The Creation and Change of Social Networks: a logical study based on group size

Sonja Smets and Fernando R. Velázquez-Quesada

Institute for Logic, Language and Computation, Universiteit van Amsterdam. Amsterdam, The Netherlands. S.J.L.Smets@uva.nl, F.R.VelazquezQuesada@uva.nl

Abstract. This paper is part of an on-going programme in which we provide a logical study of social network formations. In the proposed setting, agent a will consider agent b as part of her network if the number of features (properties) on which they differ is small enough, given the constraints on the size of agent a's 'social space'. We import this idea about a limit on one's social space from the cognitive science literature. In this context we study the creation of new networks and use the tools of Dynamic Epistemic Logic to model the updates of the networks. By providing a set of reduction axioms we are able to provide sound and complete axiomatizations for the logics studied in this paper.

1 Introduction

While the study of social interactions has received a lot of attention in logic and AI, the existence of a specific social group or network on which these studies are based is typically taken for granted. So what is left mostly unexplored is the way a social group is formed or the way in which a social network is created. This is exactly the topic we address in this paper. As such, this proposal complements our previous work in [1] which provides a *threshold* based approach to social network formation. In the threshold setting, an agent a considers agent b as part of her social network if and only if the number of features in which they differ is smaller or equal than a given threshold θ. This paper follows a different approach by using an idea that arises from the cognitive science literature: focus not on a similarity threshold, but rather on the size of the agent's 'social space'. In real life, agents may be willing to keep expanding their social network with people who are decreasingly less similar from them, as long as there is still 'enough space' in their social environment.[1] This is famously known as the *Dunbar's number*: a suggested cognitive limit to the number of people with whom one can maintain stable social relationships (see, e.g., [2]).

In the next section we first introduce the social network models as a context in which we can specify a distance between agents in a network. This distance is then used to create a layered structure of an agent's possible social contacts,

[1] Think, for example, how we establish conversations with relatively 'distant' acquaintances mostly only when our close friends are not around.

which is an essential ingredient in the mechanism that allows agents to form a new social network or to extend a given one when they are asked to take into account the bound on their 'social space'. We study a logical system that can express such network creations, giving a sound and complete axiomatisation for it. Finally we focus on the representation of more refined scenarios in which not all features play the same important role in the network creation/formation process. We conclude with a series of ideas for possible generalizations and/or alternative settings that can be explored in future work.

2 Modelling Social Networks

Similar to [1], our starting point is the basic setting of [3] in which we work with a relational 'Kripke' model in which the domain is interpreted as the set of agents, the accessibility relation represents a social connection from one agent to another, and the atomic valuation describes the features (behavior/opinions) that each agent has. Let A denote a countable set of agents, and P (with $\mathsf{A} \cap \mathsf{P} = \varnothing$) a countable set of features that agents might or might not have:

Definition 2.1 (Social Network Model). *A* social network model (SNM) *is a tuple* $M = \langle \mathsf{A}, S, V \rangle$ *where* $S \subseteq \mathsf{A} \times \mathsf{A}$ *is the* social *relation (*Sab *indicates that agent* a *is socially connected to agent* b*) and* $V : \mathsf{A} \to \wp(\mathsf{P})$ *is a* feature *function (*$p \in V(a)$ *indicates that agent* a *has feature* p*).*

Note how the social relation S does not need to satisfy any specific property (in particular, it is not required to be irreflexive, and neither symmetric), and thus it differs from the *friendship* relation of other approaches (e.g., [4,5,3,6]). Given a social network model, we define a notion of 'distance' between agents based on the number of features in which they differ.

Definition 2.2 (Distance). *Let* $M = \langle \mathsf{A}, S, V \rangle$ *be a* SNM*. Let* $\textsc{msmtch}_M(a, b)$ *be the set of features distinguishing agents* $a, b \in \mathsf{A}$ *in* M*:*

$$\textsc{msmtch}_M(a, b) := \mathsf{P} \setminus \{p \in \mathsf{P} : \ p \in V(a) \textit{ iff } p \in V(b)\}$$

Then, the distance *between* a *and* b *in* M *is given by*

$$\textsc{dist}_M(a, b) := |\textsc{msmtch}_M(a, b)|$$

As discussed in [1], $\textsc{dist}$ is a mathematical distance: for any agents $a, b \in \mathsf{A}$ and any SNM, ***(i)*** the distance from a to b is non-negative (non-negativity: $\textsc{dist}_M(a, b) \geq 0$), ***(ii)*** the distance from a to b is equal to that from b to a (symmetry: $\textsc{dist}_M(a, b) = \textsc{dist}_M(b, a)$), and ***(iii)*** the distance from an agent to herself is 0 (reflexivity: $\textsc{dist}_M(a, a) = 0$). Moreover, $\textsc{dist}$ is a *semi-metric*, as it also satisfies subadditivity: 'going directly' from a to c is 'faster' than 'going' via another agent ($\textsc{dist}_M(a, c) \leq \textsc{dist}_M(a, b) + \textsc{dist}_M(b, c)$). Still, $\textsc{dist}$ is not a *metric*, as it does not satisfy *identity of indiscernibles*: $\textsc{dist}_M(a, b) = 0$ does not imply $a = b$, as two different agents may have exactly the same features.[2]

[2] See [7, Chapter 1] for more details on mathematical distances.

