Multi-Agent Topological Evidence Logics*

Alexandru Baltag[1], Nick Bezhanishvili[1], and Saúl Fernández González[2]

[1] ILLC, Universiteit van Amsterdam
[2] IRIT, Université de Toulouse

1 Introduction

In [BBÖS16] a topological semantics for evidence-based belief and knowledge is introduced, where epistemic sentences are built in a language $\mathcal{L}_{\forall KB\Box\Box_0}$, which includes modalities allowing us to talk about defeasible knowledge (K), infallible knowledge ($[\forall]$), belief (B), basic evidence ($\Box_0$) and combined evidence ($\Box$).

Definition 1 (The dense interior semantics)**.** Sentences of $\mathcal{L}_{\forall KB\Box\Box_0}$ are read on *topological evidence models* (topo-e-models), which are tuples (X, τ, E_0, V) where (X, τ) is a topological space, E_0 is a subbasis of τ and $V : \mathsf{Prop} \to 2^X$ is a valuation.

The semantics of a formula ϕ is as follows: $\|p\| = V(p)$; $\|\phi \wedge \psi\| = \|\phi\| \cap \|\psi\|$; $\|\neg\phi\| = X \setminus \|\phi\|$; $\|\Box\phi\| = \mathrm{Int}\,\|\phi\|$; $x \in \|K\phi\|$ iff $x \in \mathrm{Int}\,\|\phi\|$ and $\mathrm{Int}\,\|\phi\|$ is dense[1]; $x \in \|B\phi\|$ iff $\mathrm{Int}\,\|\phi\|$ is dense; $x \in \|[\forall]\phi\|$ iff $\|\phi\| = X$; $x \in \|\Box_0\phi\|$ iff there is $e \in E_0$ with $x \in e \subseteq \|\phi\|$; $x \in \|\Box\phi\|$ iff $x \in \mathrm{Int}\,\|\phi\|$.

Crucially, using topological spaces to model epistemic sentences grants us an *evidential* perspective of knowledge and belief. Indeed, we can see the opens in the topology as the pieces of evidence the agent has (and thus our modality $\Box$, which encodes "having evidence", becomes the topological interior operator). For some proposition ϕ to constitute (defeasible) knowledge, we demand that the agent has a factive *justification* for ϕ, i.e. a piece of evidence that cannot be contradicted by any other evidence the agent has. In topological terms, a justification amounts to a *dense* piece of evidence. Having a (not necessarily factive) justification constitues belief. The set X encodes all the possible worlds which are consistent with the agent's information, thus for the agent to know ϕ infallibly ($[\forall]\phi$), ϕ needs to hold throughout X.

The fragment of this language that only contains the Booleans and the K modality, $\mathcal{L}_K$, has $\mathsf{S4.2}$ as its logic.

The framework introduced in [BBÖS16] is single-agent. A multi-agent generalisation is presented in this text, along with some "generic models" and a notion of group knowledge. Our proposal differs conceptually from previous multi-agent approaches to the dense interior semantics [Ö17, Ram15].

2 Going Multi-Agent

For clarity of presentation we work in a two-agent system.[2] Our language now contains modalities $K_i, B_i, [\forall]_i, \Box_i, \Box_i^0$ for $i = 1, 2$, each encoding the same notion as in the single-agent system.

*This paper compiles the results contained in Chapters 3 to 5 of Saúl Fernández González's Master's thesis [FG18]. The authors wish to thank Guram Bezhanishvili for his input.

[1] A set $U \subseteq X$ is dense whenever $\mathrm{Cl}\,U = X$, or equivalently when it has nonempty intersection with every nonempty open set.

[2] Extending these results to $n \geq 2$ agents is straightforward, see [FG18, Section 6.1].

The Problem of Density. The first issue one comes across when defining a multi-agent semantics is that of accounting for the notion of defeasibility, which, as we have seen, is closely tied to density. A first (naive) approach would be to consider two topologies and a valuation defined on a common space, (X, τ_1, τ_2, V) and simply have: $x \in \|K_i\phi\|$ iff there exists some τ_i-dense open set such that $x \in U \subseteq \|\phi\|$. This does not work, neither conceptually (for we are assuming that the set of worlds compatible with each agent's information is the same for both agents) nor logically (adopting this semantics gives us highly undesirable theorems such as $\neg K_1 \neg K_1 p \to K_2 \neg K_1 \neg K_1 p$). Seeing as each agent's knowledge is an $\mathsf{S4.2}$ modality and no interaction between the agents is being assumed, one would expect the two-agent logic to simply combine the $\mathsf{S4.2}$ axioms for each of the agents.

