A First Approach to Belief Dynamics in Complex Social Networks

Fabio R. Gallo[1], Gerardo I. Simari[1], Maria Vanina Martinez[1], Natalia Abad Santos[2], and Marcelo A. Falappa[1]

[1] Institute for Computer Science and Engineering (UNS–CONICET)
Dept. of Comp. Science and Eng., Universidad Nacional del Sur (UNS), Argentina
[2] Department of Mathematics, Universidad Nacional del Sur (UNS), Argentina
{fabio.gallo,gis,mvm,mfalappa}@cs.uns.edu.ar,nasantos@uns.edu.ar

1 Introduction and Related Work

Many people today participate in more than one social network on a daily basis; such networks can be seen as dynamic environments containing knowledge about their members. From a KR perspective, these environments are complex in many dimensions; for instance, it is not clear how to treat news being communicated through the network, especially when it contradicts a user's own knowledge, or when contradicting news is received from different sources. We are interested in reasoning about the *diffusion of beliefs* in social media. Given the variety of platforms, the underlying communication structure can be seen as a *complex network*—the potential to integrate many different data sources is enormous. The complex network is represented by a graph, and users have access to a stream of *news items* that are produced by others. Local revisions are performed first, where each user responds to news items that show up in their feeds; since each local revision is carried out in parallel, the result of this stage could violate global integrity constraints. A global revision process is then carried out in order to return to a consistent state. In this paper we focus solely on local revisions.

Related Work. Networks have been used to model diffusion processes in real-world domains [4]; more ad hoc models have also been central to world events like the US elections and the Brexit vote [7, 1]. Our research continues the work of [6], where the authors propose a general formalism to model complex networks. Whereas their work focuses on cascading processes, they do not contemplate individual knowledge bases for each agent. Our end goal is to build on the first steps—modelling local revisions—to eventually model the combination of these processes as they give rise to cascades. This work is also part of the research line started in [2], where we proposed a model called *Social Knowledge Bases*.

2 Formalizing Local Revisions

We assume a language $\mathcal{L}$ of ground literals over symbols in *Pred*. Two literals l_1, l_2 are *contradictory* iff $l_1 = \neg l_2$. We also assume disjoint sets *VP* and *EP* of vertex and edge predicate symbols, resp., such that $VP \cap Pred = \emptyset$ and

$EP \cap Pred = \emptyset$. Vertex predicates have arity 1, and edge predicates have arity 2; we use $\mathcal{L}_{\mathcal{V}}$ (resp., $\mathcal{L}_{\mathcal{E}}$) to denote literals over VP (resp., EP).

Definition 1 ([5]). *A* Social network *is a 4-tuple* $(V, E, l_{vert}, l_{edge})$ *where:* V *is a finite set of* vertices, $E \subseteq V \times V$ *is a finite set of* edges, $l_{vert} : V \to 2^{\mathcal{L}_{\mathcal{V}}}$ *is a* vertex labelling function, *and* $l_{edge} : E \to 2^{\mathcal{T}}$ *is an* edge labelling function, *where* $\mathcal{T} = \{\langle b, w\rangle \mid b \in \mathcal{L}_{\mathcal{V}}, w \in [0, 1]\}$.

Definition 2 ([3]). *A* Network KB *(NKB, for short) is a 5-tuple* $(V, E, l_{vert}, l_{edge}, K)$, *where* $(V, E, l_{vert}, l_{edge})$ *is a social network, and* $K : V \to 2^{\mathcal{L}}$ *is a mapping;* $K(v_i)$ *is called the* KB *associated with* v_i.

We enrich networks with a set of constraints that condition and relate the structural part and users' KBs; satisfaction and consistency are defined as expected.

Definition 3 ([3]). *A* constraint C *over an NKB* $(V, E, l_{vert}, l_{edge}, K)$ *is a pair* (S, B) *where, given* $v_1, ..., v_n \in V$, *and* $e_1, ..., e_m \in E \cap \{v_1, ..., v_n\} \times \{v_1, ..., v_n\}$*:* S, *called the* structural part, *contains a conjunction of conditions that can be of either of the following forms:* $l_{vert}(v) = a$, $a \in VP$, $v \in \{v_1, ..., v_n\}$*; or* $\langle b, w\rangle \in l_{edge}(e)$, $\langle b, w\rangle \in \mathcal{T}$, $w \in [s_1, s_2]$, *for some* $0 \leq s_1 \leq s_2 \leq 1$. B *is called the* belief part *and contains a conjunction of conditions about* $K(v_1), ..., K(v_n)$.

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.

