Coming Upon the Classic Notion of Implicit Knowledge Again

Bernhard Heinemann

Faculty of Mathematics and Computer Science,
University of Hagen,
58084 Hagen, Germany
bernhard.heinemann@fernuni-hagen.de

Abstract. Subsequently, we introduce a novel semantics for the bimodal logic of subset spaces, denoted by LSS. This system was originally invented by Moss and Parikh for the purpose of clarifying the intrinsic relationship between the epistemic notion of knowledge and the geometric concept of topology. Focussing on the knowledge-theoretic side in this paper, we re-adjust LSS to multi-agent scenarios. As a result, a particular dynamic logic of implicit knowledge is obtained. This finds expression in the technical outcome of the paper, which covers soundness, completeness, decidability, and complexity issues regarding the arising system.

Keywords: epistemic logic, implicit knowledge, subset space semantics, topological reasoning.

1 Introduction

Reasoning about knowledge constitutes an important foundational issue in Artificial Intelligence. We concentrate on some of its logical aspects in this paper. In particular, we are concerned with the idea of *implicit* (or *distributed*) *knowledge* of a group of agents.

Meanwhile, several instructive and very readable treatises on the diverse logics of knowledge are available. While more recent publications rather stress the *dynamics* of informational perceptions, including aspects of *belief, desire,* or *intention* (*BDI*) (see, e.g., [3], as well for further references), the classic textbooks [7] and [13] can thoroughly serve as a common ground for the fundamentals of epistemic logic needed here. Accordingly, given a finite collection G of agents, a binary *accessibility relation* R_A connecting *possible worlds* or *conceivable states of the world,* is associated with every agent $A \in G$. The *knowledge of* A is then defined through the *validity* of the corresponding formulas at all states the agent considers possible at the actual one. Now, *collecting together* such 'locally allocated' knowledge means ruling out those worlds that are inconceivable to some of the agents in G. To put it another way, the implicit knowledge of the agents under discussion is represented exactly by *intersecting* the respective sets of accessible states; see [13], Sect. 2.3, or [7], Sect. 2.2 and Sect. 3.4. (Throughout this paper, the term *implicit knowledge* is used, as in [13]; on the other hand,

R. Buchmann et al. (Eds.): KSEM 2014, LNAI 8793, pp. 1–12, 2014.

the term *distributed knowledge* is employed in the latter reference, since the idea of *awareness* (and, therefore, that of *explicit* knowledge) enters the field there.)

Moss and Parikh's bi-modal logic of subset spaces, LSS (see [14], [5], or Ch. 6 of [1]), may be rated as a cross-disciplinary framework for dealing with topological as well as epistemic scenarios. This is exemplified in the single-agent case subsequently. The *epistemic state* of an agent in question, i.e., the set of all those states that cannot be distinguished by what the agent topically knows, can be viewed as a *neighborhood* U of the actual state x of the world. Formulas are then interpreted with respect to the resulting pairs x, U called *neighborhood situations.* Thus, both the set of all states and the set of all epistemic states constitute the relevant semantic domains as particular subset structures. The two modalities involved, K and $\Box$, quantify over all elements of U and 'downward' over all neighborhoods contained in U, respectively. This means that K captures the notion of knowledge as usual (see [7] or [13] again), and $\Box$ reflects *effort to acquire knowledge* since gaining knowledge goes hand in hand with a shrinkage of the epistemic state. In fact, knowledge acquisition is this way reminiscent of a topological procedure. The appropriate logic of 'real' topological spaces as well as that of further computationally interesting spatial structures (viz tree-like ones) were examined by Georgatos rather promptly; see [8], [9]. The ongoing research into subset and topological spaces, respectively, is reported in the handbook [1]. More recent developments include the papers [12], [2], and [15], with the last two forging links between subset spaces and *Dynamic Epistemic Logic (DEL);* see [6].

