Logic of Visibility in Social Networks

Rustam Galimullin, Mina Young Pedersen*, and Marija Slavkovik

Dept. of Information Science and Media Studies, University of Bergen, Norway
{rustam.galimullin, mina.pedersen, marija.slavkovik}@uib.no

Abstract. Social media is not a neutral channel for news consumption. How visible information posted online is, depends on many factors such as the network structure, the emotional volatility of the content and the design of the social media platform. In this paper, we use formal methods to study the visibility of agents and information in a social network. We introduce a modal logic to reason about a social network of agents that can follow each other, post and share information. We show that by imposing some simple rules on the system, a potentially malicious agent can take advantage of the network construction to post an unpopular opinion that may reach many agents. The network is presented both in a static and dynamic form. We prove completeness, expressivity and model checking problem complexity results for the corresponding logical systems.

Keywords: Social networks · Visibility · Modal logic · Reachability.

1 Introduction

An overwhelming majority of people across the globe consume news on social media [1]. Social media, however, is not a neutral channel. How visible information posted online is, and how many users in a social network it can reach, depends on many factors. These include the network structure [16], the emotional volatility of the content [8], past exposure to similar information [21] and the design and recommendation algorithms of the particular social media platform [20].

This paper contributes to the study of social networks and the measurable impact social media platforms have on their users. Social networks have been studied using numerous methods in numerous disciplines. Formal logic methods for representing and reasoning about social networks have been used to analyze opinion diffusion and social influence [9–11, 19], social bots [22], group polarization [24, 25], gatekeepers [5], echo chambers [23] and informational cascades [3], among other phenomena. Our work is positioned within this literature.

We are here concerned with the problem of applying **formal methods to study the visibility of agents and information in a social network**. In addition to having structural properties, number of agents and how they are connected, a social network also has other properties connected to the visibility of an agent, such as: which interests and opinions the agents have, what they are

* Corresponding author.

communicating and how the network changes through time. It is our position that logic based methods are needed to complement empirical methods to reach a full understanding of these properties of social networks.

The notion of visibility is rooted in the idea of being seen. One of our main motivations is to present an analysis of visibility that captures a complex view of what it means to be visible in a social network, one that extends merely counting the followers of an agent. To do this, we introduce a modal logic for representing agents, their opinions and interactions in a social network.

Our social network consists of a set of agents and two sets of relations between them: one represents followers, the other represents posts that pass through the network. We turn to Figure 1 for an intuitive explanation of the network.

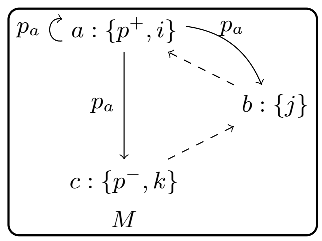

Fig. 1. Model M with the followership relation depicted by dashed arrows.

The network M consists of three agents a, b and c. Dashed arrows represent a followership relation: c follows b and b follows a. The situation concerns a post on a particular topic, called p. Agent a is in favour of, *pro*, p, denoted p^+, whereas c is *contra* p, denoted p^-. Agent b has no opinion about p. Furthermore, agent a has posted on p, represented by a reflexive loop denoted p_a and agents b and c have seen the post, denoted by p_a-arrows from a.

The intuition behind our models is to observe a situation of posting and sharing a post, after it has happened. Posting, sharing, following and unfollowing adhere to some simple rules of the system:

1. When an agent posts, all her followers can see the post.
2. If an agent sees a post on a topic she likes, she will reshare the post and follow the original poster.
3. If an agent sees a post on a topic she dislikes, she does not reshare it and unfollows the agent from whom she has seen the post.
4. If an agent sees a post on a topic she is indifferent to, she does not do anything.

These rules are an oversimplification of a real-life network, but we believe they capture some key notions of a social network which we can use to analyze situations that may occur in an actual network setting. Knowing the rules of the system, we can return to M in Figure 1 and observe that c likely has unfollowed a after a posted on p.

We first present a logic that specifies a static network, as seen in Figure 1. The purpose of this logic is not to define what visibility is, but to allow us to discuss different qualitative and quantitative measures of visibility and formalise some of them in the logic. Next, we extend the framework into a dynamic setting where we step-wise observe what happens when information is posted in the network. We show that according to the rules of the system, the interests of the agents' followers matter a lot to what information is shared and seen. We also show that a malicious agent could take advantage of the network construction to post an unpopular opinion that will reach many agents. We believe these observations can be useful to understand how agents in a network contribute to spreading controversial information such as misinformation.

We are also interested in the mathematical properties of the system itself. We give formulas corresponding to the rules of the system, and claim that the logic is complete with respect to the models with these rules. The model checking problem for the static logic is proved to be in P. In the dynamic extension, we show that the language with the dynamic modality is strictly more expressive than the static language without. We also prove that the model checking problem for the dynamic logic is $PSPACE$-complete.

The contribution of the paper is the following:

- We introduce a novel logic to analyze posting and sharing information in a social network, and prove mathematical results about this formal system.
- We propose quantitative and qualitative measures of visibility and reachability, and formalize some of the properties as logical formulas.
- We use our formal system to reason about mechanisms that might occur in real-life online social networks, specifically we formalize how a potentially malicious agent could take advantage of the network construction to post a controversial opinion that will reach many agents.

The paper is structured as follows. In Section 2, we give an overview of the related work in social network analysis on reachability and visibility. In Section 3, we present static visibility logic (SVL). We specify some mathematical properties of the logic, propose measures of visibility and give some corresponding logical formulas. In Section 4, we extend SVL with a dynamic operator and name it visibility logic (VL). We give a motivating example where we show that one can exploit the network structure to expose more agents to a controversial opinion. In Section 5, we summarize our paper and outline directions for future work.

2 Related Work: Visibility and Reachability

Visibility in social networks is yet to be explicitly explored from a formal logical perspective. The concept has however been researched in the social network analysis literature. We present a selected collection of this work to learn how this related field has attempted to measure visibility. There seem to be no consensus in the literature of what it means to be visible in a social network, which motivates the usefulness of a further study on this topic.

Closely related to visibility in social networks is the notion of reachability. What exactly reachability is, or how closely related it is to visibility, is not agreed upon. This is illustrated by the different measures seen in this section. Visibility and reachability are presented as properties of both networks, agents and posts, when relevant we specify which in the following.

[13], known to be part of the canonical literature in social network analysis, describe the reachability properties of a network in terms of identifying which agents are reachable from which others through connected paths of edges.

[28] distinguishes reachability and visibility in an online social network, where the first measure is dependent on the second. The network is represented as an undirected graph where nodes represent agents and the relation between them represent one of three non-overlapping relations: trusted friends, acquaintances or distrusted agents. Agents can post information with four different visibility settings: trusted friends, trusted friends and acquaintances, all friends and public. The visibility of an agent is therefore measured with respect to what relation the viewers of the post has to the agent that posts.