Definition 4 ([3]). *A* news item *is a triple* $\langle o, l, d\rangle$, *where* $o \in V$ *is the* origin, $l \in \mathcal{L}$ *is a literal, and* d *is an indication of its new status in* o*'s KB:* a, r, *or* f, *which represent* added, removed, *or* flipped, *resp. A set of news items is* consistent *if it does not contain items with the same origin and literal, but a different decision or the same origin and conflicting literals.*

We assume that all sets of news items are consistent; function χ changes news items into *canonical form* without flips—each flip is changed to add. We denote with $\mathcal{NKB}$ the universe of NKBs, and $\mathcal{P}$ the universe of news items. Given NKB $= (V, E, l_{vert}, l_{edge}, K)$, set of news items $P = \{p_1, ..., p_n\}$, $v \in V$, we use $P_v = \{p = \langle o, l, d\rangle \in P \mid (o, v) \in E\}$ to refer to the set of news items seen by v. Given $p = \langle o, l, d\rangle \in \mathcal{P}$, let $src(p) = o$, $lit(p) = l$, and $dec(p) = d$. Posts made by users are seen by all their connections; when this occurs, each user adopts a position regarding the information that they receive from their friends. Users generally have many possible sources of information that can make mutually inconsistent posts, or posts that contradict their own local belief base.

Postulates for Local NKB Belief Revision

Definition 5. Local BR operators *are functions* $\circledast : \mathcal{NKB} \times V \times 2^{\mathcal{P}} \to \mathcal{NKB}$.

Let $v \in V$, $\circledast$ be a local NKB revision operator, and $\textsc{nkb}' = \circledast(\textsc{nkb}, v, P) = (V', E', l'_{vert}, l'_{edge}, K')$. We propose a set of postulates as reasonable properties for local NKB revision; *note that not all postulates are meant to hold for all operators—we will discuss this further below.*

Inclusion: $K'(v) \subseteq K(v) \cup lit(P_v)$. **Success:** $lit(P_v) \subseteq K'(v)$.

Weak Success: $lit(P_v) \subseteq K'(v)$ when $lit(P_v)$ is consistent. **Consistency:** $K'(v) \nvdash \perp$.

Vacuity 1: If $l \notin K(v), lit(p) = l$ implies $dec(p) = r$, for all $p \in P_v$, then $l \notin K'(v)$.

Vacuity 2: If $l \in K(v), lit(p) = \neg l$ implies $dec(p) = r$, for all $p \in P_v$, then $l \in K'(v)$.

Weak Vacuity 1: If $l \notin K(v)$ and $lit(p) \neq l$ for all $p \in P_v$, then $l \notin K'(v)$.

Weak Vacuity 2: If $l \in K(v)$ and $lit(p) \neq \neg l$ for all $p \in P_v$, then $l \in K'(v)$.

Strong Congruence: Let P_v^* be a set of news item and $\text{NKB}^* = \circledast(\text{NKB}, v, P_v^*) = (V^*, E^*, l_{vert}^*, l_{edge}^*, K^*)$, if $\chi(P_v) = \chi(P_v^*)$ then $K'(v) = K^*(v)$.

Weak Congruence: Let $P^+ = \{p \in \chi(P) \mid dec(p) = a\}$ be a set of all positive/added news items of $\chi(P)$, $P^- = \{p \in P \mid dec(p) = r\}$ be a set of all negative/removed news items of P and let P_v^* be a set of news item and $\text{NKB}^* = \circledast(\text{NKB}, v, P_v^*) = (V^*, E^*, l_{vert}^*, l_{edge}^*, K^*)$. If $P_v^+ = P_v^{*+}$, and $P_v^- \subseteq P_v^{*-}$ then $K^*(v) \subseteq K'(v)$.

Majority: Let $Pro = \{p \in \mathcal{P} \mid lit(p) = l, dec(p) \neq r\} \subseteq P_v$ and $Con = \{p \in P_v \mid lit(p) = \neg l, dec(p) \neq r\} \subseteq P_v$. Given $l \in \mathcal{L}$ s.t. $l \notin K(v)$ and $\neg l \notin K(v)$, if $|Pro| > |Con|$ then $\neg l \notin K'(v)$, and if $|Pro| < |Con|$ then $l \notin K'(v)$.

Local Effect: $\forall w \in V$ s.t. $w \neq v, K'(w) = K(w)$.

Structural Preservation: $V = V', E = E', l_{vert} = l'_{vert}$, and $l_{edge} = l'_{edge}$.

Weighted Majority: For NKBs that have one label per edge, given $l \notin K(v)$ and $\neg l \notin K(v)$, let $sumWPos$ and $sumWNeg$ be two values calculated as follows:
$sumWPos = \sum_{e \in I} weight(e)$; where $I = \{e \in E \mid e = (v, src(p)), lit(p) = l, p \in \mathcal{P}\}$.
$sumWNeg = \sum_{e \in J} weight(e)$; where $J = \{e \in E \mid e = (v, src(p)), lit(p) = \neg l, p \in \mathcal{P}\}$.
If $sumWPos > sumWNeg$ then $\neg l \notin K'(v)$; if $sumWPos < sumWNeg$ then $l \notin K'(v)$.

The following definition characterizes different types of local revisions.

Definition 6. *A local NKB revision operator is* basic *if it satisfies Structural Preservation, Local Effect, Consistency, Uniformity, and Inclusion. Let $\circledast$ be a basic local NKB revision operator:*
$\circledast$ *is* restrained *if it satisfies Strong Congruence, Vacuity 1, and Vacuity 2.*
$\circledast$ *is* weakly restrained *if it satisfies W. Congruence, and W. Vacuity 1 and 2.*
$\circledast$ *is* social *if it satisfies Weak Success, and either Majority or Weighted Majority.*

Towards Constructions based on User Types. We now briefly discuss how the operator types presented above could be used, along with characterizations of different user types, to construct concrete operators. Though the actual constructions are not discussed here, we take an important first step in the next section by showing how different user types can be detected in real-world datasets.