Most papers on LSS deal with the single-agent case. Notwithstanding this, a multi-agent version was suggested in [10] (see also [11]). The key idea behind these papers is as follows: incorporate the agents in terms of additional modalities and, apart from this variation of the *logic,* let the original semantics be unchanged. However, what happens when, in contrast, the *semantics* is modified, and even in a way suggesting itself, namely to the effect that the agent structure is reflected in the *atomic* semantic entities already? – It turns out that the scope of the modality K has to be restricted then, but fortunately in a quite acceptable manner: K hereby mutates to an *implicit knowledge operator* (and, as will become apparent later, the logic remains the same in this case).

This idea will be implemented in the rest of this paper. Our aim is to give precise definitions as related to the underlying language, state the axioms and rules of the arising logic, prove soundness and completeness with respect to the intended class of domains, and reason about the intrinsic effectiveness and efficiency properties. (However, we must omit elaborate examples, due to the lack of space; in this respect, the reader is referred to the quoted literature.) The outcome we strive for is, in fact, an alternative modal description of implicit knowledge, and in the presence of a rather general operator describing increase of individual knowledge. (Thus, we are not ambitious in producing a system 'beating' others (in particular, more differentiated ones) here. But note that it is very desirable to have to hand distinct (e.g., differently fine-grained) ways of seeing a subject: this would allow one to react on varying problems flexibly; hence such a broadening of the horizon is a widespread practice in many mathematically oriented fields.)

The subsequent technical part of the paper is organized as follows. In the next section, we recapitulate the basics of multi-agent epistemic logic. The facts we need from the logic of subset spaces are then listed in Section 3. Section 4 contains the new multi-agent setting of LSS. Our main results, including some proof sketches, follow in the next two sections. Finally, we conclude with a summing up and a couple of additional remarks. – An attempt has been made to keep the paper largely self-contained. However, acquaintance of the reader with basic modal logic has to be assumed. As to that, the textbook [4] may serve as a standard reference.

2 Revisiting the Most Common Logic of Knowledge

All languages we consider in this paper are based on a denumerably infinite set $\mathsf{Prop} = \{p, q, \ldots\}$ of symbols called *proposition variables* (which should represent the basic facts about the states of the world). Let $n \in \mathbb{N}$ be given (the number of agents under discussion). Then, our modal language for knowledge contains, among other things, a one-place operator K_i representing the i-th agent's knowledge, for every $i \in \{1, \ldots, n\}$. The set KF of all *knowledge formulas* is defined by the rule

$$\alpha \ ::= \ \top \mid p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \mathsf{K}_i\alpha \mid \mathsf{I}\,\alpha,$$

where $i \in \{1, \ldots, n\}$. The missing boolean connectives will be treated as abbreviations, as needed. The connective I is called the *implicit knowledge operator.* Moreover, the modal duals of K_i and I are denoted by L_i and J, respectively.

As was indicated right at the outset, each of the operators K_i comes along with a binary relation R_i on the set X of all states of the world. The kind of knowledge we would like to model should certainly be mirrored in the characteristics of these relations. Having multi-agent systems à la [7] in mind where 'accessibility' means 'indistinguishability of the local states of the other agents', one is led to *equivalence relations* actually. Furthermore, the *intersection* of these equivalences is the relation associated with the implicit knowledge operator. Thus, the multi-modal frames for interpreting the above formulas are tuples $F = (X, R_1, \ldots, R_n, R_{\mathsf{I}})$, where X is a non-empty set, $R_i \subseteq X \times X$ is an equivalence relation for every $i \in \{1, \ldots, n\}$, and $R_{\mathsf{I}} = \bigcap_{i=1,\ldots,n} R_i$. And a model M based on such a frame is obtained by adding a valuation to the frame, i.e., a mapping V from Prop into the powerset of X, determining those states where the respective proposition variables become valid. Satisfaction of formulas is then defined *internally,* i.e., in models at particular states. We here remind the reader of the case of a modal operator, say the one for implicit knowledge:

$$M, x \models \mathsf{I}\,\alpha :\Longleftrightarrow \text{ for all } y \in X : \text{if } (x, y) \in R_{\mathsf{I}}, \text{ then } M, x \models \alpha,$$

for all $x \in X$ and $\alpha \in \mathsf{KF}$. – The just described semantics is accompanied by a logic which is a slight extension of the multi-modal system $\mathrm{S5}_{n+1}$. This means that we have, in particular, the well-known and much-debated knowledge and introspection axioms for each of the modalities involved, for example, those relating to the I-operator:

