

The Egocentric Logic of Preferences

Junli Jiang¹ and Pavel Naumov²

¹Institute of Logic and Intelligence, Southwest University, China

²University of Southampton, the United Kingdom

walk08@swu.edu.cn, p.naumov@soton.ac.uk

Abstract

The paper studies preferences of agents about other agents in a social network. It proposes a logical system that captures the properties of such preferences, called “likes”. The system can express nested constructions “agent likes humbled people”, “agent likes those who like humbled people”, etc. The main technical results are a model checking algorithm and a sound, complete, and decidable axiomatization of the proposed system.

1 Introduction

INVESTIGATOR: Have your student ever advocated the overthrow of the United States Government?

As a young faculty who just started to teach in the United States, the second author of this paper did not expect this question. He instantly understood that answering “yes” would jeopardize the student’s chances to get security clearance required for the student’s job with the US Government. However, he did not know that answering “yes” to this question and expressing a sympathy for the student could disqualify the author from ever working for the US Department of Defense. Indeed, under US law, “. . . sympathy with persons or organizations that advocate the overthrow of the United States Government”¹ is a disqualifying factor for a security clearance at the US Department of Defense.

In this paper we propose a logical system for reasoning about preferences (or “likes”) of agents about other agents which is based on *egocentric logic* [Prior, 1968]. The egocentric approach has also been used in [Grove and Halpern, 1991; Grove and Halpern, 1993; Grove, 1995; Seligman *et al.*, 2013; Epstein and Naumov, 2021]. In an egocentric logic, the statements capture properties of the agents. As a result, the satisfaction is a relation between an agent and a formula:

the student \Vdash advocated the overthrow.

In general, we write $a \Vdash \varphi$ if property φ is true *about* agent a . In this paper, we propose modality L (“likes those who”) that captures agents’ preferences about other agents. For example,

it is clear from the opening example that the investigator preferred job applicants who have not advocated the overthrow of the US Government:

the investigator $\Vdash L\neg$ (“advocated the overthrow”).

In general, we write $a \Vdash L\varphi$ if agent a likes (prefers) any agent with property φ to an agent without property φ .

At this point, placement of an agent on the left-hand-side of the relation \Vdash seems to be not very justified. The reason for doing this appears clear only when one starts *nesting* modalities. Indeed, because the agent is on the left-hand-side of \Vdash , formula $L\varphi$ captures a *property of an agent*, just like formula φ does. As a result, it can be nested. Statement $LL\varphi$ means “likes those who like φ ” and statement $L\neg L\varphi$ means “likes those who do not like φ ”. For example,

the investigator $\Vdash L\neg L$ (“advocated the overthrow”). (1)

means that the investigator likes those who do not sympathise with the people advocating the overthrow of the US Government. Note that statements like the one above can not be expressed in traditional multiagent modal language where the agent is placed in subscript. This is because in such a language, the operator L_a in statement L_a (“advocated the overthrow”) takes a statement about agents as an argument and returns a statement about the world (“agent a likes those who advocate the overthrow”) as the value. Such operators cannot be nested.

Statement (1) illustrates the capability of an egocentric logic, but it does not capture correctly the security clearance requirements in the United States. Indeed, imagine a bit unusual situation when a job candidate does not sympathise with the people advocating the overthrow of the US Government, but yet advocates for this herself. Cited above Title 32 – National Defence of US code of federal regulations disqualifies such person from being employed by the Department of Defense as well. Thus,

the investigator $\Vdash L(\neg$ (“advocated the overthrow”) $\wedge \neg L$ (“advocated the overthrow”)).

It is interesting to note that the above regulation does *not* disqualify those who *sympathise with the sympathisers*:

the investigator $\not\Vdash L(\neg$ (“advocated the overthrow”) $\wedge \neg L$ (“advocated the overthrow”) $\wedge \neg LL$ (“advocated the overthrow”)).

¹<https://www.law.cornell.edu/cfr/text/32/147.3>

In this paper, we propose a formal semantics for modality L in the egocentric setting, describe a model checking algorithm for formulae that use this modality, and give a sound, complete, and decidable axiomatization of the properties of this modality. To illustrate the capability of our logical system, we add to it “for all friends” modality F as well as constants humbled η and egotistic ε . We assume that friends do not have opposite preferences, that humbled agents do not think that they are better than their friends, and that egotistic people do not think that they are worse than their friends.

