A Degree-Theoretic Framework for Feasible Knowledge

LIBOR BĚHOUNEK[1]

Abstract: A semantics for epistemic logic is defined that retains the epistemic agent's logical rationality while avoiding logical omniscience. The proposed solution to the logical omniscience paradox is based on resource-awareness implemented by means of t-norm fuzzy logics.

Keywords: epistemic logic, logical omniscience, feasibility, t-norm fuzzy logic, possible-world semantics

1 The logical omniscience paradox

Standard epistemic logic (see, e.g., Meyer, 2003) renders the epistemic modality *the agent knows that* φ as a propositional modal operator K. The operator is supposed to validate some of the standard axioms of normal modal logics, which formalize several principles of epistemic reasoning:

(K)	$\vdash \mathrm{K}(\varphi \to \psi) \to (\mathrm{K}\varphi \to \mathrm{K}\psi)$	logical rationality
(T)	$\vdash \mathrm{K}\varphi \to \varphi$	truth of knowledge
(4)	$\vdash \mathrm{K}\varphi \to \mathrm{KK}\varphi$	positive introspection
(5)	$\vdash \neg\mathrm{K}\varphi \to \mathrm{K}\neg\mathrm{K}\varphi$	negative introspection
(Nec)	$\varphi \vdash \mathrm{K}\varphi$	necessitation

The axioms (4) and (5) of positive and negative introspection are often considered optional, depending on the introspective abilities of the epistemic agent. The doxastic variant with the modality *the agent believes that* φ replaces the truth axiom (T) by the consistency axiom (D):

(D)	$\vdash \mathrm{K}\neg\varphi \to \neg\mathrm{K}\varphi$	consistency of beliefs

[1]The work was supported by project No. LQ1602 "IT4I XS" in the program NPU II of the Ministry of Eductation, Youth and Sports of the Czech Republic.

Being a normal modal logic, standard epistemic (or doxastic) logic is sound and complete with respect to the usual Kripke semantics, where K is evaluated $\Box$-like and the intended meaning of the accessibility relation R is that of epistemic indistinguishability: in a possible world w, the agent only knows (or believes) that the actual world is one of the worlds w' such that wRw', but cannot distinguish which one of these.

A well known defect of standard epistemic logic is the paradox of *logical omniscience,* or the fact that by the inference rule (Nec) and the axiom schema (K), the agent automatically knows all propositional consequences of their actual knowledge, including all propositional tautologies. This is, of course, unrealistic for real-world agents.

Since the standard axioms and rules of doxastic logic likewise include (K) and (Nec), the problem of logical omniscience applies to standard doxastic logic as well. We will therefore deal with both of these variants at once, for the most part referring to epistemic logic only, but *mutatis mutandis* applying our considerations to doxastic logic too.

2 Solutions based on resource-awareness

The logical omniscience paradox is obviously caused by over-idealization of deductive abilities of epistemic agents in standard epistemic logic. In particular, the modal axioms neglect the costs of logical derivations, or the fact that the agent needs to spend some of their limited resources, such as computation time, memory, or energy, in order to perform the derivation. Consequently, an important class of solutions to the logical omniscience paradox is based on *resource-awareness,* i.e., on modifying the standard account by acknowledging that some resources need be spent on logical deductions.

Several resource-aware solutions to logical omniscience have been proposed, i.a., by Duc (1997, 2001) and Artemov and Kuznets (2014). One of these proposals is to introduce time-awareness by prepending the temporal modality $\langle \mathrm{F} \rangle$ (standing for "at some future point") to the conclusions of epistemic principles. The axioms (K) and (4) are thus modified to read:

$$\begin{aligned} &\vdash \mathrm{K}(\varphi \to \psi) \to \big(\mathrm{K}\varphi \to \langle \mathrm{F} \rangle \mathrm{K}\psi\big) \\ &\vdash \mathrm{K}\varphi \to \langle \mathrm{F} \rangle \mathrm{K}\mathrm{K}\varphi, \end{aligned}$$

meaning: "if the agent knows both φ and $\varphi \to \psi$, then *at some future point* (namely, after applying modus ponens), the agent will know ψ" and "if φ is known, then *at some future point* (namely, after performing introspection), $\mathrm{K}\varphi$ will be known"; and similarly for the axiom (5) and the rule (Nec).