The reachability factor of a post is defined in terms of a function: $d(v_1, v_2) = \frac{|e(v_1,v_2)|}{\sqrt{|v_1|} \times \sqrt{|v_2|}}$. In this function, v_1 is the set of agents in the network that have seen the post and v_2 is the set of agents that have not seen the post. $e(v_1, v_2)$ is the set of relations between agents across v_1 and v_2 specified with respect to the relations in the network graph. The reachability factor is dependent on the visibility settings of the agent who posts; the set v_1 increases and v_2 decreases when the visibility settings include a higher number of agents.

[29] presents a temporal characterisation of reachability. In this work, the reachability is measured between two given nodes in a time interval in the network. The network is presented as a series of undirected graphs which represent how a network changes through time. The nodes in the network can be regarded as agents and the relation between them as information channels. Node j is reachable in the time interval $[t_{min}, t_{max}]$ from node i if a message can be delivered through the information channel in that time interval.

[27] defines two types of visibility of an agent in an online social network: topological and behavioral visibility. Although it is mentioned that this could be a generic social network, the examples refer to the microblogging network Twitter, which is represented as a directed graph of agents who can post and follow each other. In the model presented in [27], a tweet embodies at least one topic from a set of interests S. Each agent in the network also has some specified interests from S. Topological visibility of an agent is calculated based on the number of followers of the agent and the clustering coefficient of the network. The clustering coefficient is usually defined in the literature in terms of directed graphs, and is meant to give a view of the network structure. The higher the number, the more highly connected the network is. It is not specified which definition of clustering coefficient is used in [27]. The behavioral visibility of an agent is defined as the average of the visibility of all the tweets that are shared by the user in a time interval Δt. The visibility of a tweet represents the number

of users influenced by the tweet, and is proportional to the number of followers whose interests matches the topics of the tweet.

3 Reasoning About Visibility in a Static Setting

In this section we present some of the main concepts underlying our intuitions about visibility in social networks.

3.1 Language and Semantics of SVL

Let $\mathsf{Nom} = \{i, j, k, ...\}$ be a countable set of nominals, and $\mathsf{Top} = \{p, q, r, ...\}$ be a countable set of topics, such that $\mathsf{Nom} \cap \mathsf{Top} = \emptyset$.

Definition 1 (Syntax). *We define the well-formed formulas of* the language of the static fragment of visibility logic $\mathcal{SVL}$ *to be generated by the following grammar:*

$$\varphi ::= p^+ \mid p^- \mid i \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Diamond_{i:p}\varphi \mid \Diamond^{-1}_{i:p}\varphi \mid \blacklozenge\varphi \mid \blacklozenge^{-1}\varphi \mid @_i\varphi$$

where $p \in \mathsf{Top}$ *and* $i \in \mathsf{Nom}$. *We define propositional connectives like* $\vee, \rightarrow$ *and the formulas* $\top, \bot$ *as usual and the duals as standard* $\Box := \neg\Diamond\neg$, $\Box^{-1} := \neg\Diamond^{-1}\neg$, $\blacksquare := \neg\blacklozenge\neg$, *and* $\blacksquare^{-1} := \neg\blacklozenge^{-1}\neg$.

In our language, similar to other approaches to logic-based analysis of social networks (see, e.g., [10]), we distinguish three possible dispositions of an agent to a topic $p \in \mathsf{Top}$. The agent may be *pro* p, which we express with p^+, *contra* p, expressed by p^-, or *indifferent to* p, if the agent is neither *pro* nor *contra* p.

Constructs $\blacklozenge\varphi$ and $\blacklozenge^{-1}\varphi$ express that 'the current agent follows an agent satisfying φ' and 'the current agent is followed by an agent that satisfies φ' respectively. Formulas $\Diamond_{i:p}\varphi$ and $\Diamond^{-1}_{i:p}\varphi$ mean that 'there is an agent satisfying φ who sees the (re)post of agent i on topic p' and 'there is an agent satisfying φ whose (re)post on topic p (originally posted by agent i) is seen by the current one'.

Formulas of $\mathcal{SVL}$ are defined on relational visibility models.

Definition 2. *A* visibility model *(or a* model*)* M *is a tuple* $(A, F, +, -, V, R)$, *where*

- A *is a non-empty set of agents;*
- $F : A \rightarrow 2^A$ *is an irreflexive followership relation,*
- $+ : A \rightarrow 2^{\mathsf{Top}}$ *assigns to each agent a set of topics she is pro,*
- $- : A \rightarrow 2^{\mathsf{Top}}$ *assigns to each agent a set of topics she is contra such that for all agents* a *it holds that* $+(a) \cap -(a) = \emptyset$,
- $V : \mathsf{Nom} \rightarrow 2^A$ *is a valuation such that for all* $i \in \mathsf{Nom}$*:* $|V(i)| = 1$,
- $R : \mathsf{Top} \times A \rightarrow 2^{A \times A}$ *is a visibility relation for each topic and each agent satisfying the following conditions, where* $p \in \mathsf{Top}$ *and* $c \in A$*:*
  1. *If* $(a, b) \in R(p, c)$*, then* $(a, a) \in R(p, c)$.

2. *If* $(a,a) \in R(p,c)$*, then* $(a,b) \in R(p,c)$ *for all* b *such that* $b \in F(a)$*.*
3. *If* $(a,b) \in R(p,c)$*,* $p \in +(b)$*, and* $b \neq c$*, then* $(b,b) \in R(p,c)$ *and* $b \in F(c)$*.*
4. *If* $(a,b) \in R(p,c)$*,* $p \in -(b)$*, and* $a \neq b$*, then* $(b,b) \notin R(p,c)$ *and* $b \notin F(a)$*.*
5. *If* $(a,b) \in R(p,c)$*,* $p \notin +(b)$*,* $p \notin -(b)$*, and* $a \neq b$*, then* $(b,b) \notin R(p,c)$*.*

A pointed visibility model M_a *is a model* M *with a distinguished point* $a \in A$ *where evaluation takes place. If necessary, we refer to the elements of the tuple as* A_M*,*F_M*,*$+_M$*,*$-_M$*,*V_M*, and* R_M*. A visibility model such that for all* $a \in A$ *there is some* $i \in \mathsf{Nom}$ *such that* $V(i) = \{a\}$ *is called* named*. All models we will be dealing with in the paper are named. A visibility model with no restrictions on* F *or* R *is called* minimal*.*

In the definition above, the first condition on R states that if agent b sees a post, which was originally posted by agent c on topic p, from agent a, then a herself can see the post. The second condition ensures that if an agent posts a post, all her followers can see the post. Condition number three specifies that if an agent sees a post on the topic she likes, she will reshare the post and follow the original poster. The fourth condition says that if an agent sees a post on the topic she dislikes, she does not reshare it and unfollows the agent from whom she has seen the post. Finally, the last condition stipulates that if an agent sees a post on a topic she is indifferent to, she does not reshare the post.