Static Language $\mathcal{L}$. Following [3], social network models are described by a *propositional* language $\mathcal{L}$, with special atoms describing the agents' features and their social relationship:

Definition 2.3 (Language $\mathcal{L}$). *Formulas φ, ψ of the language $\mathcal{L}$ are given by*

$$\varphi, \psi ::= p_a \mid \mathrm{S}_{ab} \mid \neg\varphi \mid \varphi \wedge \psi$$

with $p \in \mathsf{P}$ and $a, b \in \mathsf{A}$. We read p_a as "agent a has feature p" *and* S_{ab} *as* "agent a is socially connected to b"*. Boolean constants ($\top, \bot$) and other Boolean operators ($\vee, \rightarrow, \leftrightarrow, \veebar$, the latter representing the exclusive disjunction) are defined as usual. Given a SNM $M = \langle \mathsf{A}, S, V \rangle$, the semantic interpretation of $\mathcal{L}$-formulas in M is given by:*

$$\begin{array}{llllll} M \Vdash p_a & \mathit{iff}_{def} & p \in V(a), & \quad M \Vdash \neg\varphi & \mathit{iff}_{def} & M \nVdash \varphi, \\ M \Vdash \mathrm{S}_{ab} & \mathit{iff}_{def} & Sab, & \quad M \Vdash \varphi \wedge \psi & \mathit{iff}_{def} & M \Vdash \varphi \text{ and } M \Vdash \psi. \end{array}$$

A formula $\varphi \in \mathcal{L}$ is valid (notation: $\Vdash \varphi$) when $M \Vdash \varphi$ holds for all models M.

Since there are no restrictions on the social relation nor on the feature function, any axiom system of classical propositional logic is fit to characterize syntactically the validities of $\mathcal{L}$ over the class of social network models.

3 Group-size-based social network creation

As mentioned before, [1] approaches social network creation by considering a similarity threshold θ, then defining each agent's new social space as all those agents that differ from her in at most $\theta \in \mathbb{N}$ features. This proposal follows a different strategy. Borrowing an idea from cognitive science [2], it considers a maximum group-size $\lambda \in \mathbb{N}$, then defining each agent's new social space as the λ agents that are closer to her, according to the above defined distance.

This section implements this idea of agents having a size-bounded social space; the following tools are used to make this idea precise.

Definition 3.1. *Given a social network model $M = \langle \mathsf{A}, S, V \rangle$ and an agent $a \in \mathsf{A}$, the quantitative notion of distance* $\textsc{dist}$ *induces a qualitative (total, reflexive, transitive and well-founded) relation $\preccurlyeq_a^M \subseteq \mathsf{A} \times \mathsf{A}$ of* distance from a. *Such relation is given by*

$$\preccurlyeq_a^M := \{(b_1, b_2) \in \mathsf{A} \times \mathsf{A} : \ \textsc{dist}_M(a, b_1) \leq \textsc{dist}_M(a, b_2)\},$$

and thus $b_1 \preccurlyeq_a^M b_2$ indicates that, in model M, agent b_1 is at least as close to agent a as agent b_2. By defining the notion of $\preccurlyeq_a^M$-minimum in the standard way (for $\mathsf{B} \subseteq \mathsf{A}$, take $\textsc{min}_a(\mathsf{B}) := \{b \in \mathsf{B} : \ b \preccurlyeq_a^M b' \text{ for all } b' \in \mathsf{B}\}$), this relation induces a sequence of layers (i.e., an ordered list of subsets) on A ($\mathrm{A}_a(-1), \mathrm{A}_a(0), \ldots, \mathrm{A}_a(n), \ldots$, for $n \geq 0$), with each set containing agents equally distant from a:

$$\mathrm{A}_a(-1) := \varnothing, \quad \mathrm{A}_a(0) := \textsc{min}_a(\mathsf{A}), \quad \mathrm{A}_a(n+1) := \textsc{min}_a(\mathsf{A} \setminus \bigcup_{k=-1}^{n} \mathrm{A}_a(k)).$$

Different agents might be 'equally distant' from a, and thus $\preccurlyeq_a^M$ is not antisymmetric: layers might have more than one element.[3] Moreover: while an initial empty layer $\mathrm{A}_a(-1)$ has been defined (its usefulness will be clear below), the layer $\mathrm{A}_a(0)$ always contains those agents that are feature-wise identical to a (including a herself). Note also how the layers are collectively exhaustive and pairwise disjoint: every agent appears in exactly one of them. Finally, when A is finite, at some point a 'first' empty layer $\mathrm{A}_a(k)$ will appear (for some $k > 0$), and from that moment on all layers will be empty too.

The layered structure of an agent's social contacts will be a helpful tool to model how agents can form a new social network or even extend a given one by performing updates on their social relations. Such agents are asked to establish new connections to agents that are close enough to them given the bound on their 'social space'. To model this we introduce the idea of a bounded similarity update operation on models, using the tools of Dynamic Epistemic Logic on how one can model such transformations on models [8,9,10].

Definition 3.2 (Bounded similarity update). *Let $M = \langle \mathsf{A}, R, V\rangle$ be a SNM; take $\lambda \in \mathbb{N}$. Denote by $\ell_a(\lambda)$ the 'last' layer of contacts an agent $a \in \mathsf{A}$ can add to her network without going above the maximum group size λ, i.e.,*

$$\ell_a(\lambda) := \max\{n \in \mathbb{N} \cup \{-1\} : |\bigcup_{k=-1}^{n} \mathrm{A}_a(k)| \leq \lambda\}$$

The bounded similarity update *on M produces the SNM $M_{\bowtie_\lambda} = \langle \mathsf{A}, S_{\bowtie_\lambda}, V\rangle$, with its social relation given by*

$$S_{\bowtie_\lambda} := \{(a, b) \in \mathsf{A} \times \mathsf{A} : b \in \bigcup_{k=-1}^{\ell_a(\lambda)} \mathrm{A}_a(k)\}$$

Since layers might have more than one element, each agent could reach a point where she should decide whether to add the next layer of friends and go above the limit λ, or stop and stay strictly below it. The definition provided above chooses the second possibility: agents will always stay below the limit, even if that means leaving some 'memory slots' empty. The extra empty layer $\mathrm{A}_a(-1)$ makes this definition work in cases in which the first layer $\mathrm{A}_a(0)$ contains already too many agents. In such situations, $\ell(\lambda) = -1$ and hence $\bigcup_{k=-1}^{\ell(\lambda)} \mathrm{A}_a(k) = \mathrm{A}_a(-1) = \varnothing$; thus, after the bounded similarity update operation, the agent will be friendless.