The first logical system for reasoning about preferences, the logic of ‘better’, was proposed by Halldén [1957]. It captures the properties of the binary modality “statement p is better than q ”. Similar logical systems have been proposed in [Von Wright, 1963; Doyle *et al.*, 1991]. In these systems one can define “likes” modality $L\varphi$ to mean that “statement φ is better than statement $\neg\varphi$ ”. Lang, van der Torre, and Weydert considered a conditional version $L(\varphi|\psi)$ of this modality, which stands for “statement φ is preferred over $\neg\varphi$ assuming that ψ is true” [2002]. Van Benthem, Girard, and Roy proposed a logical system for a related modality “statement φ holds in all worlds better than current” [2009]. The same modality is also investigated in [Liu, 2011; Christoff *et al.*, 2021]. Lorini and Schwarzenrüber proposed to define desirability without using a preference relation. They say that a statement is desirable if it is true in all worlds indistinguishable from the current world that are labeled as “good” [2011]. Unlike us, none of these works used the egocentric approach. As a result, they cannot express statements “like those who like”, “like those who do not like”, etc.

The paper is structured as follows. First, in sections 2 through 4 we propose the formal semantics of modalities L and F as well as constants η and ε and illustrate them with two examples. Second, in Section 5, we describe a model checking algorithm for our logical system and analyse its complexity. Third, in Section 6, we give a sound, complete, and decidable logical system that captures the interplay between modalities L, F and constants η , ε . The proofs of the soundness and the completeness are omitted due to the space constraint.

2 Social Network

In this section, we introduce the notion of a social network that serves as a foundation for the semantics of our logic. In the paper, we assume a fixed set of propositional variables \mathcal{P} .

Definition 1. A tuple $(\mathcal{A}, \{\prec_a\}_{a \in \mathcal{A}}, \mathcal{F}, \mathcal{H}, \mathcal{E}, \pi)$ is called a social network if

1. \mathcal{A} is a (possibly empty) finite set of “agents”,
2. \prec_a is a strict partial order “preference” relation on \mathcal{A} , for each agent $a \in \mathcal{A}$,
3. $\mathcal{F}(a) \subseteq \mathcal{A}$ is a set for any agent $a \in \mathcal{A}$, called a “circle of friends”, such that
 - (i). $a \in \mathcal{F}(a)$ for all agents $a \in \mathcal{A}$,
 - (ii). if $b \prec_a c$, then $c \not\prec_f b$, for all agents $a, b, c \in \mathcal{A}$ and each agent $f \in \mathcal{F}(a)$,

4. $\mathcal{H} \subseteq \mathcal{A}$ is a set of “humbled” agents; it will be assumed that $f \not\prec_h h$ for any humbled agent $h \in \mathcal{H}$ and any agent $f \in \mathcal{F}(h)$,
5. $\mathcal{E} \subseteq \mathcal{A}$ is a set of “egotistic” agents; it will be assumed that $e \not\prec_e f$ for any egotistic agent $e \in \mathcal{E}$ and any agent $f \in \mathcal{F}(e)$,
6. $\pi(p)$ is a subset of \mathcal{A} for each variable $p \in \mathcal{P}$.

Item 3(i) of the above definition states that each agent is her own friend. We make this assumption to simplify one of the axioms of our logical system. Our results in this paper can be easily adjusted if this assumption is removed. Item 3(ii) states that friends cannot have opposite preferences. The similarity of preferences between friends has been observed in psychology [Selfhout *et al.*, 2009] and it is the underlying assumption of the balance theory [Harary, 1953]. Note that we do not assume that friendship is symmetric: if $a \in \mathcal{F}(b)$, then it is not necessarily that $b \in \mathcal{F}(a)$. Items 4 and 5 capture our assumptions that humbled agents do not think that they are better than their friends and that egotistic agents do not think that they are worse than their friends.

In this paper, we primarily use terms “friend”, “humbled”, and “egotistic” to show-off the expressive power of our language. We do not imply that items 3, 4, and 5 capture the full depth of human friendship or the notions of humbleness and egotism.

Note that propositional variables, just like all formulae in our language, represent noun-free fragments of propositions such as “is humbled” or “is rich”. Prior calls them *subjectless predicates* [1968]. Grove and Halpern use term *relative propositions* because they are relative to an agent [Grove and Halpern, 1991; Grove and Halpern, 1993; Grove, 1995]. To reflect this, item 6 of Definition 1 specifies the valuation $\pi(p)$ of a proposition variable p as a set of agents. Informally, these are the agents for which propositional variable p is true.

Let us now turn to an example of a social network. Recently, a lot has been said about polarisation of opinions in the society. From Brexit debates [Del Vicario *et al.*, 2017], ethnic integration in Netherlands [Oosterwaal and Torenvlied, 2010], to politics in the US [Sides and Hopkins, 2015], it has been observed that the people split into two groups that do not interact and do not value the opinions of each other. In reality, there is nothing new about such split. The same situation existed in 18th century England, as highlighted in political satire book Gulliver’s Travels [Swift, 1995]. Following Swift, we capture such polarisation in Lilliput’s Social Network depicted in Figure 1. All agents in the network are divided into big endians and little endians. We assume that there is at least one agent in each group. Directed edges labeled with a group of agents represent the preferences of all agents in the group. Thus, each big endian prefers each big endian over each little endian and each little endian prefers each little endian over each big endian. We also assume that each agent a in this setting has a set of friends $\mathcal{F}(a)$, not depicted in the figure.