Note that our definition of R does not preclude situations where agents may have seen a post on the topic they dislike from an agent they do not follow. How such situation may come about will be the focus of the next section.

Definition 3 (Semantics). *Let* $M = (A, F, +, -, V, R)$ *be a model,* $a, b, c \in A$*,* $p \in \mathsf{Top}$*,* $i \in \mathsf{Nom}$*, and* $\varphi, \psi \in \mathcal{SVL}$*.* The semantics of SVL *is recursively defined as follows:*

$$
\begin{array}{lll}
M_a \models p^+ & \textit{iff } p \in +(a) \\
M_a \models p^- & \textit{iff } p \in -(a) \\
M_a \models i & \textit{iff } a \in V(i) \\
M_a \models \neg\varphi & \textit{iff } M_a \not\models \varphi \\
M_a \models \varphi \wedge \psi & \textit{iff } M_a \models \varphi \textit{ and } M_a \models \psi \\
M_a \models \Diamond_{i:p}\varphi & \textit{iff } \exists b, c \in A : (a,b) \in R(p,c) \textit{ and } V(i) = \{c\} \textit{ and } M_b \models \varphi \\
M_a \models \Diamond_{i:p}^{-1}\varphi & \textit{iff } \exists b, c \in A : (b,a) \in R(p,c) \textit{ and } V(i) = \{c\} \textit{ and } M_b \models \varphi \\
M_a \models \blacklozenge\varphi & \textit{iff } \exists b \in A : a \in F(b) \textit{ and } M_b \models \varphi \\
M_a \models \blacklozenge^{-1}\varphi & \textit{iff } \exists b \in A : b \in F(a) \textit{ and } M_b \models \varphi \\
M_a \models @_i\varphi & \textit{iff } M_b \models \varphi \textit{ and } \{b\} = V(i)
\end{array}
$$

Observe that if $M_a \not\models p^+$ then we have that either $p \in -(a)$ or not. This corresponds to the intuition that agent a is not *pro* p if she actively dislikes the topic (she is *contra* p), or she is indifferent to it. Similarly, for $M_a \not\models p^-$.

Recall the example from Figure 1. In the figure, we have that $M_a \models p^+$, $M_c \models p^-$, and $M_b \models \neg p^+ \wedge \neg p^-$, meaning that agent a is *pro* topic p, agent b is indifferent towards the topic, and c is *contra* p. Moreover, we have, for example, that $M_c \models \Diamond_{i:p}^{-1}\top \wedge \blacksquare\neg p^+$, meaning that agent c has seen a post by the agent with name i on topic p, and that all agents that c follows are not *pro* p.

Definition 4. *A formula* φ *is called* valid *if for all models* M_a*, we have that* $M_a \models \varphi$*. Formulas* φ *and* ψ *are* equivalent*, if for all models* M_a*, it holds that* $M_a \models \varphi$ *if and only if* $M_a \models \psi$*.*

Definition 5. *Let* $\mathcal{L}_1$ *and* $\mathcal{L}_2$ *be two languages. We say that* $\mathcal{L}_2$ is more expressive than $\mathcal{L}_1$ *if for each* $\varphi \in \mathcal{L}_1$ *there is an equivalent* $\psi \in \mathcal{L}_2$*, and there is a* $\chi \in \mathcal{L}_2$ *for which there is no equivalent* $\tau \in \mathcal{L}_1$*.*

The following notion of bisimulation is based on hybrid bisimulation [2] and on bisimulation for logics with 'backwards-looking' modalities (see, e.g., [18]).

Definition 6. *Let* $M = (A_M, F_M, +_M, -_M, V_M, R_M)$ *and* $N = (A_N, F_N, +_N, -_N, V_N, R_N)$ *be visibility models, and* $Q \subseteq \mathsf{Nom}$*. We say that* M *and* N *are* Q-bisimilar *(denoted* $M \leftrightarrows_Q N$*) if there is a non-empty relation* $B \subseteq A_M \times A_N$*, called* Q-bisimulation *such that the following conditions are satisfied:*

Atoms$^+$ *If* $B(a,b)$*, then for all* $p \in \mathsf{Top}$*:* $p \in +_M(a)$ *iff* $p \in +_N(b)$*,*
Atoms$^-$ *If* $B(a,b)$*, then for all* $p \in \mathsf{Top}$*:* $p \in -_M(a)$ *iff* $p \in -_N(b)$*,*
Nominals 1 *If* $B(a,b)$*, then for all* $i \in Q$*:* $a \in V_M(i)$ *iff* $b \in V_N(i)$*,*
Nominals 2 *For all* $i \in Q$*, if* $V_M(i) = \{a\}$ *and* $V_N(i) = \{b\}$*, then* $B(a,b)$*,*
Forth $\Diamond$ *If* $B(a,b)$ *and* $(a,a') \in R_M(p,c)$*, then there is a* $b' \in A_N$ *such that* $(b,b') \in R_N(p,c)$ *and* $B(a',b')$*,*
Back $\Diamond$ *If* $B(a,b)$ *and* $(b,b') \in R_N(p,c)$*, then there is an* $a' \in A_M$ *such that* $(a,a') \in R_M(p,c)$ *and* $B(a',b')$*,*
Forth $\Diamond^{-1}$ *If* $B(a,b)$ *and* $(a',a) \in R_M(p,c)$*, then there is a* $b' \in A_N$ *such that* $(b,b') \in R_N(p,c)$ *and* $B(a',b')$*,*
Back $\Diamond^{-1}$ *If* $B(a,b)$ *and* $(b',b) \in R_N(p,c)$*, then there is an* $a' \in A_M$ *such that* $(a',a) \in R_M(p,c)$ *and* $B(a',b')$*,*
Forth $\blacklozenge$, Back $\blacklozenge$, Forth $\blacklozenge^{-1}$, Back $\blacklozenge^{-1}$ *Similar to the cases of* $\Diamond$ *and* $\Diamond^{-1}$ *substituting* R *for* F *taking care of arguments.*

We say that M_a *and* N_b *are* Q*-bisimilar and denote this by* $M_a \leftrightarrows_Q N_b$ *if there is a bisimulation linking agents* a *and* b*.*

The following theorem is a standard result in modal logic [2].