- $\mathsf{I}\,\alpha \rightarrow \alpha$
- $\mathsf{I}\,\alpha \rightarrow \mathsf{I}\,\mathsf{I}\,\alpha$
- $\mathsf{J}\,\alpha \rightarrow \mathsf{I}\,\mathsf{J}\,\alpha$,

where $\alpha \in \mathsf{KF}$. The schemata

- $\mathsf{K}_i\alpha \rightarrow \mathsf{I}\,\alpha$, for every $i \in \{1, \ldots, n\}$,

designated ($\mathbb{I}$), constitute the extension of $\mathrm{S5}_{n+1}$ addressed a moment ago. The following is taken from [7], Theorem 3.4.1 (see also [13], Theorem 2.3.2).

Theorem 1. *The logic* $\mathrm{S5}_{n+1} + (\mathbb{I})$ *is sound and complete with respect to the class of models described above.*

While sketching a proof of this theorem, the authors of [7] point to the difficulties related to the intersection property (i.e., $R_{\mathsf{I}} = \bigcap_{i=1,\ldots,n} R_i$) on the way towards completeness. We, too, shall encounter this problem, in Section 5 (albeit in weakened form).

3 The Language and the Logic of Subset Spaces

In this section, we first fix the language for subset spaces, $\mathcal{L}$. After that, we link the semantics of $\mathcal{L}$ with the common relational semantics of modal logic. (This link will be utilized later in this paper.) Finally, we recall some facts on the logic of subset spaces needed subsequently. – The proceeding in this section is a bit more rigorous than that in the previous one, since $\mathcal{L}$ and LSS are assumed to be less established.

To begin with, we define the syntax of $\mathcal{L}$. Let the set SF of all *subset formulas*[1] over Prop be defined by the rule

$$\alpha \ ::= \ \top \mid p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \mathsf{K}\alpha \mid \Box\alpha.$$

Here, the duals of K and $\Box$ are denoted by L and $\Diamond$, respectively. In view of our considerations in the introduction, K is called the *knowledge operator* and $\Box$ the *effort operator*.

Second, we examine the semantics of $\mathcal{L}$. For a start, we define the relevant domains. We let $\mathcal{P}(X)$ designate the powerset of a given set X.

Definition 1 (Semantic Domains)

1. *Let X be a non-empty set (of* states*) and $\mathcal{O} \subseteq \mathcal{P}(X)$ a set of subsets of X. Then, the pair $\mathcal{S} = (X, \mathcal{O})$ is called a* subset frame.

[1] The prefix 'subset' will be omitted provided there is no risk of confusion.

2. *Let* $\mathcal{S} = (X, \mathcal{O})$ *be a subset frame. Then the set* $\mathcal{N}_{\mathcal{S}} := \{(x, U) \mid x \in U \text{ and } U \in \mathcal{O}\}$ *is called the* set of neighborhood situations *of* $\mathcal{S}$.
3. *Let* $\mathcal{S} = (X, \mathcal{O})$ *be a subset frame. An* $\mathcal{S}$-valuation *is a mapping* $V : \mathsf{Prop} \to \mathcal{P}(X)$.
4. *Let* $\mathcal{S} = (X, \mathcal{O})$ *be a subset frame and* V *an* $\mathcal{S}$*-valuation. Then,* $\mathcal{M} := (X, \mathcal{O}, V)$ *is called a* subset space (based on $\mathcal{S}$).

Note that neighborhood situations denominate the semantic atoms of the bi-modal language $\mathcal{L}$. The first component of such a situation indicates the actual state of the world, while the second reflects the uncertainty of the agent in question about it. Furthermore, Definition 1.3 shows that values of proposition variables depend on states only. This is in accordance with the common practice in epistemic logic; see [7] or [13] once more.