The simple temporal modality $\langle \mathrm{F} \rangle$ can further be refined into dynamic modalities whose programs represent deduction steps: e.g., composing the programs $\langle \mathrm{mp} \rangle$ for modus ponens and $\langle \mathrm{pi} \rangle$ for positive introspection, the appropriately modified axioms (K) and (4) yield:

$$\vdash \mathrm{K}(\varphi \to \psi) \to \big(\mathrm{K}\varphi \to \langle \mathrm{mp}; \mathrm{pi} \rangle \mathrm{K}\mathrm{K}\psi\big)$$

Another approach to resource-awareness consists in syntactic stratification of the epistemic modality K, e.g., by counting the steps needed for the logical derivation. The stratified axioms (K) and (4) then read:

$$\begin{aligned}&\vdash \mathrm{K}^n(\varphi \to \psi) \to (\mathrm{K}^m\varphi \to \mathrm{K}^{n+m+1}\psi)\\&\vdash \mathrm{K}^n\varphi \to \mathrm{K}^{n+1}\mathrm{K}^n\varphi,\end{aligned}$$

and similarly for the stratified axiom (5) and rule (Nec).

All of the aforementioned solutions avoid logical omniscience by deferring the knowledge deduced from the agent's actual knowledge to some future point. An unsatisfactory feature of these solutions, though, is that they solve the paradox not by treating the modality K itself, but by substituting some modification thereof, such as $\langle \mathrm{F} \rangle \mathrm{K}$ or K^n. The aim of the present paper is to solve the logical omniscience paradox by a resource-aware treatment of the very modality K. We will achieve this goal by distinguishing the actual, potential, and feasible knowledge of an epistemic agent and employing suitable non-classical logics for resource-aware reasoning about the costs of logical derivation.

3 The logics of costs and resources

As has been argued by the present author in one of the previous volumes of *The Logica Yearbook* (Běhounek, 2009), most kinds of typical resources exhibit the structure of a *(semi)linear residuated lattice.*[2] In more detail, it can be observed that many common kinds of resources come to us in amounts that can be compared and added or subtracted. The comparability establishes a linear lattice order on amounts; the addition and subtraction of amounts then constitute the residuated lattice's operations of fusion and residual. Admittedly, in some of the more complex types of resources (imagine, for

[2] In algebraic parlance, *semilinear* means subdirectly decomposable into linear components.

instance, the lists of cooking ingredients in recipe books), the order may no longer be linear. Nevertheless, even such resources can often be decomposed into some components (oil, flour, sugar, salt, etc.) whose amounts are fully comparable, and so the order is at least semilinear.

The logics that are sound and complete with respect to the algebraic semantics of (semi)linear residuated lattices are known as *fuzzy logics.*[3] Thus, arguably, logics that are generally suitable for resource-aware reasoning, epistemic or otherwise, can be found among fuzzy logics.[4]

Before we proceed, let us recall a few properties of propositional fuzzy logics (for more details see, e.g., Hájek, 1998; Běhounek, Cintula, & Hájek, 2011). In this paper, we will only need the following facts:

1. The standard set of truth values for fuzzy logics is the real unit interval $[0,1]$. Propositional connectives are interpreted truth-functionally, by certain well-behaved operations on $[0,1]$. Particular fuzzy logics differ in the choice of these truth functions for connectives.

2. The prominent fuzzy logics G, Ł, and Π interpret conjunction (or *fusion,* denoted by $\otimes$) by the following truth functions on $[0,1]$:

$$\begin{aligned} \text{Gödel–Dummett logic G:} \quad & \|\varphi \otimes \psi\| = \min(\|\varphi\|, \|\psi\|) \\ \text{Łukasiewicz logic Ł:} \quad & \|\varphi \otimes \psi\| = \max(0, \|\varphi\| + \|\psi\| - 1) \\ \text{Product fuzzy logic } \Pi\text{:} \quad & \|\varphi \otimes \psi\| = \|\varphi\| \cdot \|\psi\| \end{aligned}$$

3. These three logics, as well as all of their relatives from the family of so-called *t-norm fuzzy logics,* validate this law for implication:

$$\|\varphi \to \psi\| = 1 \quad \text{iff} \quad \|\varphi\| \leq \|\psi\|$$

Let us now give a brief description of the cost-based interpretation of t-norm fuzzy logics; for more details see the original paper (Běhounek,

[3] The delimitation of the class of fuzzy logics as the logics of classes of (semi)linear residuated lattices is due to Cintula (2006).

[4] By virtue of having the algebraic semantics of residuated lattices, fuzzy logics fall within the broader family of *substructural logics* (see, e.g., Ono, 2003; Kowalski & Ono, 2010). A substructural logic that is often regarded as the logic of resources is *linear logic.* However, as argued in the previous paper (Běhounek, 2009), using linear logic for resources disregards the semilinear structure manifested by most kinds of resources—in other words, neglects their decomposability into components with linearly ordered amounts. Consequently, linear logic as the logic of resources is unnecessarily weak, and under-generates resource-wise, compared to suitable fuzzy logics.

