied. ■

We are now ready to define *consistency* of an NKB.

Definition 4. Given a finite set of constraints IC , we say that an NKB $= (V, E, l_{vert}, l_{edge}, K)$ is consistent *w.r.t.* IC if it satisfies all constraints in IC .

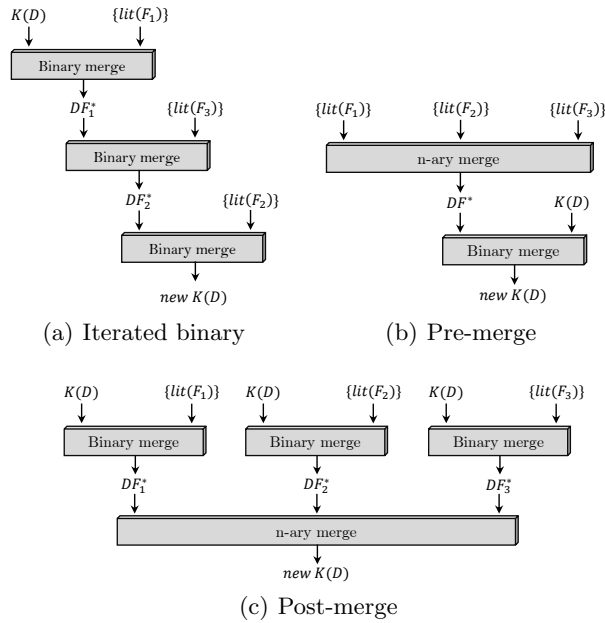


Fig. 2. Illustration of the possible solutions proposed.

The *network belief dynamics problem* arises when the epistemic input is comprised of a set of what we call *news items* coming from the neighbor nodes. A news item consists of a triple $\langle o, l, d \rangle$, where o is the *origin* (the node that originates the change), l is a literal representing new knowledge, and d is an indication of its new status in the origin node's belief base: $+$, $-$, or f , which represent *added* (l was added to $K(o)$), *flipped*, ($\neg l$ was changed to l in $K(o)$) or *forgotten* (l was removed from $K(o)$), respectively.

3 Possible Solutions using Existing Operators

In this section we discuss three possible solutions to deal with integration of information coming from news items in a complex social network; we use our running example to illustrate each one. Consider then the NKB from Figure 1(a), and suppose user D had a bad experience with a certain product brand S —he now believes that products of brand S should not be purchased (denoted with $\neg\alpha$). On the other hand, he believes that brand T offers excellent products (denoted with β). Independently, users A , B , and C make three different posts about brands S and T , at moments t_1 , t_3 and t_2 , respectively (cf. Figure 1(b)¹). Then, user D logs in at some future moment t_5 and sees the posts made by his friends, and realizes that they express different opinions about brand S . User D

¹ Time points are used for illustrative purposes; they are not currently formalized.

needs to decide how to react to those *news items*—in particular, he must decide how they impact his current knowledge base ($K(D) = \{-\alpha, \beta, -\delta\}$).

We first introduce some notation and definitions; in the following let $NKB = (V, E, l_{vert}, l_{edge}, K)$, $v_i \in V$, and $F = \{F_1, \dots, F_n\}$ be a set of n news items—for simplicity, in this work we will assume that $IC = \emptyset$. Given $NI = \langle o, l, d \rangle$, let $lit(NI) = l$ be the function that returns the literal involved in NI . We say that the set F is *consistent* if and only if $\bigwedge_{F_i \in F} lit(F_i) \not\models \perp$. Given NKB and set of news items F , the following properties must hold:

- For each $v_i \in V$, $K(v_i) \not\models \perp$; individual knowledge bases are consistent.
- $\{lit(F_i)\} \cup K(v_i) \not\models \perp$; when someone shares something, they cannot believe the contrary.

In the following, let “ $*$ ” be a binary operator that can stand for either a belief merging or belief revision operator; thus, it is a function of the form $* : 2^{\mathcal{L}} \times 2^{\mathcal{L}} \rightarrow 2^{\mathcal{L}}$. Similarly, let \otimes be an n -ary belief merging operator of the form $\otimes : 2^{\mathcal{L}} \times \dots \times 2^{\mathcal{L}} \rightarrow 2^{\mathcal{L}}$.

The first possible solution we propose to the network belief change problem is called *iterated binary*.

Solution 1 (Iterated binary) Let Γ be a function taking an NKB, a node in the network, and a set of news items, such that:

$$\Gamma(NKB, v, \{F_1, \dots, F_n\}) = \Gamma(NKB', v, \{F_1, \dots, F_{n-1}\})$$

where NKB' is identical to NKB except that $K(v)$ is replaced with $*(K(v), \{lit(F_n)\})$. Finally, the base case is simply $\Gamma(NKB, v, \emptyset) = NKB$.

Function Γ emulates a user reading the posts one by one and iteratively applying operator $*$ to update their knowledge base.

Example 3 (Function Γ). To update user D 's knowledge base according to the news items in Figure 1(b), we apply function Γ in the following manner: $\Gamma(NKB, D, \{F_1, F_2, F_3\}) = (*((* (K(D), \{lit(F_1)\}), \{lit(F_3)\}), \{lit(F_2)\})$. Here, we are assuming that posts are shown to user D sorted by *time* (thus the sequence $\langle F_1, F_3, F_2 \rangle$); different sequences are of course possible, and the result can be different depending on the sequence used. Figure 2(a) depicts this process.

Another criterion to choose a sequence could be based on the *weight* of the relation between user D and each user who posted a news item involved in the belief change operation. Going back to Figure 1(a), weights between D - A , D - B and D - C are 0.2, 0.9, and 0.5 respectively. According to this criterion, the sequence would be $\langle F_2, F_3, F_1 \rangle$. Then, function Γ would be: $\Gamma(NKB, D, \{F_1, F_2, F_3\}) = *((* (* (K(D), \{lit(F_2)\}), \{lit(F_3)\}), \{lit(F_1)\})$.