For a given subset space $\mathcal{M}$, we now define the relation of *satisfaction,* $\models_{\mathcal{M}}$, between neighborhood situations of the underlying frame and formulas from SF. Based on that, we define the notion of *validity* of formulas in subset spaces. In the following, neighborhood situations are often written without parentheses.

Definition 2 (Satisfaction and Validity). *Let* $\mathcal{S} = (X, \mathcal{O})$ *be a subset frame.*

1. *Let* $\mathcal{M} = (X, \mathcal{O}, V)$ *be a subset space based on* $\mathcal{S}$*, and let* $x, U \in \mathcal{N}_{\mathcal{S}}$ *be a neighborhood situation. Then*

$$\begin{array}{lll}
x, U \models_{\mathcal{M}} \top & & \text{is always true} \\
x, U \models_{\mathcal{M}} p & :\Longleftrightarrow & x \in V(p) \\
x, U \models_{\mathcal{M}} \neg\alpha & :\Longleftrightarrow & x, U \not\models_{\mathcal{M}} \alpha \\
x, U \models_{\mathcal{M}} \alpha \wedge \beta & :\Longleftrightarrow & x, U \models_{\mathcal{M}} \alpha \text{ and } x, U \models_{\mathcal{M}} \beta \\
x, U \models_{\mathcal{M}} \mathsf{K}\alpha & :\Longleftrightarrow & \forall\, y \in U : y, U \models_{\mathcal{M}} \alpha \\
x, U \models_{\mathcal{M}} \Box\alpha & :\Longleftrightarrow & \forall\, U' \in \mathcal{O} : [x \in U' \subseteq U \Rightarrow x, U' \models_{\mathcal{M}} \alpha],
\end{array}$$

where $p \in \mathsf{Prop}$ *and* $\alpha, \beta \in \mathsf{SF}$*. In case* $x, U \models_{\mathcal{M}} \alpha$ *is true we say that* α holds in $\mathcal{M}$ at *the neighborhood situation* x, U.
2. *Let* $\mathcal{M} = (X, \mathcal{O}, V)$ *be a subset space based on* $\mathcal{S}$*. A subset formula* α *is called* valid in $\mathcal{M}$ *iff it holds in* $\mathcal{M}$ *at every neighborhood situation of* $\mathcal{S}$.

Note that the idea of knowledge and effort described in the introduction is made precise by Item 1 of this definition. In particular, knowledge is defined as validity at all states that are indistinguishable to the agent here, too.

Subset frames and spaces can be considered from a different perspective, as is known since [5] and reviewed in the following. Let a subset frame $\mathcal{S} = (X, \mathcal{O})$ and a subset space $\mathcal{M} = (X, \mathcal{O}, V)$ based on $\mathcal{S}$ be given. Take $X_{\mathcal{S}} := \mathcal{N}_{\mathcal{S}}$ as a set of worlds, and define two accessibility relations $R^{\mathsf{K}}_{\mathcal{S}}$ and $R^{\Box}_{\mathcal{S}}$ on $X_{\mathcal{S}}$ by

$$\begin{array}{l}
(x, U)\, R^{\mathsf{K}}_{\mathcal{S}}\, (x', U') :\Longleftrightarrow\ U = U' \text{ and} \\
(x, U)\, R^{\Box}_{\mathcal{S}}\, (x', U') :\Longleftrightarrow\ (x = x' \text{ and } U' \subseteq U),
\end{array}$$

for all $(x, U), (x', U') \in X_{\mathcal{S}}$. Moreover, let $V_{\mathcal{M}}(p) := \{(x, U) \in X_{\mathcal{S}} \mid x \in V(p)\}$, for every $p \in \mathsf{Prop}$. Then, bi-modal Kripke structures $S_{\mathcal{S}} := \big(X_{\mathcal{S}}, \{R^{\mathsf{K}}_{\mathcal{S}}, R^{\Box}_{\mathcal{S}}\}\big)$ and $M_{\mathcal{M}} := \big(X_{\mathcal{S}}, \{R^{\mathsf{K}}_{\mathcal{S}}, R^{\Box}_{\mathcal{S}}\}, V_{\mathcal{M}}\big)$ result in such a way that $M_{\mathcal{M}}$ is equivalent to $\mathcal{M}$ in the following sense.