Properties and variations. The social network created by the threshold approach of [1] is reflexive (hence serial) and symmetric, though it might not be neither transitive nor Euclidean. In contrast, a social network created by the group-size bounded similarity update does not guarantee any of such properties. First, for reflexivity,

[3] In such case, and if no additional criteria is used to distinguish agents in the same layer, all of them should 'stand together': the decision of whether they will become part of a's social network should be of a 'either all or else none' nature.

Proposition 3.1. *Let $M = \langle \mathsf{A}, S, V\rangle$ be a* SNM *and $a \in \mathsf{A}$ be an agent; take $M_{\bowtie_\lambda} = \langle \mathsf{A}, S_{\bowtie_\lambda}, V\rangle$ (Definition 3.2). Then, a considers herself as part of her new social network ($S_{\bowtie_\lambda} aa$) if and only if the amount of people that are feature-wise identical to her is at most the limit λ ($|\mathrm{A}_a(0)| \leq \lambda$).*

Note how $|\mathrm{A}_a(0)| > \lambda$ implies not only that $S_{\bowtie_\lambda} aa$ will fail, but also that $S_{\bowtie_\lambda}[a] = \varnothing$ (so a will be friendless after the operation).

For symmetry, transitivity and Euclideanity,

Fact 3.1. Let $M = \langle \mathsf{A}, S, V\rangle$ be a SNM; take $M_{\bowtie_\lambda} = \langle \mathsf{A}, S_{\bowtie_\lambda}, V\rangle$. Then, $S_{\bowtie_\lambda}$ might not be neither symmetric, nor transitive nor Euclidean.

Proof. Here are counterexamples to each one of these properties.

- Symmetry fails for a and b if, despite a having 'enough space' for b, there is some c that is both closer to b than a ($\text{DIST}_M(b, c) \leq \text{DIST}_M(b, a)$, so b would pick c over a), and farther away from a than b ($\text{DIST}_M(a, b) \leq \text{DIST}_M(a, c)$, so a would chose b over c). By taking $\lambda = 2$, the SNM below on the left shows such situation, with the SNM on the right being the result of the update.[4]

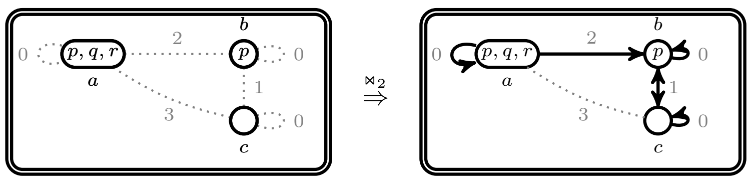

  More generally, symmetry fails if a high occurrence of similar agents produces a fully connected cluster, leaving dissimilar ones with asymmetric edges.
- The failure of transitivity also relies on b being close enough to a (so $S_{\bowtie_\lambda} ab$ holds) and c being both close enough to b (so $S_{\bowtie_\lambda} bc$ holds) and further away from a 'in b's direction' (so $S_{\bowtie_\lambda} ac$ fails). The models above showing the failure of symmetry also show how transitivity might fail.
- Finally, the relation is not Euclidean if, even though b_1 and b_2 are both close enough to a for the latter to call them her friends, they are different enough from each other to allow somebody else to take their supposed place by being more similar to each one of them (while also being very different from a). Such slightly convoluted situations are described better graphically, and the SNM below on the left is an example (take $\lambda = 3$).

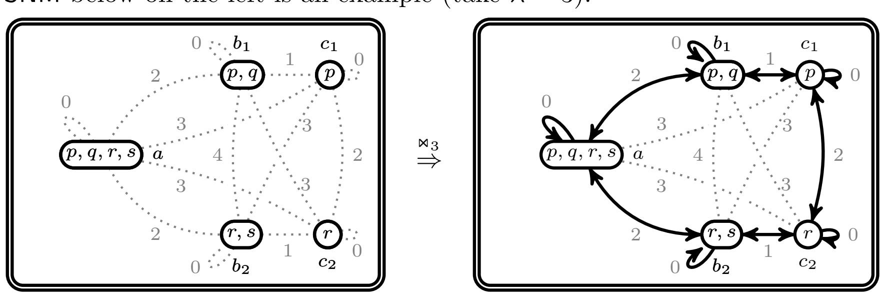

[4] Numbers over edges indicate distance. Edges in black are actual pairs in the social network relation, and dotted grey edges are shown only for distance information.

Characterising those situations in which the group-size approach produces symmetric, transitive or Euclidean social networks is not straightforward. Obviously, a λ larger or equal than $|\mathsf{A}|$ will produce fully connected (hence symmetric, transitive and Euclidean) relations; still, these properties might be achieved under other circumstances. For example, symmetry can be achieved also when the agents are 'similarly dissimilar' (i.e., their differences are 'uniformly distributed'), as the update might yield ring-like structures with symmetric edges (see the above 'Euclideanity' counterexample).