2009). Under this interpretation, the truth values of fuzzy logic represent *costs,* or amounts of a particular resource, such as money, time, or memory. The designated truth value 1 represents zero costs and all other truth values represent non-zero costs; the smaller the truth value, the larger the cost. The truth functions of fuzzy logic perform certain operations on costs. For instance, the operation of fusion $a \otimes b$ yields the value of the two costs a and b put together, while the implication $a \to b$ expresses the surcharge to a that would yield at least the cost b; the remaining propositional connectives of fuzzy logic possess a natural meaning in terms of costs as well. Particular fuzzy logics differ in the manner of combining costs: e.g., the way costs are combined by the standard truth function for $\otimes$ in Łukasiewicz logic Ł is that of *bounded additivity* (where 0 represents the maximal cost); in product fuzzy logic Π, *unbounded additivity* (via the logarithm, with 0 representing an infinite cost); and in Gödel–Dummett logic G, *maxitivity* (as is appropriate, e.g., for various kinds of capacity). The formulae of fuzzy logic then represent various ways of combining costs assigned to atoms.[5] Specifically, tautological implications $\varphi \to \psi$ express valid laws of cost preservation, of the form: "The cost ψ never exceeds the cost φ." Fuzzy logics can thus be viewed as calculi that formalize reasoning *salvis expensis,* in a similar manner as classical logic formalizes reasoning *salva veritate.*[6]

In the next sections I propose a solution to logical omniscience based on resource-aware reasoning modeled in fuzzy logic. The approach has already been briefly sketched in an earlier short paper (Běhounek, 2013). The present paper elaborates the solution in more detail, taking in part advantage of the recently developed fuzzy intensional semantics (Běhounek & Majer, 2018). Due to limited space, most of the mathematical apparatus is omitted here and deferred to a later full exposition.

[5]Formulae thus directly represent combinations of costs, rather than propositions or states of affairs. Compare this *formulae-as-costs* interpretation, e.g., with *formulae-as-types* in the Curry–Howard correspondence or the categorial grammar interpretation of the Lambek calculus. Alternatively, the logical atoms can be interpreted as gradual propositions of the form "the item x is inexpensive"; however, that requires an additional assumption on the correspondence between combining costs and combining degrees of truth by the propositional connectives.

[6]The general algebraic semantics of t-norm fuzzy logics in terms of semilinear residuated lattices generalizes this interpretation even to resources with non-linearly ordered amounts (yet decomposable into linearly ordered components).

4 Three kinds of knowledge

When discussing solutions to logical omniscience based on resource awareness, it is generally useful to distinguish the following three kinds of a given agent's knowledge:

- *Actual knowledge,* or the explicit knowledge that is immediately available to the agent. In artificial agents, this can be the contents of their database. In humans, it is the sum of all the facts the person knows without needing to infer them from other known facts.
- *Potential knowledge,* or the implicit knowledge that is, at least in principle, logically derivable from the actual knowledge. In other words, the set of all logical consequences of the agent's actual knowledge.
- And *feasible knowledge,* or that part of potential knowledge that the agent can feasibly derive from actual knowledge, taking into account the agent's limited resources (such as time, memory, etc.).

It can be noticed that logical omniscience is only troublesome for the notion of *feasible* knowledge, since the *actual* knowledge of a non-idealized real-world agent is never closed under logical consequence (one reason being its finiteness); while any agent's *potential* knowledge does indeed include all logical truths by definition.

It can furthermore be observed that whereas actual knowledge can be viewed as crisp and finite (and potential knowledge as crisp and infinite), feasible knowledge is apparently *gradual,* since long and complex logical derivations will require more of the agent's limited resources (e.g., time, memory, or energy) than shorter or simpler ones—and so can be *less feasibly* performed by the agent. The fact that fuzzy logics are, by design, tailored to deal with gradual notions just further underscores their suitability for resource-sensitive treatment of feasible knowledge.

As a matter of fact, the paradox of logical omniscience can be construed as an instance of the sorites paradox: a *single* additional step of logical derivation is always feasible and within the capability of a logically rational agent.[7] However, just like in the sorites paradox, it does not follow that arbitrarily long logical derivations would be feasible, with however many steps. In consequence of this correspondence, any solution to the sorites paradox generates a solution to the logical omniscience paradox. Our proposed solution

[7] Except when hard limits (e.g., on time or memory) are imposed. This case can be set aside, though, as then neither the sorites series of steps nor logical omniscience arise.

to logical omniscience thus, besides making use of resource-awareness, also parallels the treatment of the sorites paradox in fuzzy logic (cf. Hájek & Novák, 2003).

Combining these ideas, we will render feasible knowledge as a fuzzy set of formulae that are logically derivable from the actual knowledge, where the membership degrees represent the derivation costs. The next sections elaborate the fuzzy semantics of feasible knowledge in more detail.