Proposition 1. *For all* $\alpha \in \mathsf{SF}$ *and* $(x, U) \in X_{\mathcal{S}}$*, we have that* $x, U \models_{\mathcal{M}} \alpha$ *iff* $M_{\mathcal{M}}, (x, U) \models \alpha$.

Here (and later on as well), the non-indexed symbol '$\models$' denotes the usual satisfaction relation of modal logic (as it was the case in Section 2 already). – The proposition is easily proved by induction on α. We call $S_{\mathcal{S}}$ and $M_{\mathcal{M}}$ the Kripke structures *induced* by $\mathcal{S}$ and $\mathcal{M}$, respectively.[2]

We now turn to the *logic* of subset spaces, LSS. Here is the appropriate axiomatization from [5], which was proved to be sound and complete in Sect. 1.2 and, respectively, Sect. 2.2 there:

1. All instances of propositional tautologies
2. $\mathsf{K}(\alpha \to \beta) \to (\mathsf{K}\alpha \to \mathsf{K}\beta)$
3. $\mathsf{K}\alpha \to (\alpha \wedge \mathsf{KK}\alpha)$
4. $\mathsf{L}\alpha \to \mathsf{KL}\alpha$
5. $(p \to \Box p) \wedge (\Diamond p \to p)$
6. $\Box (\alpha \to \beta) \to (\Box\alpha \to \Box\beta)$
7. $\Box\alpha \to (\alpha \wedge \Box\Box\alpha)$
8. $\mathsf{K}\Box\alpha \to \Box\mathsf{K}\alpha$,

where $p \in \mathsf{Prop}$ and $\alpha, \beta \in \mathsf{SF}$. – The last schema is by far the most interesting in this connection, as the interplay between knowledge and effort is captured by it. The members of this schema are called the *Cross Axioms* since [14]. Note that the schema involving only proposition variables is in accordance with the remark on Definition 1.3 above.

As the next step, let us take a brief look at the effect of the axioms from the above list within the framework of common modal logic. To this end, we consider bi-modal Kripke models $M = (X, R, R', V)$ satisfying the following four properties:

- the accessibility relation R of M belonging to the knowledge operator K is an equivalence,
- the accessibility relation R' of M belonging to the effort operator $\Box$ is reflexive and transitive,
- the composite relation $R' \circ R$ is contained in $R \circ R'$ (this is usually called the *cross property*), and
- the valuation V of M is constant along every R'-path, for all proposition variables.

Such a model M is called a *cross axiom model* (and the frame underlying M a *cross axiom frame*). Now, it can be verified without difficulty that LSS is sound with respect to the class of all cross axiom models. And it is also easy to see that every induced Kripke model is a cross axiom model (and every induced Kripke frame a cross axiom frame). Thus, the completeness of LSS for cross axiom models follows from that of LSS for subset spaces (which is Theorem 2.4 in [5]) by means of Proposition 1. This completeness result will be used below, in Section 6.

[2] It is an interesting question whether one can identify the induced Kripke structures amongst all bi-modal ones; see the paper [12] for an answer to this.

4 A Multi-agent Semantics Based on Subset Spaces

In this section, subset spaces for multiple agents are shifted into the focal point of interest. We first introduce the class of domains we consider relevant in this connection. The members of this class turn out to be slightly different from those ‘multi-agent structures’ that were taken as a basis in the paper [10]. The main difference, however, concerns the semantic atoms, into which the actual knowledge states of the agents are incorporated now. In what follows, we discuss the possible ways of interpreting knowledge or subset formulas within the new framework. We argue why we should confine ourselves to formulas from SF here, and how implicit knowledge comes into play then. Finally in this section, we prove that the logic LSS is sound with respect to the novel semantics.