Theorem 1. *Let* M_a *and* N_b *be two models. If* $M_a \leftrightarrows_Q N_b$*, then for all* $\varphi \in \mathcal{SVL}$ *such that* φ *includes only nominals from* Q*,* $M_a \models \varphi$ *if and only if* $N_b \models \varphi$*.*

3.2 Soundness, Completeness and Model Checking of SVL

To argue that SVL is complete, we first present an axiomatisation for *the minimal version* of SVL. Recall that a visibility model is called minimal, if there are no restrictions on relations F or R.

Definition 7. *The proof system of* the minimal version of SVL, **MSVL***, comprises axioms and rules of inference of basic hybrid logic* $\mathcal{H}(@)$ *and additional axioms and rules of inference dealing with the converse relation.*

$(A0)$	*Propositional tautologies*	$(A13)$	$@_i\Box_{i:p}\Diamond^{-1}_{i:p}i$
$(A1)$	$\Box_{i:p}(\varphi \to \psi) \to (\Box_{i:p}\varphi \to \Box_{i:p}\psi)$	$(A14)$	$@_i\Box^{-1}_{i:p}\Diamond_{i:p}i$
$(A2)$	$\Box^{-1}_{i:p}(\varphi \to \psi) \to (\Box^{-1}_{i:p}\varphi \to \Box^{-1}_{i:p}\psi)$	$(A15)$	$@_i\blacksquare\blacklozenge^{-1}i$
$(A3)$	$\blacksquare(\varphi \to \psi) \to (\blacksquare\varphi \to \blacksquare\psi)$	$(A16)$	$@_i\blacksquare^{-1}\blacklozenge i$
$(A4)$	$\blacksquare^{-1}(\varphi \to \psi) \to (\blacksquare^{-1}\varphi \to \blacksquare^{-1}\psi)$	$(A17)$	$p^+ \land p^- \to \bot$
$(A5)$	$@_i(\varphi \to \psi) \to (@_i\varphi \to @_i\psi)$	$(R0)$	*From* $\varphi \to \psi$ *and* φ, *infer* ψ
$(A6)$	$@_i\varphi \leftrightarrow \neg @_i\neg\varphi$	$(R1)$	*From* φ, *infer* $\Box_{i:p}\varphi$
$(A7)$	$@_i i$	$(R2)$	*From* φ, *infer* $\Box^{-1}_{i:p}\varphi$
$(A8)$	$@_i@_j\varphi \leftrightarrow @_j\varphi$	$(R3)$	*From* φ, *infer* $\blacksquare\varphi$
$(A9)$	$\Diamond_{i:p}@_i\varphi \to @_i\varphi$	$(R4)$	*From* φ, *infer* $\blacksquare^{-1}\varphi$
$(A10)$	$\Diamond^{-1}_{i:p}@_i\varphi \to @_i\varphi$	$(R5)$	*From* φ, *infer* $@_i\varphi$
$(A11)$	$\blacklozenge @_i\varphi \to @_i\varphi$	$(R6)$	*From* φ, *infer* φ^σ
$(A12)$	$\blacklozenge^{-1}@_i\varphi \to @_i\varphi$	$(R7)$	*From* $@_i\varphi$ *with* $i \notin \varphi$, *infer* φ

In the proof system, σ *is a* substitution *that uniformly replaces nominals by nominals, and topics with formulas.*

The axiomatisation of the minimal version of SVL is essentially an axiomatisation of hybrid tense logic [6] with additional axioms and rules of inference for followership. In particular, axioms $(A13)-(A16)$ capture the interaction between modalities and their converses. Axiom $(A17)$ specifies that agents cannot be both *pro* and *contra* the same topic. The completeness of the axiomatisation can be shown via a standard canonical model construction for hybrid logics (see, e.g., [7, Section 7.3])

Theorem 2. MSVL *is sound and complete with respect to the class of minimal visibility models.*

To get the axiomatisation of full SVL, we extend the axiomatisation from Definition 7 with an axiom for irreflexivity of followership $@_i\neg\blacklozenge i$ and five axioms for conditions on the visibility relation. We call the resulting proof system **SVL**.

Proposition 1. *The following formulas capture the conditions on the visibility relation (from 1 to 5 in Definition 2):*

1. $@_i(\Diamond_{j:p}\top \to \Diamond_{j:p}i)$
2. $@_i(\Diamond_{j:p}i \to \blacksquare^{-1}\Diamond^{-1}_{j:p}i)$
3. $@_i((p^+ \land \Diamond^{-1}_{j:p}\top \land \neg j) \to (\Diamond_{j:p}i \land \blacklozenge j))$
4. $@_i((p^- \land \Diamond^{-1}_{j:p}(\neg i \land k)) \to (\neg\Diamond_{j:p}\top \land \neg\blacklozenge k))$
5. $@_i((\neg p^- \land \neg p^+ \land \Diamond^{-1}_{j:p}\neg i) \to \neg\Diamond_{j:p}\top)$

We believe that completeness of **SVL** can be shown by the canonical model construction [7, Section 7.3], where our model will additionally have two different valuation functions, one for $+$ and one for $-$, and two types of relations for F and R. That the model satisfies conditions 1–5 from Definition 2 is guaranteed by axioms from Proposition 1 and the irreflexivity of followership axiom. We omit full details due to the lack of space.

Theorem 3. SVL *is complete with respect to visibility models.*

We also mention here that the complexity of the model checking problem of SVL is in P. This result follows trivially from the fact that model checking hybrid tense logic with universal modality is in P [15].

Theorem 4. *Model checking SVL is in P.*

3.3 Expressing Visibility

We look at measures of visibility and reachability, and formalize some of them as formulas in SVL. Let $M = (A, F, +, -, V, R)$ be a model. Some quantitative amounts related to visibility that we can count in finite models are:

- How many followers the agent called i has: $|\{a \in A \mid M_a \models \blacklozenge i\}|$
- How many agents have seen the agent called i's post on p: $|\{a \in A \mid M_a \models \Diamond_{i:p}^{-1}\top\}|$
- How many agents that are pro p have seen the agent called i's post on p: $|\{a \in A \mid M_a \models p^+ \wedge \Diamond_{i:p}^{-1}\top\}|$

We also present some formulas corresponding to qualitative properties of agents in the network. The following formulas are forced at an agent iff the property holds of that agent:

- The current agent i is the original poster of a post on p: $i \wedge \Diamond_{i:p}\top$
- The current agent has seen i's post on p: $\Diamond_{i:p}^{-1}\top$
- All the followers of the current agent i have shared i's post on p: $i \wedge \blacksquare^{-1}\Diamond_{i:p}\top$
- The current agent i shared a post to a follower j, but j also saw the post from another source: $i \wedge \blacklozenge^{-1}(j \wedge \Diamond_{i:p}^{-1}i \wedge \Diamond_{i:p}^{-1}(\neg i \wedge \neg j))$
- The current agent i has gained a follower who is *pro* p, after i posted on p: $i \wedge \Diamond_{i:p}\top \wedge \blacklozenge^{-1}(p^+ \wedge \Diamond_{i:p}^{-1}i)$
- The current agent i has reached the agent j with i's post on p in no more than 3 steps: $i \wedge \Diamond_{i:p}\Diamond_{i:p}\Diamond_{i:p}j$

4 Visibility Logic

To reason about the effects of agents posting on various topics, we introduce a dynamic extension of SVL that we call *visibility logic* (VL). Compared to SVL, VL is enriched with dynamic operators $[\pi]\varphi$, where π is an action of the current agent making a post. While defining VL, we follow *dynamic epistemic logics* (DELs) [12], and in particular *action model logic* [4, 12]. We begin with a motivating example.