Proposition 1. A little endian cannot be a friend of a big endian. A big endian cannot be a friend of a little endian.

Proof. Suppose that there is a little endian l and a big endian




Figure 1: Lilliput's Social Network.

b such that $l \in \mathcal{F}(b)$. Thus, $l \prec_b b$ and

$$b \prec_l l, \quad (2)$$

see Figure 1. Statement $l \prec_b b$ implies that $b \not\prec_l l$ by item 3(ii) of Definition 1 and the assumption $l \in \mathcal{F}(b)$, which contradicts to statement (2). The proof of the other statement is similar. \square

Although the Lilliput's Social Network is based on a fictional book, castes define a similar social structure in India. There, similar to our claim in Proposition 1, "caste was and is a defining factor in friendships" [Dhanaraj, 2016]. In the United States, about forty percent of registered voters do not have a *single* close friend who supports the other major party candidate [Dunn, 2020].

3 Syntax and Semantics

The language Φ of our system is defined by the grammar:

$$\varphi := \eta \mid \varepsilon \mid p \mid \neg\varphi \mid \varphi \rightarrow \varphi \mid F\varphi \mid L\varphi.$$

We read η as "is humbled", ε as "is egotistic", F as "for all friends", and L as "likes those who". We assume that \bar{F} is an abbreviation for $\neg F\neg$. We read \bar{F} as "has a friend who". We assume that disjunction \vee , conjunction \wedge , and biconditional \leftrightarrow are defined in our language in the standard way.

Next, we define the semantics of our logical system.

Definition 2. For any agent $a \in \mathcal{A}$ of a social network $(\mathcal{A}, \{\prec_a\}_{a \in \mathcal{A}}, \mathcal{F}, \mathcal{H}, \mathcal{E}, \pi)$ and any formula $\varphi \in \Phi$, the satisfaction relation $a \Vdash \varphi$ is defined recursively as follows.

1. $a \Vdash \eta$, if $a \in \mathcal{H}$,
2. $a \Vdash \varepsilon$, if $a \in \mathcal{E}$,
3. $a \Vdash p$, if $a \in \pi(p)$,
4. $a \Vdash \neg\varphi$, if $a \not\Vdash \varphi$,
5. $a \Vdash \varphi \rightarrow \psi$, if $a \not\Vdash \varphi$ or $a \Vdash \psi$,
6. $a \Vdash F\varphi$, if $f \Vdash \varphi$ for all agents $f \in \mathcal{F}(a)$,
7. $a \Vdash L\varphi$, when for all agents $b, c \in \mathcal{A}$, if $b \not\prec_a c$ and $c \Vdash \varphi$, then $b \prec_a c$.

Some of the existing works on preferences are based on *ceteris paribus* principle [Von Wright, 1963]. If this principle is applied to modality "like", then one would say that an agent a likes φ if she prefers agents for whom φ is true to those for whom φ is not true, *everything else being equal*. Capturing

"everything else being equal" assumption semantically is a non-trivial task. One of possible ways to do it is proposed in [Van Benthem *et al.*, 2009]. In this paper, we do *not* use *ceteris paribus* principle. Item 7 of the above definition states that an agent a likes φ if she prefers *each* agent for whom φ is true to *each* agent for whom φ is not true.

We use $A\varphi$ as an abbreviation for formula $\varphi \wedge L\varphi \wedge L\neg\varphi$. Because of the next lemma, we read A as "for all agents".

Lemma 1. $a \Vdash A\varphi$ iff $b \Vdash \varphi$ for all agents $b \in \mathcal{A}$.

Proof. (\Rightarrow): Suppose that there is an agent $b \in \mathcal{A}$ such that $b \not\Vdash \varphi$. Note that the assumption $a \Vdash A\varphi$ of the lemma implies that $a \Vdash \varphi$ and $a \Vdash L\varphi$. Then, by item 7 of Definition 2, the assumption $b \not\Vdash \varphi$ implies that $b \prec_a a$.

At the same time, the assumption $b \not\Vdash \varphi$ implies $b \Vdash \neg\varphi$ by item 4 of Definition 2. Note that the assumption $a \Vdash A\varphi$ of the lemma also implies that $a \Vdash \varphi$ and $a \Vdash L\neg\varphi$. Then, $a \prec_a b$ again by item 7 of Definition 2.

Finally, note that statements $b \prec_a a$ and $a \prec_a b$ are inconsistent because relation \prec_a is a strict partial order by item 2 of Definition 1.

(\Leftarrow): Suppose that $a \not\Vdash \varphi \wedge L\varphi \wedge L\neg\varphi$. Thus, one of the following cases takes place.

Case I: $a \not\Vdash \varphi$. Then, there is an agent b such that $b \not\Vdash \varphi$.