5 Fuzzy modal logic of feasible knowledge: syntax

Following the previous considerations, we are going to formalize feasible knowledge as a unary fuzzy modality K governed by a suitable t-norm fuzzy logic $\mathcal{L}$. As explained in Section 3, the choice of the fuzzy logic depends on the way we intend to compound costs (for example, Łukasiewicz logic is suitable for bounded addition).

The epistemic modality K can, in principle, be applied to formulae of any language $\mathcal{E}$ in which the agent's knowledge is formalized.[8] Our syntax will, therefore, be two-layered, allowing:

(i) The modality K to be applied to a formula φ in the language $\mathcal{E}$ of the agent's knowledge representation, and

(ii) Modal atoms of the form $\mathrm{K}\varphi$ to be combined by the connectives of a propositional t-norm fuzzy logic $\mathcal{L}$ of our choice.

We will denote the logic with this two-layered syntax by $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$. For example, if the agent's knowledge is represented in first-order multi-agent doxastic logic $\mathrm{KD45}\forall^n_{\mathrm{B}}$ with unary modalities $\mathrm{B}_1, \ldots, \mathrm{B}_n$ and unary predicates P, Q, then $\mathrm{K}\big((\forall x)\mathrm{B}_1(Px \rightarrow \mathrm{B}_n Qx)\big) \otimes \mathrm{K}\big((\exists y)\neg \mathrm{B}_1 Py\big)$ is a well-formed formula of the logic $Ł_{\mathrm{K}}[\mathrm{KD45}\forall^n_{\mathrm{B}}]$.

It can be noticed that in $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$, the modality K of feasible knowledge cannot be nested. This is because the 'outer' logic $\mathcal{L}$ serves for *our* reasoning about the agent's feasible knowledge, and its formulae are not part of the agent's knowledge. It is, nevertheless, possible that the language $\mathcal{E}$ of the agent's reasoning contains its own epistemic modality that refers to the agent's own knowledge: for instance, the agent might use the standard epistemic logic $\mathrm{S5}_{\mathrm{K}}$ for their epistemic reasoning. Strictly speaking, the

[8]The language is usually equipped with a set of derivation rules the agent can apply, i.e., a logic of the agent's reasoning. We will therefore regard $\mathcal{E}$ rather as a logic than just its language.

agent's ('inner') epistemic modality and our ('outer') epistemic modality K should always be denoted by different symbols. Nevertheless, since they are distinguished by the level of nesting ('our' K's being the outermost ones), in logics such as $\text{Ł}_{\mathrm{K}}[\mathrm{S5}_{\mathrm{K}}]$ we will take the liberty of using the same symbol K on both syntactic levels. This will make our syntax closer to that of standard epistemic logic and allow us, for example, to write positive introspection in its traditional form $\mathrm{K}\varphi \rightarrow \mathrm{KK}\varphi$.

Besides nesting, the two-layered syntax of $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$ also prohibits combining modal and non-modal atoms by the connectives of $\mathcal{L}$. This restriction appears natural, as the connectives of $\mathcal{L}$ are intended to just combine the costs of knowledge. It is, nevertheless, reasonable to relax this constraint and permit forming mixed formulae too. One reason is that non-modal atoms can as well be employed to denote costs (as, e.g., in Section 8 below), and then they become meaningfully combinable with the costs of knowledge. Moreover, the costs themselves can be regarded as the degrees of truth of certain propositions (e.g., when we interpret $\mathrm{K}\varphi$ as the graded statement *"to infer φ from the actual knowledge is inexpensive"*), and thus as freely combinable with any other graded propositions of the logic $\mathcal{L}$. So, provided that $\mathcal{E}$-formulae can be assigned truth values of $\mathcal{L}$ (as is the case, e.g., if the logic $\mathcal{E}$ is bivalent, since all $\mathcal{L}$-algebras contain the truth values 0 and 1), we will also consider the extension of $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$ that admits mixed formulae and denote it by $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$. This will allow us, for instance, to discuss the truth axiom $\mathrm{K}\varphi \rightarrow \varphi$ within the framework of such logics as $\text{Ł}_{\mathrm{K}}(\mathrm{S5})$.[9]

6 Fuzzy modal logic of feasible knowledge: semantics

We will introduce our cost-sensitive fuzzy modal logic $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$ of feasible knowledge by means of a fuzzy-relational possible-world semantics, where possible worlds represent the agent's epistemic states. Each possible world is assigned a set of $\mathcal{E}$-formulae that form the agent's actual knowledge in that state. Transitions between the states correspond to inference steps the agent is able to perform; i.e., to changes of the actual knowledge in consequence of deductions performed by the agent. As usual, possible transitions are encoded by an accessibility relation; in our case, the relation is weighted ($\mathcal{L}$-valued, or fuzzy) and the weights represent the costs of the transitions.