4.1 Example: Taking the Advantage to be Seen by Many

In some networks, the best tactic for exposing more agents to a controversial opinion is to first post on a popular topic. Consider the follower-network M in Figure 2. For simplicity, we do not include nominals in the Figure. This network consists of 6 agents named in alphabetical order from a to f. Agent a has two followers b and c. Agent b has three followers d, e and f. We assume that agents d, e and f might have some followers that we do not have information about, noted in the figure with dots. Furthermore, agent a is positive in favor of vaccination (abbreviated v in the figure) which is a controversial topic amongst the agents: all agents have an opinion about vaccines and three of the agents a, c and f are *pro* vaccination, whereas b, d and e are *contra* vaccination. The topic of dogs (abbreviated d) on the other hand, is widely liked. All agents like dogs, except a who is indifferent: $d \notin +(a)$ and $d \notin -(a)$.

Imagine that agent a wants to post on vaccines, and want as many as possible of the other agents in the network to see the post. We show that the best tactic for agent a is to first post on dogs, even though a is indifferent about dogs, and then later post on vaccines. Consider first the scenario in Figure 2 where agent a posts v from the initial outset. An update happens in two steps. First, we add visibility arrows corresponding to posting and resharing. In the second step, we update the followership relation based on whether the agents who has seen a post are *pro* or *contra* the post. The resulting update is $M^{a:v}$ in Figure 2, where agents b and c has seen a's vaccine-post and only c remains as a's follower.

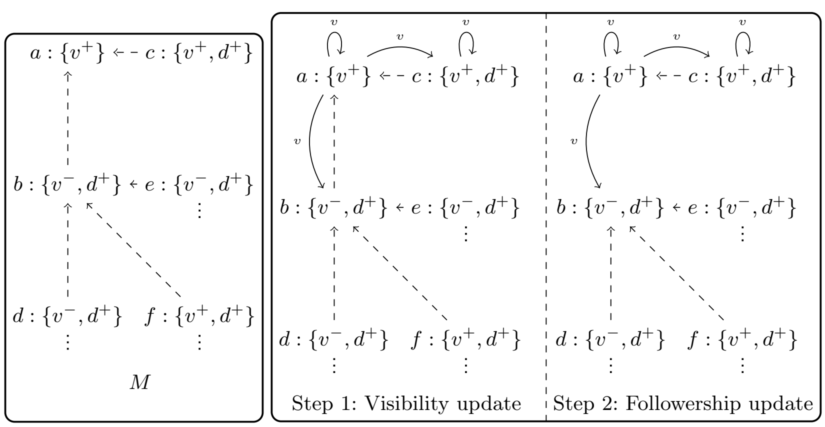

Fig. 2. A follower-network M (left) where vaccination is a controversial topic, and update $M^{a:v}$ (right) after agent a posts in favor of vaccines.

Then, consider instead the situation where agent a posts on dogs in the update $M^{a:d}$ before posting on vaccines in the update $M^{a:d,a:v}$, both in Figure 3. Note that to make the situation easier to read, we omit the followership arrows in the visibility update and the visibility arrows in the followership update in both figures. After $M^{a:d}$, all agents have seen agent a's post on dogs. Since they like dogs, all agents also follow a after the update.

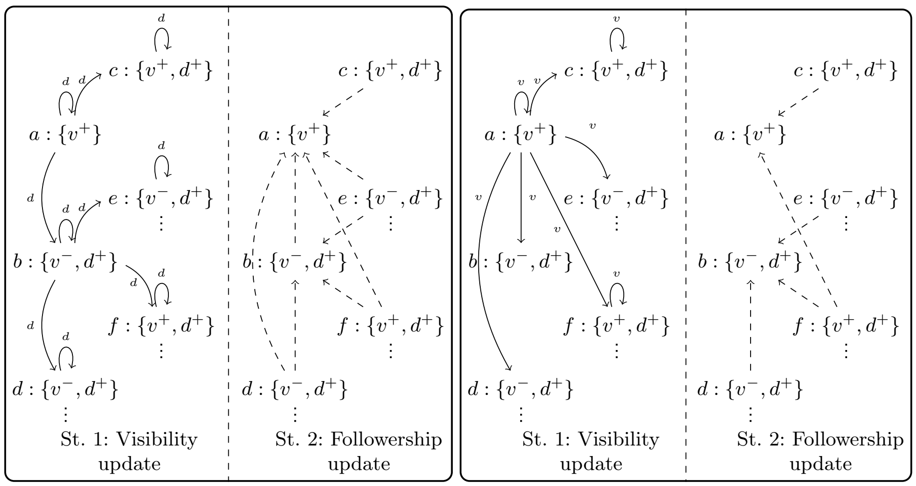

Fig. 3. Updates $M^{a:d}$ (left) and $M^{a:d,a:v}$ (right) after agent a posts on dogs before posting on vaccination.

In $M^{a:d,a:v}$, we see the results after agent a first posted on dogs, and then vaccines. All the agents have now seen a's vaccine-post. Most of them did not like it and unfollowed a, but only after they were exposed to the post. Interestingly, we also notice that agent f, who was not originally a follower of a in the initial network outset, now follows a and has shared the vaccine post to their followers.

There are two tactical reasons for agent a to post on dogs before their more controversial post on vaccines. Firstly, a larger portion of the agents now saw a's post on vaccines since they followed agent a after the dog-post. Secondly, a has been able to reach out and expand their network: agent f who is also pro vaccines, has shared the vaccination post to their, for us unknown, followers.

The reason behind a phenomenon such as this is directly connected to an underlying notion of trust between agents in the network. In our setting, agents follow other agents when the former is exposed to content that they like by the latter. In the example, we imagine agents likely followed a because they wanted to see more dog-friendly content. Agent a misuses the trust of their followers by pretending to be interested in dogs before posting on vaccination.