Case II: $a \not\Vdash L\varphi$. Hence, by item 7 of Definition 2, there are agents $b, c \in \mathcal{A}$ such that $b \not\Vdash \varphi$, $c \Vdash \varphi$, and $b \not\prec_a c$. So, there is an agent b such that $b \not\Vdash \varphi$.

Case III: $a \not\Vdash L\neg\varphi$. Thus, again by item 7 of Definition 2, there are agents $b, c \in \mathcal{A}$ such that $b \not\Vdash \neg\varphi$, $c \Vdash \neg\varphi$, and $b \not\prec_a c$. Therefore, $c \not\Vdash \varphi$ by item 4 of Definition 2. \square

The lemma below follows from item 7 of Definition 2.

Lemma 2. For any social network, if $a \Vdash \varphi \leftrightarrow \psi$ for each agent $a \in \mathcal{A}$, then $a \Vdash L\varphi \leftrightarrow L\psi$ for each agent $a \in \mathcal{A}$. \square

4 Examples

In this section, we illustrate the formal semantics of modalities L , F given in Definition 2.

4.1 Lilliput's Social Network

First, let us return to the Lilliput's Social Network from Figure 1 and consider several statements with modality L . The proposition below shows that an agent likes little endians if and only if she herself is a little endian.

Proposition 2. $a \Vdash L(\text{"is a little endian"})$ iff agent a is a little endian.

The assumption on page 2 that the set of little endians is nonempty is important for Proposition 2. Indeed, by item 7 of Definition 2, the statement $a \Vdash L(\text{"is a little endian"})$ is vacuously true for any agent a if the set of little endians is empty.

The next proposition states that an agent likes big endians if and only if she herself is a big endian.

Proposition 3. $a \Vdash L(\text{"is a big endian"})$ iff agent a is a big endian.

The proposition below states that an agent *likes those who like little endians* if and only if she herself is a little endian.

Proposition 4. $a \Vdash \text{LL}(\text{“is a little endian”})$ iff agent a is a little endian.

Finally, the proposition states that an agent *likes those who do not like little endians* if and only if she herself is a big endian.

Proposition 5. $a \Vdash \text{L}\neg\text{L}(\text{“is a little endian”})$ iff agent a is a big endian.

4.2 The World According to Steve Jobs

This example originates from the quote

“Steve Jobs has a saying that *A* players hire *A* players; *B* players hire *C* players; and *C* players hire *D* players.” [Kawasaki, 2004, p. 101]

We capture this quote through the diagram depicted in Figure 2. Here all agents are partitioned into groups *A*, *B*, *C*, and *D*. As with our previous example, we assume that each group is not empty. The hiring preference relation of each group is depicted in the diagram using directed edges.




Figure 2: The World According to Steve Jobs.

The following propositions about *A*-, *B*-, *C*-, and *D*-players are true in the context of this example.

Proposition 6. *A*-players can not be friends with *B*-players or *C*-players. *B*-players cannot be friends with *C*-players.

Proof. To prove that *A*-players can not be friends with *B*-players, consider any *A*-player a and any *B*-player b . Recall that we assume that each group contains at least one player. Let “third player” be any *C*-player c . Then, $c \prec_a a$ and $a \prec_b c$, see Figure 2. Therefore, $a \notin \mathcal{F}(b)$ and $b \notin \mathcal{F}(a)$ by item 3(ii) of Definition 1. The proofs of the other parts of the proposition are similar, but one needs to consider a *D*-player as a “third player”. \square

Proposition 7. Humbled *A*-players cannot have friends among *D*-players.

Proof. Suppose a and d are, respectively, an *A*-player and a *D*-player such that $a \in \mathcal{H}$ and $d \in \mathcal{F}(a)$. Then, $d \not\prec_a a$ by item 4 of Definition 1, which contradicts to Figure 2. \square

Note that $a \in \mathcal{F}(b)$ is not equivalent to $b \in \mathcal{F}(a)$ per Definition 1. In particular, in spite of the previous proposition, a humbled *A*-player can be a friend of a *D*-player.

Proposition 8. $a \Vdash \text{L}(\text{“is an A-player”})$ iff agent a is an *A*-player.

Proof. (\Rightarrow): By our assumption, each of the groups of players is nonempty. Let agent a' be any *A*-player and agent n be any non-*A*-player. Assumption $a \Vdash \text{L}(\text{“is an A-player”})$ implies that $n \prec_a a'$ by item 7 of Definition 2. Therefore, agent a is an *A*-player, see Figure 2.

(\Leftarrow): Suppose that agent a is an *A*-player. Consider any *A*-player a' and non-*A*-player n . By item 7 of Definition 2, it suffices to show that $n \prec_a a'$. Note that $n \prec_a a'$ is true because agents a and a' are *A*-players and agent n is not an *A*-player, see Figure 2. \square

The proofs of the next two propositions are similar.