These results might suggest that the networks created by Definition 3.2 are relatively 'arbitrary': compared with the threshold approach of [1], which guarantees reflexivity and symmetry, the group-size approach might seem to produce random social networks. This is actually not the case. In the threshold approach, what matters for deciding whether b will become part of a's social network (besides the threshold itself) is only the distance between a and b. However, in the group-size approach, what matters for deciding whether b will become part of a's social network (besides the group-size itself) is the distance between a and *all agents*. Indeed, the distance between a and b is, by itself, not enough: it is possible for a and b to be extremely similar (say, $\textsc{dist}_M(a,b) = 1$), and still b will not be in a's social network if the number of agents feature-wise identical to a is high. Even more: a and b might be feature-wise identical, and yet they will not be socially connected if the number of agents feature-wise identical to them is larger than the group-size.

The group-size approach is context-sensitive: the new social networks are built not in terms of how fit is each candidate individually, but rather on how fit is each candidate *compared with the rest*. In other words, it is not about similarity, but rather about *relative* similarity. A detailed study of the group-conditions that guarantee the social network will have specific properties is left for future work.

Still, the operation might be defined in slightly alternative ways. As mentioned before, in some cases the agent will have empty 'memory slots' because the next layer would have put her social network above the size limit. One could make it possible for an agent to take on exactly λ contacts by asking for additional criteria to distinguish agents in the same layer (e.g., in an appropriate setting, considering not only the agents' features but also their preferences/beliefs). Still, one can also assume that λ is a loose limit, allowing the agent to go above it when she cannot tell the members of a group apart. For this, a small change in the definition of the upper limit ℓ_a (Definition 3.2) is enough:

$$\ell_a(\lambda) := \max\{n+1 \in \mathbb{N} : |\bigcup_{k=-1}^{n} \mathrm{A}_a(k)| \leq \lambda\}.$$

Readers interested in irreflexive friendship relations (as those in [4,5,3,6]) can achieve this property by defining agent a's sequence of layers not in terms of the full set of agents A, but rather in terms of $\mathsf{A}^{-a} := \mathsf{A} \setminus \{a\}$.

Finally, a further variation is to allow for each agent to have a *personal group-size* [11,12]. This can be represented by a function $\Lambda : \mathsf{A} \to \mathbb{N}$ indicating how many friends each agent can handle, which can be then used to define the new updated relation as $S_{\bowtie_\lambda} := \{(a,b) \in \mathsf{A} \times \mathsf{A} \,:\, b \in \bigcup_{k=-1}^{\ell_a(\Lambda(a))} \mathrm{A}_a(k)\}$, the only difference being the use of $\Lambda(a)$ instead of λ when defining the agents who will join a's social group.

Dynamic Language $\mathcal{L}_{\bowtie_\lambda}$. To express how the bounded update changes a social network, we define the language $\mathcal{L}_{\bowtie_\lambda}$.

Definition 3.3 (Language $\mathcal{L}_{\bowtie_\lambda}$). *The language $\mathcal{L}_{\bowtie_\lambda}$ extends $\mathcal{L}$ with a modality $[\bowtie_\lambda]$ to build formulas of the form $[\bowtie_\lambda]\,\varphi$ (* "after a bounded similarity update, φ is the case" *). The semantic interpretation of this modality refers to the bounded similarity updated model of Definition 3.2 as follows. Let M be a* SNM*; then,*

$$M \Vdash [\bowtie_\lambda]\,\varphi \quad \textit{iff}_{def} \quad M_{\bowtie_\lambda} \Vdash \varphi.$$

Note that no precondition is required for a bounded similarity update. Because of this and the functionality of the model operation, the dual modality $\langle \bowtie_\lambda \rangle\,\varphi := \neg\,[\bowtie_\lambda]\,\neg\varphi$ is such that $\Vdash [\bowtie_\lambda]\,\varphi \leftrightarrow \langle \bowtie_\lambda \rangle\,\varphi$.

The axiom system characterising validities of $\mathcal{L}_{\bowtie_\lambda}$ in SNM is built via the *DEL* technique of *recursion axioms*. As such, it makes crucial use of the fact that the basic 'static' language $\mathcal{L}$ is already expressive enough to characterise the changes that the bounded similarity update operation brings about. The crucial axiom, the one characterising the way in which the social network relation changes, will be built up step by step.

First, note that, when P is *finite*, the following $\mathcal{L}$-formula is true in a model M if and only if agents a and b differ in exactly $t \in \mathbb{N}$ features:

$$\mathrm{Dist}^t_{a \cdot b} := \bigvee_{\{\mathsf{P}' \subseteq \mathsf{P} : |\mathsf{P}'| = t\}} \Big(\bigwedge_{p \in \mathsf{P}'} (p_a \veebar p_b) \wedge \bigwedge_{p \in \mathsf{P} \setminus \mathsf{P}'} (p_a \leftrightarrow p_b) \Big) \qquad {}^{5}$$

The second step consists in defining the $\mathcal{L}$-formula $\mathrm{Closer}_{a \cdot b_1 \cdot b_2}$, which is true in a model M if and only if agent b_2 is at most as close to agent a as agent b_1 (i.e., $\textsc{dist}_M(a, b_1) \leq \textsc{dist}_M(a, b_2)$):

$$\mathrm{Closer}_{a \cdot b_1 \cdot b_2} := \bigvee_{j_1=0}^{|\mathsf{P}|} \bigvee_{j_2=j_1}^{|\mathsf{P}|} \Big(\mathrm{Dist}^{j_1}_{a \cdot b_1} \wedge \mathrm{Dist}^{j_2}_{a \cdot b_2} \Big) \qquad {}^{6}$$

[5] More precisely, the formula states that there is at least one set of features P', of size t, such that a and b differ in all features in P' and coincide in all features in $\mathsf{P} \setminus \mathsf{P}'$. There can be a most one such set; therefore the formula is true exactly when a and b differ in exactly t features.