What becomes clear in this example, is that in our simplified setting of posting and sharing in a social network, the interests of an agents' followers matter a lot to what information is shared and seen. Secondly, the system is vulnerable to exploitation by a potentially malicious agent: there are opportunities to tactically post on popular topics to later expose more agents to a controversial opinion. To reason about dynamic situations such as these, we introduce VL.

4.2 Language, Semantics, and Logical Properties of VL

The language of VL is an extension of the language of SVL.

Definition 8 (Syntax). The language of visibility logic $\mathcal{VL}$ *is defined recursively by the following BNF:*

$$\varphi ::= p^+ \mid p^- \mid i \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Diamond_{i:p}\varphi \mid \Diamond^{-1}_{i:p}\varphi \mid \blacklozenge\varphi \mid \blacklozenge^{-1}\varphi \mid @_i\varphi \mid [\pi]\varphi$$
$$\pi ::= p \mid (\pi \cup \pi)$$

where $[\pi]\varphi$ *is read 'after the current agent executes action* π*,* φ *holds'.*

Union of actions $(\pi \cup \tau)$ was inherited by DELs from propositional dynamic logic [14], and in the context of visibility formulas, $[p \cup q]\varphi$ mean 'whichever topic the current agent posts on, p or q, φ will be true (in both cases)'.

Definition 9 (Semantics). *Let* $M = (A, F, +, -, V, R)$ *be a visibility model,* $a \in A$*, and* $p, q \in \mathsf{Top}$. The semantics of VL *is the same as in Definition 3 with the following additions:*

$$\begin{array}{ll} M_a \models [p]\varphi & \textit{iff } M_a^{a:p} \models \varphi \\ M_a \models [\pi \cup \tau]\varphi & \textit{iff } M_a \models [\pi]\varphi \textit{ and } M_a \models [\tau]\varphi \end{array}$$

where $M_a^{a:p}$ is defined in the following two steps. First, let $M^* = (A,F,+,-, V, R^*)$, where $R^*(p, a)$ is the least fixed point of function $f : 2^A \to 2^A$ defined as

$$\begin{aligned} f(X) =& X \cup \{(a, a)\} \cup \{(b, c) \mid (b, b) \in X \text{ and } c \in F(b)\} \cup \\ & \cup \{(c, c) \mid p \in +(c) \text{ and } \exists b : (b, c) \in X\}. \end{aligned}$$

Informally, intermediate model M^* differs from M only in R in such a way that R^* now contains the fact that a has posted on p, that her post has reached all her followers, and that all followers who are *pro* p reshare the post further to their followers. Secondly, we construct $M^{a:p}$ out of M^* by updating F:

1. $F^{a:p}(a) = F(a) \cup \{b\}$, if $a \neq b$, $p \in +(b)$, and $\exists c : (c, b) \in R^*(p, a)$,
2. $F^{a:p}(b) = F(b) \setminus \{c\}$, if $p \in -(b)$ and $(c, b) \in R^*(p, a)$.

Intuitively, agent b will follow the original poster a if she has seen the post, maybe not even from a, and if she is *pro* the topic. Agent c will stop following anyone from whom she has seen a post on a topic she dislikes.

To give a taste of VL, let us provide some properties that are valid or not valid on visibility models. All the validities can be shown by an application of the definition of the semantics.

Proposition 2. *Let* $p, q \in \mathsf{Top}$ *and* $\varphi \in \mathcal{VL}$*.*

1. $\neg[p]\varphi \leftrightarrow [p]\neg\varphi$ *is valid.*
2. $[\pi \cup \tau]\varphi \leftrightarrow [\pi]\varphi \wedge [\tau]\varphi$ *is valid.*
3. $\Diamond^{-1}_{i:p}\varphi \leftrightarrow [q]\Diamond^{-1}_{i:p}\varphi$ *is valid.*
4. $[p]\varphi \to [p][p]\varphi$ *is not valid.*
5. $[p][p]\varphi \to [p]\varphi$ *is not valid.*

The first formula states that the operator of posting on a topic is its own dual. The second property shows how to eliminate non-deterministic choice. The third item claims that once an agent has seen a post of agent with name i on topic p, no further post can revoke this. The fact that formulas four and five are not valid indicates that consecutive posting on the same topic yields different results. A counterexample showing this would include the current agent with name i posting on p and gaining new followers. Additional post on the same topic by the same agent will add new p-arrows to those new followers thus resulting in a different updated model that is not guaranteed to satisfy φ.

4.3 Expressivity and Model Checking

Now, we state that VL is more expressive than its static fragment SVL. This result is quite interesting, since many of DELs, for example public announcement logic [26], arrow update logic [17], and action model logic [12, Chapter 6], are equally expressive as the static logic they are built upon. Those expressivity results are usually obtained with the use of so-called reduction axioms that allow one to equivalently rewrite formulas of dynamic extensions to formulas of the static fragment. Thus, the fact that VL is more expressive than SVL also entails that no reduction axioms for VL are possible.

Theorem 5. *$\mathcal{SVL} < \mathcal{VL}$.*

Proof. Consider a $\mathcal{VL}$ formula $[p]\blacklozenge^{-1}\blacksquare^{-1}\bot$, and assume towards a contradiction that there is an equivalent formula ψ of $\mathcal{SVL}$. Since ψ has a finite size $n = |\psi|$, there is a set of nominals $Q = \{j_1, \ldots, j_{n+1}\}$ that are not present in ψ.

Consider models M and N in Figure 4. The models are chains of length $n+2$ that start with agent a and with each next agent following the previous one. The only difference between the models is that the last agent in the chain in model M is *pro* topic p, and the last agent in the chain in model N is neither *pro* nor *contra* topic p.

Now we will argue that $[p]\blacklozenge^{-1}\blacksquare^{-1}\bot$ distinguishes M_a and N_a. In particular, $M_a \models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$ and $N_a \not\models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$. Indeed, agent a posting on topic p results in the updated visibility model $M_a^{a:p}$ presented in the figure. In the updated model, it holds that $M_{b_{n+1}}^{a:p} \models \blacklozenge i$, i.e. that agent b_{n+1} follows agent a. Moreover, agent b_{n+1} does not have any followers, so $M_{b_{n+1}}^{a:p} \models \blacksquare^{-1}\bot$ is vacuously true. Hence, $M_a \models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$. To see that $N_a \not\models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$, it is enough to notice that agent b_{n+1} is not *pro* topic p, and thus they do not follow agent a in the updated model.