Proposition 9. $a \Vdash \text{L}(\text{“is a C-player”})$ iff agent a is a *B*-player. \square

Proposition 10. $a \Vdash \text{L}(\text{“is a D-player”})$ iff agent a is a *C*-player. \square

Proposition 11. $a \not\vdash \text{L}(\text{“is a B-player”})$ for any agent a .

Proof. By our assumption, each of the groups of players is nonempty. Let agent b be any *B*-player and agent n be any non-*B*-player. Thus, $n \not\prec_a b$, see Figure 2. Therefore, we have $a \not\vdash \text{L}(\text{“is a B-player”})$ by item 7 of Definition 2. \square

Proposition 12. $a \Vdash \text{LL}(\text{“is an A-player”})$ iff agent a is an *A*-player.

Proposition 13. $a \Vdash \text{LL}(\text{“is a D-player”})$ iff agent a is a *B*-player.

We conclude this example with the observation that in our setting *D*-agents do not like any non-trivial group of people. In other words, they only can like the group of all agents and the empty group of agents.

Proposition 14. $a \Vdash \text{L}\varphi \rightarrow \text{A}\varphi \vee \text{A}\neg\varphi$ for each *D*-player a .

5 Model Checking

In this section, we describe a model checking algorithm for formulae in language Φ .

5.1 Model Representation

Before describing the algorithm, let us discuss how preference relation \prec_a is represented in the input. One possible way to describe a transitive relation is to give the set of pairs of agents in the relation. Another, a more compact one, is to define the relation as a transitive closure of the given set of pairs [Aho *et al.*, 1972].




Figure 3: A Social Network.

For example, to describe the relation depicted in Figure 3 the first way, one should list pairs (a_1, a_2) , (a_1, a_3) , (a_1, a_4) ,

$(a_1, a_5), (a_2, a_3), (a_2, a_4), (a_2, a_5), (a_3, a_4), (a_3, a_5),$ and (a_4, a_5) . To describe the same relation the second way, it is sufficient only to list pairs $(a_1, a_2), (a_2, a_3), (a_3, a_4),$ and (a_4, a_5) . In general, the size of the first representation could be as large as big O of square of the size of the second one. Since the efficiency of an algorithm is measured as a function of the size of its input, the larger is the size of the input representation, the more efficient the algorithm *appears* to be. An honest algorithm analysis should use the most compact representation of the data. Following this principle, we assume that the preference relation in a social network is defined as a transitive closure of the set of pairs given in the description of the relation.

5.2 Multipath Problem

We describe the model checking algorithm in the next section. Here we discuss a graph problem whose solution is used by the model checking algorithm.

Let us first review some terminology. A directed graph (V, E) is defined by a set of vertices V and a set of directed edges E . By a cut of the graph we mean a partitioning of the set V into two disjoint sets. We allow either of these two sets to be empty. An example of a cut is depicted in Figure 4, where the vertices are divided into white and black.




Figure 4: Graph Partitioning

Problem. Given a cut (X, Y) of a directed acyclic graph (DAG), decide if there is a directed path from each vertex in set X to each vertex in set Y .

For example, if X is the set of all white vertices in Figure 4 and Y is the set of black vertices, then the answer is negative because there is no directed path from vertex b to vertex d .

The brute-force approach to multipath problem is to use depth-first-search algorithm to verify for each vertex $x \in X$ and each vertex $y \in Y$ that there is a directed path from x to y . The depth-first-search takes $O(|V| + |E|)$ time, where $|V|$ is the number of vertices and $|E|$ is the number of edges in the directed acyclic graph (V, E) . Thus, the brute-force approach execution time is $O(|V|^3 + |V|^2 \cdot |E|)$.

The multipath problem can be solved in $O(|V|^2 + |V| \cdot |E|)$ using a different approach, consisting of two steps.

On the first step, topological sorting algorithm is used to generate a linear ordering of the vertices of the graph: v_1, \dots, v_n .

On the second step, dynamic programming is used to recursively compute for each vertex v_i the set of all vertices $R(v_i) \subseteq X$ from which vertex v_i is reachable by a directed

path:

$$R(v_i) = \begin{cases} \{v_i\} \cup R'(v_i), & \text{if } v_i \in X, \\ R'(v_i), & \text{if } v_i \notin X, \end{cases} \quad (3)$$

where

$$R'(v_i) = \bigcup_{(v_j, v_i) \in E} R(v_j). \quad (4)$$




Figure 5: Topological Sort

Figure 5 illustrates this algorithm on the example of the DAG depicted in Figure 4. We assume that the topological sorting algorithm returned ordering c, e, b, d, a, h, f . In Figure 5, sets $R(v)$ are depicted in the rectangles under each vertex. Thus, for example, $R(d) = \{c, e\}$, because $\{c, e\}$ is the set of all white vertices from which vertex d can be reached. The content of the rectangles is computed recursively using equation (3). For example, $R(h) = R(b) \cup R(d)$ because vertex h does not belong to set X and it has only two incoming edges: (b, h) and (d, h) .