For a start, we modify Definition 1 accordingly. Let $n \in \mathbb{N}$ be the number of the involved agents again.

Definition 3 (Multi-agent Subset Spaces)

1. *Let X be a non-empty set and $\mathcal{O}_i \subseteq \mathcal{P}(X)$ a set of subsets of X, for every $i \in \{1, \ldots, n\}$. Then, the tuple $\mathcal{S} = (X, \mathcal{O}_1, \ldots, \mathcal{O}_n)$ is called a* multi-agent subset frame.
2. *Let $\mathcal{S} = (X, \mathcal{O}_1, \ldots, \mathcal{O}_n)$ be a multi-agent subset frame. Then the set*

$$\mathcal{K}_{\mathcal{S}} := \{(x, U_1, \ldots, U_n) \mid x \in U_i \text{ and } U_i \in \mathcal{O}_i \text{ for all } i = 1, \ldots, n\}$$

is called the set of knowledge situations *of $\mathcal{S}$.*
3. *The notion of $\mathcal{S}$-*valuation *is the same as in Definition 1.*
4. *Let $\mathcal{S} = (X, \mathcal{O}_1, \ldots, \mathcal{O}_n)$ be a multi-agent subset frame and V an $\mathcal{S}$-valuation. Then, $\mathcal{M} := (X, \mathcal{O}_1, \ldots, \mathcal{O}_n, V)$ is called a* multi-agent subset space (based on $\mathcal{S}$)*.*

The second item of Definition 3 deserves a comment. Clearly, the *meaning* of every component of a knowledge situation remains unaltered in principle; but each individual agent is taken into account now. The *name,* however, is changed because the epistemic aspect, compared to the spatial one, comes more to the fore here.

Now, we would like to evaluate formulas in multi-agent subset spaces $\mathcal{M}$. For that purpose, let $x, U_1, \ldots, U_n$ be a knowledge situation of some multi-agent subset frame $\mathcal{S}$ (on which $\mathcal{M}$ is based). As no difficulties are raised in the propositional cases, we may proceed to the modalities directly. First, the case $\mathsf{K}_i\alpha$ is considered, where $i \in \{1, \ldots, n\}$. In order to retain the intended meaning, K_i should quantify across all the states that agent i considers possible at the world x, i.e., across U_i. But for some $y \in U_i$ it could be the case that $y, U_1, \ldots, U_n$ does not belong to $\mathcal{K}_{\mathcal{S}}$, for the simple reason that $y \notin U_j$ for some $j \in \{1, \ldots, n\}$. Thus, such a quantification is impossible in general. We conclude that we must drop formulas of the type $\mathsf{K}_i\alpha$ because of the new semantics (unless we add additional agent-specific modalities as in [10]). However, note that the knowledge of the individual agents is still represented, namely by the corresponding domains and, in particular, the semantic atoms.

Fortunately, the just detected problem does not appear in the case of the operator I. In fact, quantification now concerns states from the intersection of all the actual knowledge states; hence every such state leads to a knowledge situation as defined above. Thus, we let

$$x, U_1, \ldots, U_n \models_{\mathcal{M}} \mathsf{I}\,\alpha :\Longleftrightarrow \forall\, y \in \bigcap_{i=1,\ldots,n} U_i : y, U_1, \ldots, U_n \models_{\mathcal{M}} \alpha.$$

Regarding formulas from KF, we have got a restricted correspondence between the syntax and the semantics that way. But what can be said in the case of SF? – An easy inspection shows that the knowledge operator K causes as much a problem as K_i : because of the semantic defaults, quantifying is only possible over those states that are common to *all* of the agents. But *doing so* obviously means *turning* K *into* I. Consequently, we really *set*

$$\mathsf{K} = \mathsf{I}$$

henceforth. This is additionally justified by the fact that both operators share the same properties of knowledge (expressed by the S5-axioms).