[6] More precisely, the formula states that there are $j_1, j_2 \in \{0, \ldots, |\mathsf{P}|\}$, with $j_1 \leq j_2$, such that j_1 is the distance from a to b_1, and j_2 is the distance from a to b_2.

By using the $\text{Closer}_{a\cdot b_1\cdot b_2}$ formula, and in those cases in which A is finite, it is possible to provide further $\mathcal{L}$-formulas characterising the agents in each one of the layers induced by the qualitative *'distance from a'* relation $\preccurlyeq_a^M$: for $n \geq 0$,

$$\text{InLay}_{a,-1}(b) := \bot, \qquad \text{InLay}_{a,0}(b) := \bigwedge_{p\in\mathsf{P}} (p_a \leftrightarrow p_b),$$

$$\text{InLay}_{a,n+1}(b) := \bigwedge_{k=0}^{n} \neg\,\text{InLay}_{a,k}(b) \ \wedge \ \bigwedge_{c\in\mathsf{A}} \left(\bigwedge_{k=0}^{n} \neg\,\text{InLay}_{a,k}(c) \ \rightarrow \ \text{Closer}_{a\cdot b\cdot c} \right).$$

It is not hard to see that each formula $\text{InLay}_{a,k}(b)$ indeed characterises each layer $\text{A}_a(k)$, i.e., for every SNM M, $a \in \mathsf{A}$ and $k \in \mathbb{N} \cup \{-1\}$,

$$\text{A}_a(k) = \{b \in \mathsf{A} \,:\, M \Vdash \text{InLay}_{a,k}(b)\}$$

The cases for $k = -1$ and $k = 0$ are straightforward: $\text{A}_a(-1)$ is always empty, and $\text{A}_a(0)$ always contains those agents that are feature-wise identical to agent a. The remaining (inductive) case is also straightforward, as a given agent b is in $\text{A}_a(n+1)$ (formula: $\text{InLay}_{a,n+1}(b)$) if and only if it is not in any 'lower' layer (formula: $\bigwedge_{k=0}^{n} \neg\,\text{InLay}_{a,k}(b)$) and every agent that is not in a 'lower' layer is at most as close to a than b herself (formula: $\bigwedge_{c\in\mathsf{A}} \left(\bigwedge_{k=0}^{n} \neg\,\text{InLay}_{a,k}(c) \ \rightarrow \ \text{Closer}_{a\cdot b\cdot c}\right)$).

Finally, given Definition 3.2, it follows that the following $\mathcal{L}_{\bowtie_\lambda}$-validity characterizes the way the social relation changes:

$$\Vdash [\bowtie_\lambda]\, \text{S}_{ab} \ \leftrightarrow \ \bigvee_{k=0}^{\ell_a(\lambda)} \text{InLay}_{a,k}(b)$$

In words, after a bounded similarity update agent a will have agent b in her social network, $[\bowtie_\lambda]\, \text{S}_{ab}$, if and only if, before the update, agent b was in some of the layers whose agents will be part of a's social network, $\bigvee_{k=0}^{\ell_a(\lambda)} \text{InLay}_{a,k}(b)$.

As only the social relation changes in the new model, we have the following.

Theorem 3.1. *The reduction axioms and the rule on Table 1 provide, together with a propositional axiom system schema, a sound and strongly complete axiom system characterising the validities of the dynamic language $\mathcal{L}_{\bowtie_\lambda}$ (for a finite set of features and a finite set of agents).*

If the relation $S_{\bowtie_\lambda}$ is forced to be irreflexive following the suggestion above, it is enough to restrict the current axiom characterizing the new social relation to cases in which a and b are different agents, and then add an additional axiom expressing that S_{aa} is never the case after the update operation.

$$\vdash [\bowtie_\lambda]\, \text{S}_{ab} \ \leftrightarrow \ \bigvee_{k=0}^{\ell_a(\lambda)} \text{InLay}_{a,k}(b) \quad \text{for } a \neq b, \qquad\qquad \vdash [\bowtie_\lambda]\, \text{S}_{aa} \leftrightarrow \bot$$

For the variation in which the new social relation relies on personal group-size restrictions, the only needed change is the upper limit of the social network axiom: instead of the 'general' $\ell_a(\lambda)$, the 'personal' $\ell_a(\Lambda(a))$ should be used.

Table 1. Axiom system for $\mathcal{L}_{\bowtie_\lambda}$ over social network models.

$\vdash [\bowtie_\lambda]\, p_a \leftrightarrow p_a$	for $a \in \mathsf{A}$	From $\vdash \varphi$ infer $\vdash [\bowtie_\lambda]\, \varphi$
$\vdash [\bowtie_\lambda]\, \mathrm{S}_{ab} \;\leftrightarrow\; \bigvee_{k=0}^{\ell_a(\lambda)} \mathrm{InLay}_{a,k}(b)$	for $a, b \in \mathsf{A}$	From $\vdash \psi_1 \leftrightarrow \psi_2$ infer $\vdash \varphi \leftrightarrow \varphi\,[\psi_2/\psi_1]$ (with $\varphi\,[\psi_2/\psi_1]$ any formula obtained by replacing one or more occurrences of ψ_1 in φ with ψ_2).
$\vdash [\bowtie_\lambda]\, \neg\varphi \leftrightarrow \neg\, [\bowtie_\lambda]\, \varphi$		
$\vdash [\bowtie_\lambda](\varphi \wedge \psi) \leftrightarrow ([\bowtie_\lambda]\, \varphi \wedge [\bowtie_\lambda]\, \psi)$		

4 A restriction to *relevant* features

Any social-network-creation operation, such as the threshold update of [1] or the bounded update of Definition 3.2, can be seen as a 'public conversation' where all agents 'discuss' their features. Then, as the 'conversation' continues, agents will form subgroups of people sharing prior common interests.