A sample model for $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$ is depicted in Figure 1. The agent's possible

[9] If $\mathcal{E}$ contains its own epistemic modality, then $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$ needs, unlike $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$, to distinguish both modalities graphically, in order to disambiguate such formulae as $\mathrm{KK}\varphi \rightarrow \mathrm{K}\varphi$.

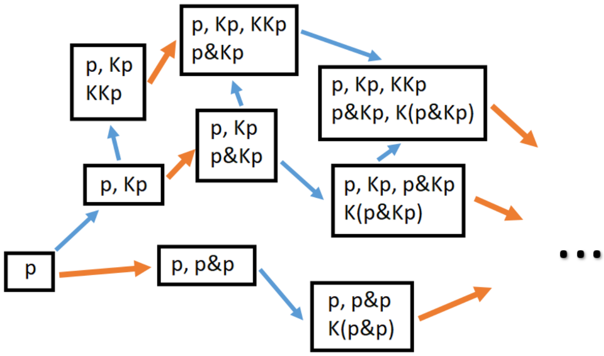

Figure 1: Example of an epistemic model.

epistemic states (represented by boxes) differ in the agent's actual knowledge (written inside the boxes). For simplicity, let us assume that the logic $\mathcal{E}$ of the agent's reasoning has only the rules of adjunction and positive introspection. Possible transitions between the states, brought about by application of these deduction rules, are indicated by arrows. The width of each arrow signifies the cost incurred by the agent for performing the inference step: in this model, applying adjunction is less costly (hence, thicker arrows). Collectively, the arrows represent the cost-weighted accessibility relation between the epistemic states.[10]

Let us have a look at how the agent's feasible knowledge is computed for any given epistemic state in Figure 1. Suppose, for instance, that the actual world w_0 is the leftmost one, in which the agent's actual knowledge consists just of the single $\mathcal{E}$-formula p, and let us calculate the degree to which the formula $\mathrm{K}(p \mathbin{\&} \mathrm{K}p)$ is part of the agent's feasible knowledge in that state; in other words, we want to determine the truth value of the $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$-formula $\mathrm{K}\big(\mathrm{K}(p \mathbin{\&} \mathrm{K}p)\big)$ in the state w_0.[11] It can be observed that in this model, the

[10]The figure omits the arrows that arise by composition of the depicted single-step arrows; their costs can be calculated as the $\mathcal{L}$-conjunction (fusion) of the costs of the single-step transitions. Also omitted are the full-width (i.e., zero-cost) arrows from each state to itself. These all, too, are part of the graded accessibility relation, which equals the $\mathcal{L}$-valued reflexive-transitive hull of the weighted single-step arrows from the picture (cf. Section 7).

[11]Recall that the outermost K is the $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$-modality governed by the fuzzy logic of costs $\mathcal{L}$,

$\mathcal{E}$-formula $\mathrm{K}(p \,\&\, \mathrm{K}p)$ belongs to the agent's actual knowledge just in the two upper-rightmost depicted states. We can see that in order to arrive at this knowledge (i.e., at one of these states), the agent needs to perform positive introspection twice and adjunction once. The cost of deriving this knowledge is thus the fusion of the costs of these three inference steps. Consequently, the truth value of $\mathrm{K}\big(\mathrm{K}(p \,\&\, \mathrm{K}p)\big)$ in the state w_0 is the $\mathcal{L}$-conjunction $\alpha \otimes \beta \otimes \alpha$, where α is the degree representing the cost of positive introspection (or the weight of the thinner arrows in Figure 1) and β the cost of adjunction (the thicker arrows) in our $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$-model of Figure 1.

It can be noticed that in this $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$-model, there are also other (longer) paths from w_0 to a state containing $\mathrm{K}(p \,\&\, \mathrm{K}p)$. These paths consist of three applications of positive introspection and one adjunction, and so are no cheaper than the shorter route described above. Since we are interested in resources that must *unavoidably* be spent to arrive at the knowledge, we take the cheapest route, or generally the infimum of costs (which is the supremum of degrees) along all possible paths.[12]

Models for $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$, which admits mixed formulae, will only need to additionally set an $\mathcal{L}$-evaluation of $\mathcal{E}$-formulae in each state. Let us summarize these ideas in a definition of fuzzy epistemic models.

Definition 1 *A* fuzzy epistemic model *for $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$ is a tuple $\mathbf{M} = (W, \boldsymbol{L}, R, A, e)$, where:*

- *W is a non-empty set (of an agent's possible epistemic states);*
- *$\boldsymbol{L}$ is an $\mathcal{L}$-algebra (of truth degrees representing costs);*
- *$R\colon W^2 \to \boldsymbol{L}$ is an $\boldsymbol{L}$-valued weighted accessibility relation on W (representing the transition costs between states);*
- *$A\colon W \times \mathit{Form}_{\mathcal{E}} \to \{0,1\}$ is a relation indicating the agent's actual knowledge $A^w = \{\varphi \in \mathit{Form}_{\mathcal{E}} \colon A(w,\varphi) = 1\}$, in each state w; and*

while the inner modalities K belong to the logic $\mathcal{E}$ of the agent's reasoning.