In order to show that $M_a \models \psi$ if and only if $N_a \models \psi$, we informally argue that M_a and N_a are $n-Q$-bisimilar. Let us first recall that ψ does not contain any of $j_1, \ldots, j_{n+1}$, and thus these nominals cannot be used to access the corresponding agents in the models. Without such an ability, the only way for ψ to spot a difference between M_a and N_a is via a sequence of $\blacklozenge^{-1}$ steps reaching agent b_{n+1}. Observe, however, that since the size of ψ is n, and that both M_a and N_a are of lengths $n+2$, there is not enough modal depth in ψ to reach the distinguishing state b_{n+1}. Hence, a contradiction. $\square$

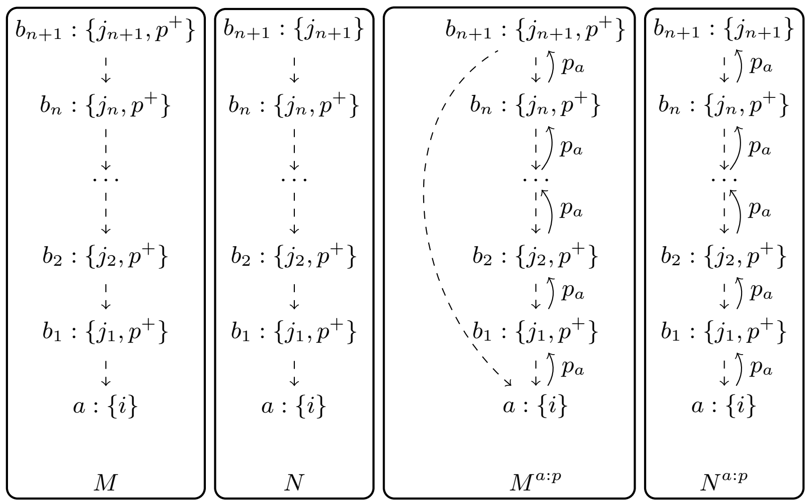

Fig. 4. Models M, N, $M^{a:p}$, and $N^{a:p}$. For models $M^{a:p}$ and $N^{a:p}$ reflexive p_a-arrows and followership arrows from b_k to a for $k \in \{1, ..., n\}$ are omitted for readability.

Not only is VL more expressive than SVL, its model checking problem is also more computationally demanding. We show this by providing a model checking algorithm for VL that runs in polynomial space. For hardness, we use the classic reduction from quantified Boolean formulas.

Theorem 6. *The model checking problem for VL is PSPACE-complete.*

Proof. To show that the model checking problem for VL is in *PSPACE*, we present Algorithm 1. For the sake of brevity, we focus on the dynamic modality, and the interested reader can find more on model checking hybrid logics in [15].

Algorithm 1 An algorithm for model checking VL

procedure MC(M, a, φ)
  case $\varphi = [p]\psi$
    return MC($M^{a:p}, a, \psi$)
  case $\varphi = [\pi \cup \tau]\psi$
    return MC($M, a, [\pi]\psi$) and MC($M, a, [\tau]\psi$)

The algorithm follows the semantics and its correctness can be shown via induction on φ. Now we argue that the algorithm requires at most polynomial space. The interesting case here is $\varphi = [p]\psi$. Without giving an explicit algorithm for constructing $M^{a:p}$, we note that the size of $M^{a:p}$ is bounded by $\mathcal{O}(|M|^2)$ (the worst-case scenario of $R(p, a)$ and F being universal). Since there are at most $|\varphi|$ symbols in φ, the total space required by the algorithm is bounded by $\mathcal{O}(|\varphi| \cdot |M|^2)$.

To show hardness of the model checking problem we use the classic reduction from the satisfiability of quantified Boolean formula: given a QBF $\Psi :=$

$Q_1 p_1 \dots Q_n p_n \psi(p_1, \dots, p_n)$, where $Q_i \in \{\forall, \exists\}$, determine whether Ψ is true. To reduce the satisfiability of QBF Ψ to the model checking of VL, we construct a model M_a and a formula Ψ' of $\mathcal{VL}$ such that Ψ is true if and only if $M_a \models \Psi'$.

More specifically, given a QBF $Q_1 p_1 \dots Q_n p_n \psi(p_1, \dots, p_n)$, we construct a visibility model $M = (A, F, +, -, V, R)$, where $A = \{a_0, \dots, a_n\}$, $F(a_i) = \{a_0\}$ for all $i \neq 0$, $+(a_i) = p_i$ for all $i \neq 0$, $-(a_i) = \emptyset$ for all i, $V(i_j) = \{a_j\}$, and $R(p, a) = \emptyset$ for all $p \in \mathsf{Top}$ and $a \in A$. Additionally, we assume that there is a topic q that no agent is either *pro* or *contra*. Intuitively, M is a model consisting of $n+1$ agents, where everyone follows agent a_0, who follows no one. Each agent, apart from a_0, is *pro* exactly one topic, and no one is *contra* anything. Finally, the translation of the QBF is done recursively as follows:

$$\psi'_0 := \psi(\Diamond_{i_0:p_1}(i_1 \wedge \Diamond_{i_0:p_1}\top), \dots, \Diamond_{i_0:p_n}(i_n \wedge \Diamond_{i_0:p_n}\top))$$

$$\psi'_k := \begin{cases} [p_k \cup q]\psi'_{k-1} & \text{if } Q_k = \forall \\ \neg[p_k \cup q]\neg\psi'_{k-1} & \text{if } Q_k = \exists \end{cases}$$

$$\psi' := \psi'_n.$$

We need to show that

$$Q_1 x_1 \dots Q_n x_n \psi(p_1, \dots, p_n) \text{ is satisfiable iff } M_a \models \psi'.$$

Agent a_0 posting on topic p_i means that the truth-value of p_i has been set to 1. If agent a_0 posts on topic q, this means that the truth-value of the corresponding p_i has been set to 0. Since there are no two agents that are *pro* the same topic, the choice of truth values is unambiguous.

We use non-deterministic choice to model quantifiers. The universal quantifier $\forall p_k$ is emulated with $[p_k \cup q]\psi'_{k-1}$ meaning that no matter what agent a_0 chooses to post on, p_k or q, formula ψ'_{k-1} will be true. Similarly, the existential quantifier $\exists p_k$ is emulated with $\neg[p_k \cup q]\neg\psi'_{k-1}$ meaning that agent a_0 can post on a topic, either p_k or q, to make ψ'_{k-1} true. Finally, propositional variable p_j is translated into the formula $\Diamond_{i_0:p_j}(i_j \wedge \Diamond_{i_0:p_j}\top)$ that is true if and only if there has been a post on p_j, and the corresponding agent a_j, who is *pro* p_j, has reposted it. For all other agents a_k, the formula will not hold. Posting on q instead of p_j results in the fact that $\Diamond_{i_0:p_j}(i_j \wedge \Diamond_{i_0:p_j}\top)$ is not satisfied anywhere in the model, thus corresponding to setting p_j to 0. $\square$

5 Conclusion and Future Work

This work was devoted to the analysis of the concept of visibility in social networks using modal logic. After discussing related work from social network analysis, we introduced a logic we named static visibility logic (SVL) and its dynamic extension, visibility logic (VL). We did not give a definite answer as to how one should measure visibility, but proposed several quantitative and qualitative measures relevant to our social network models. To motivate VL, we presented an example where we show how, given some simple rules of the system, a potential malicious agent can take advantage of the network to expose more agents to a controversial opinion. On the mathematical side, we showed soundness and completeness of SVL with respect to social networks that follow our given rules. We also proved that the language of VL is strictly more expressive than the language of SVL and that the complexity of the model checking problem for VL is *PSPACE*-complete.