When looking at social network creation from this perspective, it becomes clear that not all features can be 'discussed' at once: just some of them will be relevant at each stage of the discussion. This is not a novel idea; in [13], the authors use a game theoretic setting to define the agreement and disagreement of agents on a specific feature (or issue), which yields a way for them to update the social relation of agents with respect to one specific feature at a time.

This section explores this idea within the bounded update operation of the previous section: only a subset of all features will be relevant for each update. The resulting setting will allow us to describe more realistic scenarios, such as the step-by-step interaction in real dialogues (when personal features are slowly revealed as the conversation goes on), or cases in which agents control when one of their features becomes visible to other agents (e.g. when agents choose to expose some 'private' information only in specific circumstances).

The crucial step in this generalisation is the definition of a notion of distance that is relative only to a subset of features $\mathsf{Q} \subseteq \mathsf{P}$.

Definition 4.1 (Q-Distance). *Let $M = \langle \mathsf{A}, S, V \rangle$ be a SNM, and let $\mathsf{Q} \subseteq \mathsf{P}$ be a set of features. The* Q-distance *between a and b in M (that is, the distance between a and b in M* relative to features in Q*) is given by*

$$\textsc{dist}_M^{\mathsf{Q}}(a, b) := |\textsc{msmtch}_M(a, b) \cap \mathsf{Q}|$$

Thus, $\textsc{dist}_M^{\mathsf{Q}}(a, b)$ returns the number of atoms in Q on which a and b differ. Then, while some agent b_1 might be strictly closer to agent a than another agent b_2 with respect to all features, agent b_2 might be strictly closer to a than b_1 with respect to some strict subset of them.[7]

[7] For an example, take a model with $V(a) = \{p, q, r\}$, $V(b_1) = \{q, r\}$ and $V(b_2) = \{p\}$. Then, $\textsc{dist}_M^{\{p,q,r\}}(a, b_1) = 1 < 2 = \textsc{dist}_M^{\{p,q,r\}}(a, b_2)$, but nevertheless $\textsc{dist}_M^{\{p\}}(a, b_2) = 0 < 1 = \textsc{dist}_M^{\{p\}}(a, b_1)$.

With this notion of Q-distance (still a semi-metric, as it satisfies non-negativity, symmetry, reflexivity and subadditivity, but not a metric, as it fails to satisfy the identity of indiscernibles), one can define 'relative to Q' variants of the qualitative *'distance from a'* relation and the sequence of layers it induces.

Definition 4.2. *Given a social network model $M = \langle \mathsf{A}, S, V\rangle$, an agent $a \in \mathsf{A}$ and a subset of features $\mathsf{Q} \subseteq \mathsf{P}$, the quantitative notion of Q-distance DIST^{Q} induces a qualitative (total, reflexive, transitive and well-founded) relation $\preccurlyeq_a^{\mathsf{Q},M} \subseteq \mathsf{A} \times \mathsf{A}$ of* Q-distance from a. *Such a relation is given by*

$$\preccurlyeq_a^{\mathsf{Q},M} := \{(b_1, b_2) \in \mathsf{A} \times \mathsf{A} : \text{DIST}_M^{\mathsf{Q}}(a, b_1) \leq \text{DIST}_M^{\mathsf{Q}}(a, b_2)\}$$

and thus $b_1 \preccurlyeq_a^{\mathsf{Q},M} b_2$ indicates that, with respect to the features in Q, agent b_1 is at least as close to agent a as agent b_2 in model M. By defining the notion of $\preccurlyeq_a^{\mathsf{Q},M}$-minimum in the standard way (for $\mathsf{B} \subseteq \mathsf{A}$, take $\text{MIN}_a^{\mathsf{Q}}(\mathsf{B}) := \{b \in \mathsf{B} : b \preccurlyeq_a^{\mathsf{Q},M} b'$ for all $b' \in \mathsf{B}\}$), this relation induces the following sequence of layers on A, each one containing agents equally distant from a:

$$\mathrm{A}_{-1}^{\mathsf{Q}}(a) := \varnothing, \qquad \mathrm{A}_0^{\mathsf{Q}}(a) := \text{MIN}_a^{\mathsf{Q}}(\mathsf{A}), \qquad \mathrm{A}_{n+1}^{\mathsf{Q}}(a) := \text{MIN}_a^{\mathsf{Q}}(\mathsf{A} \setminus \bigcup_{k=-1}^{n} \mathrm{A}_k^{\mathsf{Q}}(a)).$$

Similarly we now restrict the definition of the bounded similarity update to a version that is set to be relative to a given subset of features Q.

Definition 4.3 (Bounded Q-similarity update). *Let $M = \langle \mathsf{A}, R, V\rangle$ be a SNM and $\mathsf{Q} \subseteq \mathsf{P}$ a subset of features; take $\lambda \in \mathbb{N}$. Denote by $\ell_a^{\mathsf{Q}}(\lambda)$ the 'last' layer of contacts an agent $a \in \mathsf{A}$ can add to her network without going above the maximum group size λ, i.e.,*

$$\ell_a^{\mathsf{Q}}(\lambda) := \max\{n \in \mathbb{N} \cup \{-1\} : |\bigcup_{k=-1}^{n} \mathrm{A}_k^{\mathsf{Q}}(a)| \leq \lambda\}$$

The bounded Q-similarity update *on M produces the SNM $M_{\bowtie_\lambda^{\mathsf{Q}}} = \langle \mathsf{A}, S_{\bowtie_\lambda^{\mathsf{Q}}}, V\rangle$, with its social relation given by*

$$S_{\bowtie_\lambda^{\mathsf{Q}}} := \{(a, b) \in \mathsf{A} \times \mathsf{A} : b \in \bigcup_{k=-1}^{\ell_a^{\mathsf{Q}}(\lambda)} \mathrm{A}_k^{\mathsf{Q}}(a)\}$$

Example 4.1. Consider the SNM models of the counterexample for Euclideanity (Fact 3.1 on page 5), drawn again below.