After sets $R(v)$ are computed for each $v \in V$, to get the answer to multipath problem, one only needs to check if $R(v) = X$ for each vertex $v \in Y$. In our example, the answer is negative because $R(d) = \{c, e\} \neq \{b, c, e\} = X$.

It is well known that topological sorting could be done in $O(|V| + |E|)$ time using, for example, depth-first-search algorithm [Cormen *et al.*, 2009, pp. 612-614]. We represent sets of vertices as Boolean arrays. Namely, if $V = \{v_1, \dots, v_n\}$ and $R \subseteq V$, then set R is represented by a Boolean array $\mathbf{a}[\cdot]$ of length n such that $\mathbf{a}[\mathbf{i}] = \tau$ iff $v_i \in R$. Under such representation, it takes $O(|V|)$ time to compute the union of any two subsets of V . Note that the total number of the unions computed during the second step of our algorithm is $O(|V| + |E|)$ because each edge (v_j, v_i) in formula (4) is used exactly once during the entire computation. Thus, the total execution time of the second step is $O(|V|(|V| + |E|))$. Therefore, the total execution time of the described solution of the multipath problem is $O(|V| + |E|) + O(|V|(|V| + |E|)) = O(|V|^2 + |V| \cdot |E|)$.

5.3 Model Checking Problem

By the model checking problem we mean deciding if $a \models \varphi$ is true for a given formula φ and a given agent a of a given social network. In this subsection, we show how the multipath problem from the previous subsection could be used to solve the model checking problem for language Φ . For an arbitrary social network $(\mathcal{A}, \{\prec_a\}_{a \in \mathcal{A}}, \mathcal{F}, \mathcal{H}, \mathcal{E}, \pi)$, by $|\mathcal{A}|$ we mean the number of agents in this network. By $|\prec_a|$ we mean the size of relation \prec_a using the second (more compact) representation. For example, the size of the strict order depicted in Figure 3 is 4. By $|\{\prec_a\}_{a \in \mathcal{A}}|$ we denote $\sum_{a \in \mathcal{A}} |\prec_a|$. Finally,

by $|\varphi|$ we mean the size (number of symbols) of a formula $\varphi \in \Phi$.

Theorem 1. *Model checking of formula φ for a given agent in a given social network could be accomplished in time $O(|\varphi|(|\mathcal{A}|^3 + |\mathcal{A}| \cdot |\{\prec_a\}_{a \in \mathcal{A}}|))$ in the worst case.*

Proof. To check $a \Vdash \varphi$, we consider a dynamic programming algorithm that for each subformula ψ of φ recursively checks for each agent $b \in \mathcal{A}$ if $b \Vdash \psi$ is true and stores the result for the future use.

If formula ψ is η , ε , or a propositional variable, then checking if $b \Vdash \psi$ is true takes a constant time. In this case, checking $b \Vdash \psi$ for each agent $b \in \mathcal{A}$ takes time $O(|\mathcal{A}|)$. If formula ψ has either the form $\neg\psi_1$ or the form $\psi_1 \rightarrow \psi_2$, then $b \Vdash \psi$ can be checked in a constant time using stored results for $b \Vdash \psi_1$ and $b \Vdash \psi_2$. Then, checking $b \Vdash \psi$ for each agent $b \in \mathcal{A}$ takes time $O(|\mathcal{A}|)$.

If formula ψ has the form $F\psi_1$, then, checking $b \Vdash F\psi_1$ takes time $O(|\mathcal{F}(b)|)$ using pre-computed answers for validity of $x \Vdash \psi_1$, where $x \in \mathcal{F}(b)$. Since $\mathcal{F}(b) \subseteq \mathcal{A}$, it takes $O(|\mathcal{A}|^2)$ to check $b \Vdash F\psi_1$ for all agents $b \in \mathcal{A}$.

If formula ψ has the form $L\psi_1$, then to check if $b \Vdash \psi$ is true it suffices to check that there is a \prec_b -path from each element of set $X = \{x \in \mathcal{A} \mid x \not\Vdash \psi_1\}$ to each element of set $Y = \{y \in \mathcal{A} \mid y \Vdash \psi_1\}$. This can be accomplished in time $O(|\mathcal{A}|^2 + |\mathcal{A}| \cdot |\prec_b|)$ using the solution of the multipath problem described in the previous subsection. Then, checking $b \Vdash \psi$ for all agents $b \in \mathcal{A}$ takes

$$\begin{aligned} & O\left(\sum_{b \in \mathcal{A}} (|\mathcal{A}|^2 + |\mathcal{A}| \cdot |\prec_b|)\right) = \\ & O\left(\sum_{b \in \mathcal{A}} |\mathcal{A}|^2 + |\mathcal{A}| \cdot \sum_{b \in \mathcal{A}} |\prec_b|\right) \\ & = O(|\mathcal{A}|^3 + |\mathcal{A}| \cdot |\{\prec_b\}_{b \in \mathcal{A}}|). \end{aligned}$$