As we mention in the paper, an implication of the result $\mathcal{SVL} < \mathcal{VL}$ is that a proof of the completeness of VL using reduction axioms is not possible. Thus one of the open problems is to find a sound and complete axiomatisation of VL.

Another direction for future work is to explore triggering in social network communication. The idea is that seeing a post on a controversial topic might trigger an agent to post a reaction. To do this, we could expand our framework such that agents can not only post on a topic, but also *pro* or *contra* a topic. This entails letting $\pi ::= p \mid p^+ \mid p^- \mid \pi \cup \pi$ in the dynamic formula $[\pi]\phi$. Then, we could specify particular controversial topics and add a rule stating that if an agent sees a post that is pro the controversial topic and they are themselves contra, then the agent will post contra the topic, or vice versa. We leave a proper implementation to a later possible extended version of the paper.

Acknowledgements We would like to thank anonymous reviewers for their insightful comments and constructive criticism.

References

1. Andi, S.: How and why do consumers access news on social media? (2021), https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2021/how-and-why-do-consumers-access-news-social-media
2. Areces, C., ten Cate, B.: Hybrid logics. In: Blackburn, P., van Benthem, J., Wolter, F. (eds.) Handbook of Modal Logic, Studies in Logic and Practical Reasoning, vol. 3, pp. 821–868. Elsevier (2007)
3. Baltag, A., Christoff, Z., Rendsvig, R.K., Smets, S.: Dynamic epistemic logics of diffusion and prediction in social networks. Stud. Log. **107**(3), 489–531 (2019)
4. Baltag, A., Moss, L.S.: Logics for epistemic programs. Synthese **139**(2), 165–224 (2004)
5. Belardinelli, G.: Gatekeepers in Social Networks: Logics for Communicative Actions. Master's thesis, ILLC, University of Amsterdam (2019)
6. Blackburn, P., ten Cate, B.: Pure extensions, proof rules, and hybrid axiomatics. Stud. Log. **84**(2), 277–322 (2006)
7. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic, Cambridge Tracts in Theoretical Computer Science, vol. 53. CUP (2001)
8. Brady, W.J., Wills, J.A., Jost, J.T., Tucker, J.A., Bavel, J.J.V.: Emotion shapes the diffusion of moralized content in social networks. Proc. of the National Academy of Sciences **114**(28), 7313–7318 (2017)
9. Christoff, Z., Hansen, J.U.: A logic for diffusion in social networks. J. Appl. Log. **13**(1), 48–77 (2015)
10. Christoff, Z., Hansen, J.U., Proietti, C.: Reflecting on social influence in networks. J. Logic Lang. Inf. **25**(3-4), 299–333 (2016)

11. Dennis, L.A., Fu, Y., Slavkovik, M.: Markov chain model representation of information diffusion in social networks. J. Log. Comput. (2022)
12. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic, Synthese Library, vol. 337. Springer (2008)
13. Easley, D., Kleinberg, J.: Networks, Crowds and Markets. CUP (2010)
14. Fischer, M.J., Ladner, R.E.: Propositional dynamic logic of regular programs. J. Comput. Syst. Sci. **18**(2), 194–211 (1979)
15. Franceschet, M., de Rijke, M.: Model checking hybrid logics (with an application to semistructured data). J. Appl. Log. **4**(3), 279–304 (2006)
16. Jackson, M.O., Rogers, B.W., Zenou, Y.: The economic consequences of social-network structure. J. Econ. Lit. **55**(1), 49–95 (2017)
17. Kooi, B., Renne, B.: Arrow update logic. Rev. Symb. Log. **4**(4), 536–559 (2011)
18. Kurtonina, N., de Rijke, M.: Bisimulations for temporal logic. J. Logic Lang. Inf. **6**(4), 403–425 (1997)
19. Lorini, E., Sartor, G.: A STIT logic for reasoning about social influence. Stud. Log. **104**(4), 773–812 (2016)
20. Napoli, P.M.: Social media and the public interest: Governance of news platforms in the realm of individual and algorithmic gatekeepers. Telecomm. Policy **39**(9), 751–760 (2015)
21. Nguyen, C.T.: Echo chambers and epistemic bubbles. Episteme **17**(2), 141–161 (2020)
22. Pedersen, M.Y., Slavkovik, M., Smets, S.: Social bot detection as a temporal logic model checking problem. In: Ghosh, S., Icard, T. (eds.) Proc. of the 8th LORI. LNCS, vol. 13039, pp. 158–173. Springer (2021)
23. Pedersen, M.Y., Smets, S., Ågotnes, T.: Analyzing echo chambers: A logic of strong and weak ties. In: Blackburn, P., Lorini, E., Guo, M. (eds.) Proc. of the 7th LORI. LNCS, vol. 11813, pp. 183–198. Springer (2019)
24. Pedersen, M.Y., Smets, S., Ågotnes, T.: Further steps towards a logic of polarization in social networks. In: Dastani, M., Dong, H., van der Torre, L. (eds.) Proc. of the 3rd CLAR. pp. 324–345. LNCS, Springer (2020)
25. Pedersen, M.Y., Smets, S., Ågotnes, T.: Modal logics and group polarization. J. Log. Comput. **31**(8), 2240–2269 (2021)
26. Plaza, J.: Logics of public communications. Synthese **158**(2), 165–179 (2007)
27. Rathore, N.C., Tripathy, S.: Effective visibility prediction on online social network. IEEE Trans. Comput. Soc. Syst. (2021)
28. Samanta, M., Pal, P., Mukherjee, A.: A novel scheme for abatement of privacy concern by controlling the reachability in online social network. In: Abraham, A., Cherukuri, A.K., Madureira, A.M., Muda, A.K. (eds.) Proc. of the 8th SoCPaR. Adv. Intell. Syst. Comput., vol. 614, pp. 686–697. Springer (2016)
29. Tang, J., Musolesi, M., Mascolo, C., Latora, V.: Characterising temporal distance and reachability in mobile and online social networks. Comput. Commun. Rev. **40**(1), 118–124 (2010)