[12] Admittedly, this is a gross idealization, as *finding* the cheapest route would generally be a non-trivial problem for the agent. (Thanks are due to Sebastian Sequoiah-Grayson for pointing this out at *Logica 2019*.) We adopt the idealization here, since—unlike the resource-obliviousness of standard epistemic logic—it does not produce logical omniscience. Indeed, a more realistic model of cost-aware epistemic reasoning ought to address it—possibly by replacing the costs of moving along the cheapest path with the costs of *searching* for whichever path to the desired knowledge. Then, besides the inferential distance in the logic $\mathcal{E}$, the feasibility degree would also depend on the agent's path-searching heuristics and algorithms.

- $e\colon W \times \mathit{Form}_{\mathcal{E}} \to \boldsymbol{L}$ *is an evaluation of* $\mathcal{E}$*-formulae in each state.*[13]

The truth value $\|\psi\|_w$ *of an* $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$*-formula* ψ *in a state* $w \in W$ *is defined by the following Tarski conditions, for all* n*-ary* $\mathcal{L}$*-connectives* c*,* $\mathcal{E}$*-formulae* φ*, and* $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$*-formulae* $\psi_1, \ldots, \psi_n$*:*

$$\begin{aligned} \|\varphi\|_w &= e(w, \varphi) \\ \|c(\psi_1, \ldots, \psi_n)\|_w &= c^{\boldsymbol{L}}(\|\psi_1\|_w, \ldots, \|\psi_n\|_w) \\ \|\mathrm{K}\varphi\|_w &= \textstyle\bigvee_{w' \in W,\, \varphi \in A^{w'}} Rww' \end{aligned}$$

The notions of intension $\|\psi\|\colon w \mapsto \|\psi\|_w$, tautologicity and (local) entailment with respect to a class of fuzzy epistemic models are defined in a standard manner.[14]

7 Constraints on fuzzy epistemic models

To achieve some level of real-world plausibility, fuzzy epistemic models need be constrained by some conditions on the weighted accessibility relation that would reflect various principles of the agent's epistemic reasoning. Some of such constraints are the following:[15]

1. Fuzzy transitivity of R:

   $$\forall w, w', w'' \in W\colon Rww' \otimes Rw'w'' \leq Rww''.$$

   This condition reflects the concatenability of $\mathcal{E}$-derivations: if an $\mathcal{E}$-derivation δ_1 takes the agent from an epistemic state w to a state w' at the cost Rww' and a derivation δ_2 makes the transition from w' to w'' at the cost $Rw'w''$, then the transition from w to w'' costs at most $Rww' \otimes Rw'w''$, or the cost of the concatenated derivation $\delta_1\delta_2$.

[13] Note that although both e and A assign values to $\mathcal{E}$-formulae, their roles are different: the value $e(w, \varphi)$ determines the (degree of) *truth* of φ in w, while $A(w, \varphi)$ indicates whether φ is part of the agent's actual *knowledge*. We assume that e, though taking values in $\boldsymbol{L}$, is also an admissible evaluation of formulae in the sense of $\mathcal{E}$ (cf. the end of Section 5). In fuzzy epistemic models for the logic $\mathcal{L}_{\mathrm{K}}[\mathcal{E}]$, which does not admit mixed formulae, this component of a model is simply omitted.

[14] As usual in t-norm fuzzy logics, only the truth degree 1 designated in $\boldsymbol{L}$. Notice that in infinite models, the existence in $\boldsymbol{L}$ of all the suprema required by the Tarski condition for K need be assumed. This can be ensured either by requiring the lattice-completeness of $\boldsymbol{L}$, or (to enable axiomatizability) by using *safe* models (cf. Hájek, 1998, ch. 5).

[15] The properties of fuzzy transitivity, fuzzy reflexivity, upper sets, preimages, and closures are well studied in fuzzy set theory (e.g., Bělohlávek, 2002; Běhounek, Bodenhofer, & Cintula, 2008). Fuzzy relations that are fuzzy transitive and fuzzy reflexive are called fuzzy preorders.

2. Fuzzy reflexivity of R (i.e., $\forall w \in W\colon Rww = 1$), which expresses the immediate (cost-free) availability of actual knowledge.

3. Persistence, or the upperness of $A_\varphi = \{w \in W\colon A(w,\varphi) = 1\}$ in R for each $\mathcal{E}$-formula φ.