Thus, in the worst case, it takes $O(|\mathcal{A}|^3 + |\mathcal{A}| \cdot |\{\prec_b\}_{b \in \mathcal{A}}|)$ time to check for each agent $b \in \mathcal{A}$ if $b \Vdash \psi$ is true. Therefore, it takes $O(|\varphi|(|\mathcal{A}|^3 + |\mathcal{A}| \cdot |\{\prec_b\}_{b \in \mathcal{A}}|))$ to do the same for all subformulae ψ of formula φ . \square

6 Axioms

In this section, we give a sound, complete, and decidable axiomatization of our logical system. In addition to tautologies in language Φ , the system contains the following axioms.

1. Distributivity: $A(\varphi \rightarrow \psi) \rightarrow (A\varphi \rightarrow A\psi)$ and $F(\varphi \rightarrow \psi) \rightarrow (F\varphi \rightarrow F\psi)$,
2. Friends are Agents: $A\varphi \rightarrow F\varphi$,
3. Self-Friendship: $F\varphi \rightarrow \varphi$,
4. Negative Introspection: $\neg A\varphi \rightarrow A\neg A\varphi$,
5. Substitution: $A(\varphi \leftrightarrow \psi) \rightarrow (L\varphi \rightarrow L\psi)$,
6. Friends Think Alike: $L\varphi \wedge \bar{L}\psi \rightarrow A(\varphi \rightarrow \psi) \vee A(\psi \rightarrow \varphi)$,
7. Humbleness: $\eta \wedge \varphi \wedge L\varphi \rightarrow F\varphi$,
8. Egotisticness: $\varepsilon \wedge \neg\varphi \wedge L\varphi \rightarrow F\neg\varphi$.

The meaning of the first five axioms above is straightforward. The Friends Think Alike axiom captures the requirement of item 3(ii) of Definition 1 that preferences of the friends must be compatible. Specifically, it states that if an agent likes those for whom φ is true and a friend of the agent likes those for whom ψ is true, then either each φ -agent is a ψ -agent or each ψ -agent is a φ -agent. This is the most non-trivial among the axioms of our system.

The Humbleness axiom is better understood in the form $\eta \wedge L\varphi \wedge \bar{F}\neg\varphi \rightarrow \neg\varphi$. In this form, it says that if a humbled agent likes those for whom φ is true and she has a friend for whom φ is false, then φ must be false for her as well. This property follows from item 4 of Definition 1.

Similarly, the Egotisticness axiom can be better understood if it is stated in the equivalent form $\varepsilon \wedge L\varphi \wedge \bar{F}\varphi \rightarrow \varphi$. In this form, it says that if an egotistic agent likes those for whom φ is true and she has a friend for whom φ is true, then φ must be true for her as well. This axiom captures the requirement of item 5 of Definition 1.

We write $\vdash \varphi$ and say that formula $\varphi \in \Phi$ is a *theorem* of our logical system if it can be derived from the above axioms using the Modus Ponens and the Necessitation inference rules:

$$\frac{\varphi, \varphi \rightarrow \psi}{\psi} \quad \frac{\varphi}{A\varphi}.$$

The proofs of the next two theorems are omitted due to the space constraint.

Theorem 2 (soundness). *If $\vdash \varphi$, then $a \Vdash \varphi$ for each agent a of each social network.*

Theorem 3 (completeness). *If $a \Vdash \varphi$ for each agent a of each social network, then $\vdash \varphi$.*

We conclude this section with an observation that our logical system is decidable, which follows from the completeness theorem.

Theorem 4. *Set $\{\varphi \in \Phi \mid \vdash \varphi\}$ is decidable.*

Proof. Set $\{\varphi \in \Phi \mid \vdash \varphi\}$ is recursively enumerable because it is axiomatizable. Set $\{\varphi \in \Phi \mid \not\vdash \varphi\}$ is recursively enumerable because our logical system is sound and complete with respect to social networks with finite sets of agents. Therefore, set $\{\varphi \in \Phi \mid \vdash \varphi\}$ is decidable. \square

7 Conclusion

We have proposed an egocentric logic of preferences. Unlike the previous works on preference and desire logics, the egocentric approach allows us to express nested statements of the form ‘‘agent a likes those who like humbled people’’. Our main technical results are a model checking algorithm and a sound, complete, and decidable logical system describing the properties of this modality.

Acknowledgments

Junli Jiang acknowledges the support of the Key Research Funds for the Key Liberal Science Research Base of Chongqing (NO.16SKB036) and the Fundamental Research Funds for the Central Universities (SWU1409412).