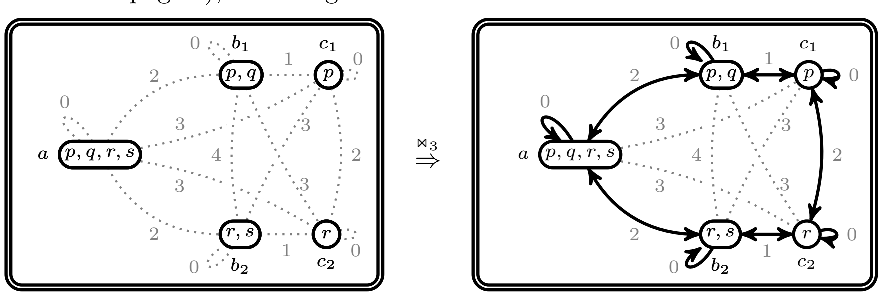

The SNM above on the left shows the resulting social network when all features $\{p, q, r, s\}$ are 'put on the table'. But suppose that this is not the case; then, the operation produces different social networks. For example, if the agents only 'talk' about features in $\{p, q\}$, then the distances are as shown in the model below on the left (underlined numbers emphasising distances that differ from the original 'fully open' situation):

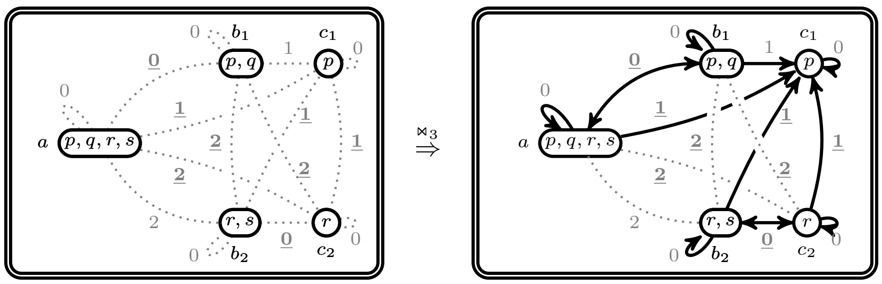

The resulting SNM is shown above on the right. As expected, the social network relation is different. Note, in particular, how while the relation is still reflexive, it is not symmetric anymore; however, it is now transitive. Also interesting is the fact that, although c_1 is considered 'a friend' by everybody, she is the only member of her own social network: all other agents are at a distance of 1, and thus adding all of them would have taken her above the limit $\lambda = 3$. In fact, by restricting the conversation to the issues in $\{p, q\}$, the resulting network can be seen as three fully connected clusters, $\{a, b_1\}$, $\{b_2, c_2\}$ and $\{c_1\}$, with the members of the firsts pointing asymmetrically to the lone member of the last.

Dynamic Language $\mathcal{L}_{\bowtie^{\mathsf{Q}}_{\lambda}}$. The language $\mathcal{L}_{\bowtie^{\mathsf{Q}}_{\lambda}}$ is similar to $\mathcal{L}_{\bowtie_{\lambda}}$; the only difference lies in the semantic interpretation of its 'dynamic' operator, $\bowtie^{\mathsf{Q}}_{\lambda}$.

Definition 4.4 (Language $\mathcal{L}_{\bowtie_{\lambda}}$). *The language $\mathcal{L}_{\bowtie^{\mathsf{Q}}_{\lambda}}$ extends $\mathcal{L}$ with a modality $[\bowtie^{\mathsf{Q}}_{\lambda}]$ to build formulas of the form $[\bowtie^{\mathsf{Q}}_{\lambda}]\varphi$ (*"after a bounded Q-similarity update, φ is the case"*). The semantic interpretation of this modality refers to the relativised bounded updated model of Definition 4.3 as follows. Let M be a SNM; then,*

$$M \Vdash [\bowtie^{\mathsf{Q}}_{\lambda}]\varphi \quad \textit{iff}_{def} \quad M_{\bowtie^{\mathsf{Q}}_{\lambda}} \Vdash \varphi.$$

With respect to an axiom system, the system presented in Table 1 can be used almost verbatim. The only change refers to the axiom characterising the new social network relation, which should now be relativised to the subset of features Q; for this, it is enough to replace P with Q in the definitions for formulas $\mathrm{Dist}^{t}_{a \cdot b}$ and $\mathrm{Closer}_{a \cdot b_1 \cdot b_2}$ (page 7), thus obtaining formulas $\mathrm{Dist}^{\mathsf{Q},t}_{a \cdot b}$, $\mathrm{Closer}^{\mathsf{Q}}_{a \cdot b_1 \cdot b_2}$ and $\mathrm{InLay}^{\mathsf{Q}}_{a,k}(b)$.

5 Conclusions and future work

Following the cognitive science literature (in particular, [2]), we have examined social networks (Section 2) by studying a *group-size* approach to social network creation based on the initial idea (see Section 3) in which *"agent a will consider agent b to be part of her social network if and only if b is within the λ closest agents to a"*. This proposal can be seen as an alternative to the approach of [1], which uses a *threshold* to establish how similar an agent should be to be incorporated to someone's social environment (i.e. *"agent a will consider agent b to be part of her social network if and only if b's distance from a is at most θ"*). Moreover, the relativised version studied in Section 4 allows for the representation of more refined scenarios where not all features play a role during the agents' interaction.