Suppose that posts are shown to user D ordered by time. He first sees F_3 and realizes that user C 's opinion is in contradiction with his own (D believes $-\alpha$ and C believes α); suppose D decides not to change his knowledge base. Then he sees F_2 , where user B agrees with C ; suppose that once more D decides not to make any changes. ■

Note that function I has no way of “remembering” what happened in previous applications; this is a drawback of the approach because perhaps if user D encounters α for a second time he could take this into account and make a different choice. Another drawback of Solution 1 is that it doesn’t model how users generally change their beliefs—when somebody reads a post, they don’t usually immediately decide what to do with the information it contains, but rather do so after a stream of several posts about a topic are seen; this approach thus also loses the *global view* of all news items together.

The second possible solution is called *pre-merge*, and is described next.

Solution 2 (Pre-merge) Let Π be a function taking an NKB, a node in the associated network, and a set of n news items, defined as follows:

$$\Pi(NKB, v, F) = * \left(K(v), \otimes(\{lit(F_1)\}, \dots, \{lit(F_n)\}) \right).$$

The function first performs a merge of the contents of all the news items, and then uses the binary revision operator to decide how the result is incorporated into the user’s knowledge base.

Example 4 (Function Π). To update D ’s knowledge base according to the news items in the running example, we apply function Π in the following manner (cf. Figure 2(b) for a depiction of the process): $\Pi(NKB, D, \{F_1, F_2, F_3\}) = *(K(D), \otimes(\{lit(F_1)\}, \{lit(F_2)\}, \{lit(F_3)\}))$. ■

Though this solution addresses the issue of considering all posts together, its main drawback is that we lose information regarding the *source* of each individual F_i . Having this information could be important because a user might want to take into account the weight of their relation when deciding if and how to change their KB (or modify the weight in the network). Another potential drawback is that the n -ary merge operation integrates knowledge coming from users even if they have no relation (such as users A and B in our example).

The last potential solution is a kind of dual of Solution 2.

Solution 3 (Post-merge) Let Ξ be a function taking an NKB, a node in the associated network, and a set of n news items, defined as:

$$\Xi(NKB, v, F) = \otimes \left(K(v) * lit(F_1), \dots, K(v) * lit(F_n) \right)$$

The operation is performed by first applying the binary operator between the user’s knowledge base and each news item’s content, and then applying the merging operator to the n results.

Example 5 (Function Ξ). Once again, let’s see how to update D ’s knowledge base with the news items from Figure 1(b) by applying function Ξ . Figure 2(c) depicts this process; we have: $\Xi(NKB, D, \{F_1, F_2, F_3\}) =$

$$\otimes \left(K(D) * \{lit(F_1)\}, K(D) * \{lit(F_2)\}, K(D) * \{lit(F_3)\} \right).$$

■

The main drawback of this solution—similarly to Solution 2—is that we lose information regarding the *source* of each individual F_i . Additionally, we cannot know which users eventually posted the *same* knowledge (as in Solution 1). For instance, we drop the fact that both B and C have posted α (cf. Figure 1(b)).

4 Conclusions and Future Work

In this paper we have formalized the problem of performing belief dynamics in a complex social network setting, in which users have their own knowledge bases but interact with each other by posting content that is visible to their connections. This formalization is part of a line of work started recently in [4], where we proposed the concept of *Social Knowledge Bases*.

We proposed three possible solutions to solving the problem of integrating multiple news items into a user’s knowledge base by leveraging existing approaches in the belief revision and merging literature; unfortunately, in each case we identified drawbacks that make the proposed solutions inadequate in general. Such drawbacks hold even in the simplified case in which there are no network integrity constraints. The main conclusion is thus that a new kind of operator is needed to address this problem: an *n-ary merging operator* capable of dealing with the unique characteristics of networked KBs and news items.

Future work involves formalizing operators by developing postulates that are adequate for this setting, proposing different constructions and proving representation theorems. We shall consider first the simplified social networks used in this paper, and gradually incorporate richer features—an advantage of this approach is that sources of computational complexity are more readily identified.

Acknowledgments. This work was supported by funds provided by Universidad Nacional del Sur and CONICET, Argentina, and by the U.S. Department of the Navy (ONR grant N00014-15-1-2742).

References

1. Chacón, J.L., Pino Pérez, R.: Logic-based merging: the infinite case. In: Proc. NMR. pp. 100–108 (2004)
2. Everaere, P., Konieczny, S., Marquis, P.: Disjunctive merging: Quota and gmin merging operators. *Artificial Intelligence* 174(12), 824–849 (2010)
3. Falappa, M., Kern-Isberner, G., Reis, M., Simari, G.R.: Prioritized and non-prioritized multiple change on belief bases. *J. Philos. Log.* 41(1), 77–113 (2012)
4. Gallo, F.R., Abad Santos, N., Simari, G.I., Falappa, M.A.: A desiderata for modeling and reasoning with social knowledge. In: Proc. CACIC (2015)
5. Konieczny, S., Pino Pérez, R.: On the logic of merging. In: Proc. KR. pp. 488–498 (1998)
6. Konieczny, S., Pino Pérez, R.: Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12(5), 773–808 (2002)
7. Liberatore, P., Schaerf, M.: Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering* 10(1), 76–90 (1998)
8. Shakarian, P., Broecheler, M., Subrahmanian, V.S., Molinaro, C.: Using generalized annotated programs to solve social network diffusion optimization problems. *ACM Trans. Comput. Logic* 14(2), 10:1–10:40 (2013)