Topological Evidence Logics: Multi-Agent Setting

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 interpreted 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 interior and closure operators of (X, τ) are denoted by Int and Cl, respectively.

The semantics of a formula ϕ is defined as follows: $\|p\| = V(p)$; $\|\phi \wedge \psi\| = \|\phi\| \cap \|\psi\|$; $\|\neg\phi\| = X \backslash \|\phi\|$; $\|\Box\phi\| = \mathrm{Int}\,\|\phi\|$; $x \in \|K\phi\|$ 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 $x \in \mathrm{Int}\,\|\phi\|$.

Modelling epistemic sentences via topological spaces 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 *dense* piece of evidence. Having a (not necessarily factive) justification constitutes 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 framework introduced in [BBÖS16] is single-agent. The fragment of this language that only contains the Booleans and the K modality, $\mathcal{L}_K$, has $\mathsf{S4.2}$ as its logic. For the analogue of the McKinsey and Tarski theorem in this single-agent setting we refer to the recent [BBFG19]. 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 in that we build it on the notion of *local density.*

2 Going Multi-Agent

For clarity of presentation we work in a two-agent system. We limit ourselves to the fragment of the language with K_i modalities for $i = 1, 2$, encoding the same notion as K in the single-agent system.

[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.

The Problem of Density. A first (naive) approach to the notion of defeasibility in the multi-agent setting (which, as we have seen, is closely tied to density) 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 approach is not promising, neither conceptually (it assumes the same set of worlds is compatible with both agents' information) nor logically (since 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; however the corresponding logic for this semantics contains undesirable theorems such as $\neg K_1\neg K_1 p \to K_2\neg K_1\neg K_1 p$).

We can fix this by making explicit, at each world $x \in X$ and for each agent, which is the set of worlds compatible with the agent's information at that state. A useful 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 $\pi \in \Pi_i[U]$.

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 be sound 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 directed (i.e. $x \leq_i y, z$ implies there exists $t \geq_i y, z$), then the collection τ_i of $\leq_i$-upsets and the set Π_i of R_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 we obtain:

Theorem 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

Following the spirit of the McKinsey-Tarski theorem [MT44], one of our aims is finding *generic models* for this logic, i.e. single topological-partitional models whose logic is precisely $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$. Generic models for the one-agent framework are given in [BBFG19].

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 $i = 1, 2$. We can define two topologies τ_i and two partitions Π_i on $\mathcal{T}_{2,2}$ by taking, respectively, the set of R_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 rational plane $\mathbb{Q} \times \mathbb{Q}$. We can define two topologies on $\mathbb{Q} \times \mathbb{Q}$ by "lifting" the open sets in the rational line horizontally or vertically [vBBtCS06]. Formally, the *horizontal topology* τ_H is generated by $\{U \times \{y\} : U \text{ is open}, y \in \mathbb{Q}\}$, and likewise for the vertical topology.

Theorem 7. *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* $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

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 $\tau_1 \vee \tau_2$, generated by $\tau_1 \cup \tau_2$), 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 this way gives us some rather unexpected features: unlike more standard notions, an agent may know a proposition but, due to the density condition on this new topology, the group may not.

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 8 (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.

Theorem 9. $\mathsf{Logic}_{K_1K_2D}$ *(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* $i = 1, 2$*) is sound and complete with respect to topological-partitional models.*

References

[BBFG19] A. Baltag, N. Bezhanishvili, and S. Fernández González. The McKinsey-Tarski Theorem for Topological Evidence Logics. In *Proceedings of WoLLIC 2019*, 2019. To appear.

[BBÖS16] A. Baltag, N. Bezhanishvili, A. Özgün, and S. Smets. Justified belief and the topology of evidence. In *Proceedings of WoLLIC 2016*, pages 83–103. Springer, 2016.

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

[MT44] J. C. C. McKinsey and A. Tarski. The algebra of topology. *Annals of mathematics*, pages 141–191, 1944.

[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.