The work presented here and in [1] form the initial steps in the study of the logical structure behind social network creation, and they already suggest interesting alternatives. While both the threshold and the group-size approaches relate agents when they are similar enough in their features, behavior, etc, one can think of an alternative scenario in which one considers the dual situation so that agents connect when they *complement each other*. In order to deal formally with this *complementary* idea, a more fine-grained setting is needed that takes into account not only the agents' features/behaviors, as in this paper, but also their doxastic state and their preferences (e.g., [14,15]).

Another straightforward generalization would be to consider not a single social network, but rather a collection of them. A slightly more realistic approach in this direction is to understand each feature not as a simple choice between *"yes"* and *"no"*, but rather as a choice among a finite range of values. Then the model can support a social network for each feature $p \in \mathsf{P}$, and agents can be grouped according to the value they assign to each such p. After all, someone who chooses football as her favourite sport and Lady Gaga as her favourite musician is bound to have different social environments in each one of these contexts.

A further route will lead us into a combined social network and epistemic study. This is another natural next step, as what matters most when establishing friendship is maybe not the agents' features and differences, but rather what one knows about them. Our work in [1] provides an initial exploration in this direction, using the threshold update approach.

In a related track, one can explore cases in which certain features are taken to be more important than others in such a way that this 'priority ordering' among features differs from agent to agent. This allows for the representation of interesting situations: the number of differences between agents a and b_1 might be very large, and yet they may agree on the feature a cares about the most. Then, a might consider that b_1 is 'closer to her' than some b_2 with whom she shares all but this most important feature. The combination of this (the explored setting in which the update is relative to only a subset of features) and an epistemic setting would allow us to describe situations where strategic behaviour plays an important role. For example, if agent a knows that she and agent b differ in some feature $p \in \mathsf{P}$, and she also knows that p is the most important feature for b, then she (a) might want to keep this topic out of the conversation, at least until it has been commonly established (i.e., it is common knowledge between a and b) that they are similar with respect to several other features. The setting becomes even more interesting when the new social network is defined not only in terms of the agents' similarities, but also in terms of existing social connections (cf. the middleman cases in [1]). In such cases, the features discussed at the beginning will define the social connections that will be available in further stages.

Finally, we observe the importance of the interplay between the social network changing operations of this proposal, and the operations that change the features (or behaviour/beliefs) in [3]. Both ideas deserve to be studied in tandem, as indeed the dynamics studied in one can affect the dynamics studied in the other.

References

1. Smets, S., Velázquez-Quesada, F.R.: How to make friends: A logical approach to social group creation. To appear in A. Baltag and J. Seligman, editors, Proceedings of Sixth International Workshop LORI 2017 (2017)
2. Dunbar, R.I.M.: Neocortex size as a constraint on group size in primates. Journal of Human Evolution **22**(6) (1992) 469–493
3. Baltag, A., Christoff, Z., Rendsvig, R.K., Smets, S.: Dynamic epistemic logics of diffusion and prediction in social networks (extended abstract). In Bonanno, G., van der Hoek, W., Perea, A., eds.: Proceedings of LOFT 2016. (2016)
4. Seligman, J., Liu, F., Girard, P.: Logic in the community. In Banerjee, M., Seth, A., eds.: Logic and Its Applications - 4th Indian Conference, ICLA 2011, Delhi, India, January 5-11, 2011. Proceedings. Volume 6521 of Lecture Notes in Computer Science., Springer (2011) 178–188
5. Liu, F., Seligman, J., Girard, P.: Logical dynamics of belief change in the community. Synthese **191**(11) (2014) 2403–2431
6. Christoff, Z., Hansen, J.U., Proietti, C.: Reflecting on social influence in networks. Journal of Logic, Language and Information **25**(3-4) (2016) 299–333
7. Deza, M.M., Deza, E.: Encyclopedia of Distances. Springer (2009)
8. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In Gilboa, I., ed.: Proceedings of TARK-98. (1998) 43–56
9. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Volume 337 of Synthese Library Series. Springer, Dordrecht, The Netherlands (2008)
10. van Benthem, J.: Logical Dynamics of Information and Interaction. Cambridge University Press (2011)
11. Roberts, S.G.B., Wilson, R., Fedurek, P., Dunbar, R.I.M.: Individual differences and personal social network size and structure. Personality and Individual Differences **44**(4) (2008) 954–964
12. Roberts, S.G.B., Dunbar, R.I.M., Pollet, T.V., Kuppens, T.: Exploring variation in active network size: Constraints and ego characteristics. Social Networks **31**(2) (2009) 138–146
13. Solaki, A., Terzopoulou, Z., Zhao, B.: Logic of closeness revision: Challenging relations in social networks. In Köllner, M., Ziai, R., eds.: Proceedings of the ESSLLI 2016 student session. (2016) 123–134

14. Baltag, A., Smets, S.: A qualitative theory of dynamic interactive belief revision. In Bonanno, G., van der Hoek, W., Wooldridge, M., eds.: Logic and the Foundations of Game and Decision Theory (LOFT7). Volume 3 of Texts in Logic and Games. Amsterdam University Press, Amsterdam, The Netherlands (2008) 13–60
15. Ghosh, S., Velázquez-Quesada, F.R.: Agreeing to agree: Reaching unanimity via preference dynamics based on reliable agents. In Weiss, G., Yolum, P., Bordini, R.H., Elkind, E., eds.: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015, ACM (2015) 1491–1499