In fuzzy epistemic models that satsify these constraints, feasible knowledge can be characterized as a closure operator on the epistemic state space:

Observation 1 *Let $\mathbf{M} = (W, \boldsymbol{L}, R, A, e)$ be a fuzzy epistemic model for $\mathcal{L}_{\mathrm{K}}(\mathcal{E})$ satisfying the conditions 1–3 above and let φ be an $\mathcal{E}$-formula. Then:*

- *The intension $A_\varphi \subseteq W$ of the actual knowledge of φ is a crisp upper set in the $\boldsymbol{L}$-valued fuzzy preorder R on W.*
- *The fuzzy intension $\|\mathrm{K}\varphi\|\colon w \mapsto \|\mathrm{K}\varphi\|_w$ of the feasible knowledge of φ is the fuzzy preimage of A_φ in R. In symbols, $\|\mathrm{K}\varphi\| = R^{\leftarrow} A_\varphi$.*

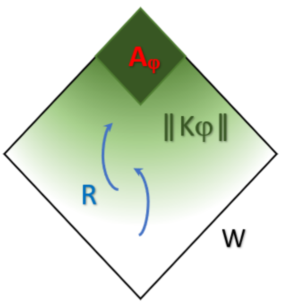

Figure 2: The fuzzy intension $\|\mathrm{K}\varphi\|$ of the feasible knowledge of φ is the fuzzy preimage of the (crisp upper) intension A_φ of the actual knowledge of φ in the fuzzy accessibility preorder R; i.e., $\|\mathrm{K}\varphi\| = R^{\leftarrow} A_\varphi$.

Moreover, since R is a fuzzy preorder, the fuzzy preimage operator $X \mapsto R^{\leftarrow} X$, where $(R^{\leftarrow} X)(w) = \bigvee_{w' \in W}(Rww' \otimes Xw')$ for all $X\colon W \to \boldsymbol{L}$ and $w \in W$, satisfies the conditions of a fuzzy closure operator.[16]

[16]Namely, is pointwise extensive and monotone w.r.t. $\leq^{\boldsymbol{L}}$, and $R^{\leftarrow}(R^{\leftarrow} X) = R^{\leftarrow} X$ for all $X\colon W \to \boldsymbol{L}$. Additionally, $R^{\leftarrow} \emptyset = \emptyset$. The (easy) proofs can be found, e.g., in (Běhounek et al., 2008).

Consequently, the feasible knowledge $\|\mathrm{K}\varphi\|$ *can also be described as the (backward) fuzzy closure of the actual knowledge* A_φ *under the cost-weighted fuzzy transition relation* R *on epistemic states.*

The list of constraints 1–3 above is neither unalterable nor complete. Reasonable conditions on fuzzy epistemic models depend largely on the assumed properties of the epistemic agent, and particularly on the logic $\mathcal{E}$ of the agent's reasoning. For example, persistence of actual knowledge (condition 3) can hardly be required if $\mathcal{E}$ is non-monotonic. On the other hand, possible additions may include the following constraints:

- The finiteness of A^w for all $w \in W$ (finite actual knowledge in each epistemic state).
- Pointwise inclusion[17] of the actual knowledge A in the model's $\mathcal{E}$-evaluation e (facticity of actual knowledge); or in doxastic variants, pointwise inclusion of A in any $\mathcal{E}$-evaluation $e'\colon W \times \mathit{Form}_{\mathcal{E}} \to \boldsymbol{L}$ (consistency of actual beliefs).
- Constant evaluation of $\mathcal{E}$-formulae across all epistemic states, i.e., $e(w,\varphi) = e(w',\varphi)$ for all $w \in W$ and $\varphi \in \mathit{Form}_{\mathcal{E}}$ (static facts).[18]
- The property of (weighted) *confluence* for transitions between epistemic states, which reflects free combinability of derivations (if valid for $\mathcal{E}$):

$$\begin{aligned}&\forall w, w_1, w_2 \in W\, \exists w_3 \in W\colon\\ &\qquad A^{w_3} = A^{w_1} \cup A^{w_2}, Rww_2 \le Rw_1w_3, Rww_1 \le Rw_2w_3.\end{aligned}$$

- If each $\mathcal{E}$-inference step produces a single $\mathcal{E}$-formula, the models are bound to satisfy the following condition for each $w \in W$ and $n \in \mathbb{N}$:

$$\bigvee\nolimits_{w' \in [w]_{n+1}} Rww' \le \bigvee\nolimits_{w' \in [w]_n} Rww',$$

  where $[w]_n = \{w \in W\colon\ \mathrm{Card}(A^{w'} \smallsetminus A^w) = n\}$.

[17]I.e., $A(w,\varphi) \le e(w,\varphi)$ for all $w \in W$ and $\varphi \in \mathit{Form}_{\mathcal{E}}$.