References

- [Aho *et al.*, 1972] Alfred V. Aho, Michael R Garey, and Jeffrey D. Ullman. The transitive reduction of a directed graph. *SIAM Journal on Computing*, 1(2):131–137, 1972.
- [Christoff *et al.*, 2021] Zoé Christoff, Norbert Gratzl, and Olivier Roy. Priority merge and intersection modalities. *The Review of Symbolic Logic*, pages 1–32, 2021.
- [Cormen *et al.*, 2009] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [Del Vicario *et al.*, 2017] Michela Del Vicario, Fabiana Zollo, Guido Caldarelli, Antonio Scala, and Walter Quattrociocchi. Mapping social dynamics on facebook: The brexit debate. *Social Networks*, 50:6–16, 2017.
- [Dhanaraj, 2016] Christina Thomas Dhanaraj. *Round Table India*, May 17, 2016. , <https://roundtableindia.co.in>.
- [Doyle *et al.*, 1991] Jon Doyle, Yoav Shoham, and Michael P Wellman. A logic of relative desire. In *International Symposium on Methodologies for Intelligent Systems*, pages 16–31. Springer, 1991.
- [Dunn, 2020] Amina Dunn. Few Trump or Biden supporters have close friends who back the opposing candidate. Pew Research Center, October 18, 2020. <https://www.pewresearch.org/fact-tank/2020/09/18/few-trump-or-biden-supporters-have-close-friends-who-back-the-opposing-candidate/>.
- [Epstein and Naumov, 2021] Sophia Epstein and Pavel Naumov. Epistemic logic of know-who. In *Proceedings of Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021.
- [Grove and Halpern, 1991] Adam J. Grove and Joseph Y. Halpern. Naming and identity in a multi-agent epistemic logic. In James F. Allen, Richard Fikes, and Erik Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR’91)*. Cambridge, MA, USA, April 22-25, 1991, pages 301–312. Morgan Kaufmann, 1991.
- [Grove and Halpern, 1993] Adam J Grove and Joseph Y Halpern. Naming and identity in epistemic logics part i: the propositional case. *Journal of Logic and Computation*, 3(4):345–378, 1993.
- [Grove, 1995] Adam J Grove. Naming and identity in epistemic logic part ii: a first-order logic for naming. *Artificial Intelligence*, 74(2):311–350, 1995.
- [Halldén, 1957] Sören Halldén. *On the Logic of “better”*. Almqvist & Wiksells, 1957.
- [Harary, 1953] Frank Harary. On the notion of balance of a signed graph. *The Michigan Mathematical Journal*, 2(2):143–146, 1953.
- [Kawasaki, 2004] Guy Kawasaki. *The art of the start: The time-tested, battle-hardened guide for anyone starting anything*. Penguin, 2004.
- [Lang *et al.*, 2002] Jérôme Lang, Leendert van der Torre, and Emil Weydert. Utilitarian desires. *Autonomous agents and Multi-agent Systems*, 5(3):329–363, 2002.
- [Liu, 2011] Fenrong Liu. *Reasoning about preference dynamics*, volume 354. Springer Science & Business Media, 2011.
- [Lorini and Schwarzenruber, 2011] Emiliano Lorini and François Schwarzenruber. A logic for reasoning about counterfactual emotions. *Artificial Intelligence*, 175(3):814–847, 2011.
- [Oosterwaal and Torenvlied, 2010] Annemarije Oosterwaal and René Torenvlied. Politics divided from society? three explanations for trends in societal and political polarisation in the netherlands. *West European Politics*, 33(2):258–279, 2010.
- [Prior, 1968] Arthur N Prior. Egocentric logic. *Noûs*, pages 191–207, 1968.
- [Selfhout *et al.*, 2009] Maarten HW Selfhout, Susan JT Branje, Tom FM ter Bogt, and Wim HJ Meeus. The role of music preferences in early adolescents’ friendship formation and stability. *Journal of adolescence*, 32(1):95–107, 2009.
- [Seligman *et al.*, 2013] Jeremy Seligman, Fenrong Liu, and Patrick Girard. Facebook and the epistemic logic of friendship. In *14th conference on Theoretical Aspects of Rationality and Knowledge (TARK ‘13), January 2013, Chennai, India*, pages 229–238, 2013.
- [Sides and Hopkins, 2015] John Sides and Daniel J Hopkins. *Political polarization in American politics*. Bloomsbury Publishing USA, 2015.
- [Swift, 1995] Jonathan Swift. Gulliver’s travels. In *Gulliver’s Travels*. Springer, 1995.
- [Van Benthem *et al.*, 2009] Johan Van Benthem, Patrick Girard, and Olivier Roy. Everything else being equal: A modal logic for ceteris paribus preferences. *Journal of philosophical logic*, 38(1):83–125, 2009.
- [Von Wright, 1963] Georg Henrik Von Wright. *The logic of preference*. Edinburgh University Press, 1963.