Simply defining two topologies on the whole space is not the right move. Instead, we want to make explicit, at each world $x \in X$, which subsets of worlds in X are compatible with each agent's information. A straightforward way to do this is via the use of partitions.

Topological-partitional models.

Definition 2. A *topological-partitional model* is a tuple $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$ where X is a set, τ_1 and τ_2 are topologies defined on X, Π_1 and Π_2 are partitions and V is a valuation.

For $U \subseteq X$ we write $\Pi_i[U] := \{\pi \in \Pi_i : U \cap \pi \neq \varnothing\}$. For $i = 1, 2$ and $\pi \in \Pi_i[U]$ we say U is *i-locally dense in* π whenever $U \cap \pi$ is dense in the subspace topology $(\pi, \tau_i|_\pi)$; we simply say U is *i-locally dense* if it is locally dense in every $\pi \in \Pi_i[U]$.

For the remainder of this text, we limit ourselves to the fragment of the language including the K_1 and K_2 modalities.

Definition 3 (Semantics)**.** We read $x \in \|K_i\phi\|$ iff there exists an i-locally dense τ_i-open set U with $x \in U \subseteq \|\phi\|$.

This definition generalises one-agent models, appears to hold water conceptually and, moreover, gives us the logic one would expectedly extrapolate from the one-agent case.

Lemma 4. *If $(X, \leq_1, \leq_2)$ is a birelational frame where each $\leq_i$ is reflexive, transitive and weakly directed (i.e. $x \leq_i y, z$ implies there exists some $t \geq_i y, z$), then the collection τ_i of $\leq_i$-upsets and the set Π_i of $\leq_i$-connected components give us a topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2)$ in which the semantics of Def. 2 and the Kripke semantics coincide.*

Now, the Kripke logic of such frames is the fusion $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$, i.e. the least normal modal logic containing the $\mathsf{S4.2}$ axioms for each K_i. As an immediate consequence:

Corollary 5. *$\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ is the $\mathcal{L}_{K_1K_2}$-logic of topological-partitional models.*

3 Generic Models

[FG18] is partially concerned with finding *generic models* for topological evidence logics, i.e. single topological spaces whose logic (relative to a certain fragment $\mathcal{L}$) is precisely the sound and complete $\mathcal{L}$-logic of topo-e-models. Let us showcase two examples of two-agent generic models for the $\mathcal{L}_{K_1K_2}$ fragment. These are particular topological-partitional models whose logic is precisely $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

The Quaternary Tree $\mathcal{T}_{2,2}$. The *quaternary tree* $\mathcal{T}_{2,2}$ is the full infinite tree with two relations R_1 and R_2 where every node has exactly four successors: a left R_i-successor and a right R_i-successor for $i = 1, 2$. Let $\leq_i$ be the reflexive and transitive closure of R_i.

We can, as we did before, define two topologies τ_i and two partitions Π_i on $\mathcal{T}_{2,2}$ in a very natural way, namely by taking, respectively, the set of $\leq_i$-upsets and the set of $\leq_i$-connected components. And we get:

Theorem 6. $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ *is sound and complete with respect to* $(\mathcal{T}_{2,2}, \tau_{1,2}, \Pi_{1,2})$.

The completeness proof uses the fact that $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ is complete with respect to finite rooted birelational Kripke frames in which both relations are reflexive, transitive and weakly directed, plus the fact proven in [vBBtCS06] that, given a preordered birelational finite frame W, there is an onto map $f : \mathcal{T}_{2,2} \to W$ which is continuous and open in both topologies.

The result then follows immediately from:

Lemma 7. *Given an* $\mathsf{S4.2} + \mathsf{S4.2}$ *frame* W *and a map* f *as described above, plus a valuation* V *on* W*, we have that* $W, V, fx \vDash \phi$ *under the Kripke semantics if and only if* $\mathcal{T}_{2,2}, V^f, x \vDash \phi$ *under the semantics of Def. 2, where* $V^f(p) = \{x \in \mathcal{T}_{2,2} : fx \in V(p)\}$.