[18]Notice that if $\mathcal{E}$-evaluation is state-dependent as in Definition 1, the truth-axiom $\mathrm{K}\varphi \to \varphi$ need not be valid in a model even if A is pointwise included in e, since the evaluation of φ can change during derivations.

8 The costs of inference steps

The general requirements on fuzzy epistemic models listed in Section 7 abstract from specific inferential abilities of epistemic agents. Still, it may sometimes be desirable to have the agent's capability of performing particular inferential steps (such as modus ponens or positive introspection) captured axiomatically, in a similar way as in standard epistemic logic. Lest we fall back into the logical omniscience paradox, though, the modal axioms of standard epistemic logic need to be adapted to reflect the agent's cost of inference. For instance, the cost-sensitive modifications of the epistemic axioms (K) and (4) might read:

(K′) $\vdash \mathrm{K}(\varphi \to \psi) \otimes \mathrm{K}\varphi \otimes \mathrm{m}_{\varphi,\psi} \to \mathrm{K}\psi$

(4′) $\vdash \mathrm{K}\varphi \otimes \mathrm{i}_{\varphi} \to \mathrm{KK}\varphi$

Here, $\mathrm{m}_{\varphi,\psi}$ and i_{φ} are new propositional constants added to $\mathcal{L}$, possibly different for each $\varphi, \psi \in \mathit{Form}_{\mathcal{E}}$ (as the costs of inference steps generally depend on the formulae involved—e.g., on their length).

The appropriate values of these propositional constants in a model are contingent on the properties of the particular agent, and for many agents (such as humans) can hardly be determined precisely. Nevertheless, independently of their exact values, their presence in the axioms ensures that the costs of inference steps are accounted for. In longer $\mathcal{E}$-derivations, these costs accumulate by the fusion connective $\otimes$ of $\mathcal{L}$, eventually making too long derivations from the actual knowledge infeasible. This eliminates logical omniscience while keeping the agent inferentially capable, by virtue of the cost-sensitivity of the axioms of logical rationality in $\mathcal{L}$.

References

Artemov, S., & Kuznets, R. (2014). Logical omniscience as infeasibility. *Annals of Pure and Applied Logic*, *165*, 6–25.

Běhounek, L. (2009). Fuzzy logics interpreted as logics of resources. In M. Peliš (Ed.), *The Logica Yearbook 2008* (pp. 9–21). College Publications.

Běhounek, L. (2013). Feasibility as a gradual notion. In A. Voronkov et al. (Eds.), *LPAR-17-short: Short Papers for 17th LPAR* (Vol. 13 of EasyChair Proceedings in Computing, pp. 15–19).

Běhounek, L., Bodenhofer, U., & Cintula, P. (2008). Relations in Fuzzy Class Theory: Initial steps. *Fuzzy Sets and Systems*, *159*, 1729–1772.

Běhounek, L., Cintula, P., & Hájek, P. (2011). Introduction to mathematical fuzzy logic. In P. Cintula, P. Hájek, & C. Noguera (Eds.), *Handbook of Mathematical Fuzzy Logic* (Vol. I, pp. 1–101). College Publications.

Běhounek, L., & Majer, O. (2018). Fuzzy intensional semantics. *Journal of Applied Non-Classical Logics*, *28*, 348–388.

Bělohlávek, R. (2002). *Fuzzy Relational Systems: Foundations and Principles* (Vol. 20). Kluwer/Plenum.

Cintula, P. (2006). Weakly implicative (fuzzy) logics I: Basic properties. *Archive for Mathematical Logic*, *45*(6), 673–704.

Duc, H. N. (1997). Reasoning about rational, but not logically omniscient, agents. *Journal of Logic and Computation*, *7*, 633–648.

Duc, H. N. (2001). *Resource-Bounded Reasoning about Knowledge* (Unpublished doctoral dissertation). University of Leipzig.

Hájek, P. (1998). *Metamathematics of Fuzzy Logic*. Kluwer.

Hájek, P., & Novák, V. (2003). The sorites paradox and fuzzy logic. *International Journal of General Systems*, *32*, 373–383.

Kowalski, T., & Ono, H. (2010). Fuzzy logics from substructural perspective. *Fuzzy Sets and Systems*, *161*, 301–310.

Meyer, J.-J. C. (2003). Modal epistemic and doxastic logic. In D. Gabbay & F. Guenthner (Eds.), *Handbook of Philosophical Logic* (2nd ed., Vol. 10, pp. 1–38). Kluwer.

Ono, H. (2003). Substructural logics and residuated lattices—an introduction. In V. F. Hendricks & J. Malinowski (Eds.), *50 Years of Studia Logica* (Vol. 21, pp. 193–228). Kluwer.

Libor Běhounek
University of Ostrava, CE IT4Innovations–IRAFM
The Czech Republic
E-mail: `libor.behounek@osu.cz`