Credulous users adopt all new knowledge that they see in their feeds, even if it contradicts their previously-held beliefs.

Incredulous users are reluctant to incorporate new knowledge appearing in their feeds, regardless of whether it has any relation with their previously-held beliefs.

Users with herd behavior accept new information as long as there are enough other users adopting it.

Blind followers accept new knowledge being shared by their contacts, as long as such contacts have a close enough relation to them.

User	Avg. Tweet Reaction
u_1	3% (61% / 36%) 2% (73% / 25%) 38% (8% / 54%)
u_2	0% (39% / 69%) 0% (53% / 47%) 20% (0% / 80%)
u_3	2% (54% / 44%) 0% (67% / 33%) 34% (4% / 62%)
u_4	23% (13% / 64%) 1% (60% / 39%) 7% (25% / 68%)

User	Avg. Tweet Reaction
u_5	0% (0% / 100%) 0% (0% / 100%) 0% (0% / 100%)
u_6	9% (2% / 89%) 0% (31% / 69%) 1% (9% / 90%)
u_7	24% (2% / 74%) 1% (55% / 44%) 1% (23% / 76%)
u_8	1% (43% / 56%) 22% (12% / 66%) 2% (23% / 75%)

User	Avg. Tweet Reaction
u_9	67% (0% / 33%) 0% (92% / 8%) 0% (62% / 38%)
u_{10}	1% (51% / 48%) 1% (64% / 35%) 32% (2% / 66%)
u_{11}	11% (0% / 89%) 0% (33% / 67%) 0% (10% / 90%)
u_{12}	9% (2% / 89%) 0% (31% / 69%) 1% (9% / 90%)

Fig. 1. Summary of user reactions to news items. The first line contains the % of positive; in parentheses we show the % of changes in sentiment and the % of times the hashtag was not used again; the other two lines are analogous for negative and neutral.

Cautious users do not adopt new knowledge unless a selected group of their connections do.

Self-confident users assign more value to their previously-held beliefs, making it difficult for them to incorporate new information that contradicts them.

Combining operator and user types, we can in theory arrive at 18 different kinds of operators. In future work, we will formally characterize their construction.

3 Experimental Evaluation

The dataset is a subset of the Twitter network, with 184,654 users connected by 66,827,454 follow relationships, and 18,292,721 tweets. Hashtags are present in 5,107,986 tweets, and there are 136,809 distinct hashtags. Our objective is to show that by analyzing information flow in a social network, one can build a *map of users* based on how they behave—the goal is to inform the construction of local NKB revision operators as discussed above. We used hashtags as proxies for news items; to obtain the sentiment associated with the use of each hashtag, we used PHPInsight. We identified the 100 most prevalent hashtags and selected 12 users. For each one, we analyzed the sentiment of the tweet via which each hashtag reached their feed (either positive, negative, or neutral), and then analyzed how they reacted—either tweet using the hashtag (again, with positive, negative, or neutral sentiment) or to not use that hashtag in any of their subsequent tweets.

Results are shown in Figure 1; we averaged the sentiment distribution of the tweets seen by each user, yielding 29% positive, 19% negative, and 52% neutral. We summarize users' reactions as an average over all hashtags seen. Users u_1–u_3, and u_{10} can be seen as *self-confident*, as they almost never reuse hashtags with the same sentiment when receiving them with positive or negative. Users u_4, u_7, u_9 behave similarly, but more for positive tweets, either reusing positive tweets or changing the sentiment of negative ones; u_8 is a dual case. User u_{10} tends to change the sentiment of what they receive when its positive or negative. This analysis also detects "dead end" users such as u_5, u_6, u_{11}, u_{12}.

Acknowledgments. We thank support from CONICET, Universidad Nacional del Sur, ONR (grant N00014-15-1-2742), and V.S. Subrahmanian for the Twitter dataset.

References

1. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. Tech. rep., National Bureau of Economic Research (2017)
2. Gallo, F.R., Abad Santos, N., Simari, G.I., Falappa, M.A.: A desiderata for modeling and reasoning with social knowledge. In: Proc. of CACIC 2015 (2015)
3. Gallo, F.R., Abad Santos, N., Simari, G.I., Martinez, M.V., Falappa, M.A.: Belief dynamics in complex social networks. In: Proc. of ASAI 2016–JAIIO 45 (2016)
4. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: Proc. of KDD '03. pp. 137–146. ACM (2003)
5. 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)
6. Shakarian, P., Simari, G.I., Callahan, D.: Reasoning about complex networks: A logic programming approach. TPLP 13(4-5-Online-Supplement) (2013)
7. Shao, C., Ciampaglia, G.L., Flammini, A., Menczer, F.: Hoaxy: A platform for tracking online misinformation. In: Proc. WWW '16 Comp. pp. 745–750 (2016)