Proof sketch. The proof of this lemma is an induction on formulas. The right to left direction for the induction step corresponding to K_i uses the fact that, if U is a connected i-upset in W with $fx \in U$, then $U' = \{z : z \geq_i y \text{ for some } y \in [x]_{\Pi_i} \text{ with } fy \in U\}$ is an i-locally dense open set in $\mathcal{T}_{2,2}$ with $x \in U \subseteq [x]_{\Pi_i}$. ∎

The rational plane $\mathbb{Q} \times \mathbb{Q}$. We can define two topologies on $\mathbb{Q}$ by "lifting" the open sets in the rational line horizontally or vertically. Formally, the *horizontal topology* τ_H is the topology generated by $\{U \times \{y\} : U \text{ is open}, y \in \mathbb{Q}\}$. Similarly, the *vertical topology* τ_V is generated by the sets $\{y\} \times U$. We have the following result:

Proposition 8. *There exist partitions* Π_H *and* Π_V *such that* $(\mathbb{Q} \times \mathbb{Q}, \tau_{H,V}, \Pi_{H,V})$ *is a topological-partitional model whose logic is* $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

Proof sketch. It is shown in [vBBtCS06] that there exists a surjective map $g : \mathbb{Q} \times \mathbb{Q} \to \mathcal{T}_{2,2}$ which is open and continuous in both topologies. Given such a map and a valuation V on $\mathcal{T}_{2,2}$, we can define a valuation V^g on $\mathbb{Q} \times \mathbb{Q}$ as above and two equivalence relations: $x \sim_H y$ iff $[gx]_{\Pi_1} = [gy]_{\Pi_1}$, and $x \sim_V y$ iff $[gx]_{\Pi_2} = [gy]_{\Pi_2}$. As we did before, we can prove that $(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V, \Pi_H, \Pi_V), V^g, x \vDash \phi$ iff $\mathcal{T}_{2,2}, V, gx \vDash \phi$, whence completeness follows. ∎

4 Distributed Knowledge

Once a multi-agent framework is defined, the obvious next step is to account for some notion of *knowledge of the group*. We will focus on *distributed* or *implicit* knowledge, i.e., a modality that accounts for that which the group of agents knows implicitly, or what would become known if the agents were to share their information.

One way to do this is to follow the evidence-based spirit inherent to the dense interior semantics. On this account, we would code distributed knowledge as the knowledge modality which corresponds to a fictional agent who has all the pieces of evidence the agents have (we can code this via the *join* topology $\tau_1 \vee \tau_2$, which is the smallest topology containing τ_1 and τ_2), and only considers a world compatible with x when all agents in the group do (the partition of this agent being $\{\pi_1 \cap \pi_2 : \pi_i \in \Pi_i\}$). Coding distributed knowledge like this gives us some rather strange results: unlike more standard notions, it can obtain that an agent knows a proposition but, due to the density condition on this new topology, the group does not (for an example, see [FG18, Example 5.2.3]).

Our proposal differs from this. Here we follow [HM92] when they refer to this notion as "that which a fictitious 'wise man' (one who knows exactly which each individual agent knows) would know". Instead of conglomerating the evidence of all the agents, we account exclusively for what they know, and we treat this information as indefeasible. Thus, our account of distributed knowledge, which is not strictly evidence-based, interacts with the K_i modalities in a more standard way, much like in relational semantics.

Definition 9 (Semantics for distributed knowledge)**.** Our language includes the operators K_1, K_2 and an operator D for distributed knowledge. In a topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$, we read $x \in \|D\phi\|$ iff for $i = 1, 2$ there exist i-locally dense sets $U_i \in \tau_i$ such that $x \in U_1 \cap U_2 \subseteq \|\phi\|$.

That is to say, ϕ constitutes distributed knowledge whenever the agents have indefeasible pieces of evidence which, when put together, entail ϕ.

As mentioned above, the logic of distributed knowledge is unsurprising:

Definition 10. $\mathsf{Logic}_{K_1K_2D}$ is the least set of formulas containing the $\mathsf{S4.2}$ axioms and rules for K_1 and K_2, the $\mathsf{S4}$ axioms and rules for D plus the axiom $K_i\phi \to D\phi$ for $i = 1, 2$.

Theorem 11. $\mathsf{Logic}_{K_1K_2D}$ *is sound and complete with respect to topological-partitional models.*

References

[BBÖS16] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. Justified belief and the topology of evidence. In *International Workshop on Logic, Language, Information, and Computation*, pages 83–103. Springer, 2016.

[FG18] Saúl Fernández González. Generic models for topological evidence logics. Master's thesis, University of Amsterdam, ILLC, 2018.

[HM92] Joseph Y Halpern and Yoram Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence*, 54(3):319–379, 1992.

[Ö17] Aybüke Özgün. *Evidence in Epistemic Logic: A Topological Perspective.* PhD thesis, University of Amsterdam, ILLC, 2017.

[Ram15] Aldo Ramírez. Topological models for group knowledge and belief. Master's thesis, University of Amsterdam, 2015.

[vBBtCS06] Johan van Benthem, Guram Bezhanishvili, Balder ten Cate, and Darko Sarenac. Multimodal logics of products of topologies. *Studia Logica*, 84(3):369–392, 2006.