The remaining case to be treated is that of the effort operator $\Box$. It becomes clear on second thought that this knowledge increasing modality should represent a *system component* here and may have an effect on each of the agents thus. For this reason, we define

$$x, U_1, \ldots, U_n \models_{\mathcal{M}} \Box\alpha :\Longleftrightarrow \begin{cases} \forall\, U_1' \in \mathcal{O}_1 \cdots \forall\, U_n' \in \mathcal{O}_n : \big[\, x \in U_i' \subseteq U_i \text{ for} \\ i = 1, \ldots, n \;\Rightarrow\; x, U_1', \ldots, U_n' \models_{\mathcal{M}} \alpha \,\big]. \end{cases}$$

In this way, the definition of the multi-agent semantics based on subset spaces is completed. The set SF has proved to be the relevant set of formulas, after identifying K and I.[3]

We are going to show that the logic LSS is sound for multi-agent subset spaces.

Proposition 2 (Soundness). *All formulas from* LSS *are valid in every multi-agent subset space.*

Proof. We only care about the Cross Axioms, since everything else is quite straightforward. Actually, we consider the dual schemata. Let $\mathcal{M}$ be an arbitrary multi-agent subset space and $x, U_1, \ldots, U_n$ a knowledge situation of the underlying frame. Suppose that $x, U_1, \ldots, U_n \models_{\mathcal{M}} \Diamond \mathsf{J}\,\alpha$. Then there are $U_1' \in \mathcal{O}_1, \ldots, U_n' \in \mathcal{O}_n$ such that $x \in U_i' \subseteq U_i$ for $i = 1, \ldots, n$ and $x, U_1', \ldots, U_n' \models_{\mathcal{M}} \mathsf{J}\,\alpha$. This means that there exists some $y \in \bigcap_{i=1,\ldots,n} U_i'$ for which $y, U_1', \ldots, U_n' \models_{\mathcal{M}} \alpha$. The world y is also contained in the intersection $\bigcap_{i=1,\ldots,n} U_i$. Thus, the tuple $y, U_1, \ldots, U_n$ is a knowledge situation satisfying $y, U_1, \ldots, U_n \models_{\mathcal{M}} \Diamond\alpha$. Consequently, $x, U_1, \ldots, U_n \models_{\mathcal{M}} \mathsf{J}\,\Diamond\alpha$. This proves the validity of $\Diamond \mathsf{J}\,\alpha \to \mathsf{J}\,\Diamond\alpha$. It follows that all Cross Axioms are valid in every multi-agent subset space.

[3] We retain the notation SF although we shall use I in place of K (and J in place of L, respectively) as from now.

The much more difficult question of completeness is tackled in the following section.

5 Completeness

As in the single-agent case, an infinite step-by-step construction is used for proving the completeness of LSS with respect to the new semantics, too; cf. [5], Sect. 2.2 (and [4], Sect. 4.6, for the method in general).[4] For that, it is natural to bring the canonical model of LSS into play in some way. Thus, we fix several notations concerning that model first. Let $\mathcal{C}$ be the set of all maximal LSS-consistent sets of formulas. Furthermore, let $\stackrel{\mathsf{I}}{\longrightarrow}$ and $\stackrel{\Box}{\longrightarrow}$ be the accessibility relations induced on $\mathcal{C}$ by the modalities I and $\Box$, respectively. Let $\alpha \in \mathsf{SF}$ be a non-LSS-derivable formula. Then, a multi-agent subset space falsifying α is to be built incrementally. In order to ensure that the resulting limit structure behaves as desired, several requirements on the approximations have to be met at every stage.

Suppose that $\neg\alpha \in \Gamma \in \mathcal{C}$, i.e., Γ is to be realized. We choose a denumerably infinite set of points, Y (the possible worlds of the desired model), fix an element $x_0 \in Y$, and construct inductively a sequence of quadruples $\big(X_m, (P_m^1, \ldots, P_m^n), (j_m^1, \ldots, j_m^n), (t_m^1, \ldots, t_m^n)\big)$ such that, for all $m \in \mathbb{N}$ and $i \in \{1, \ldots, n\}$,

1. X_m is a finite subset of Y containing x_0,
2. P_m^i is a finite set carrying a partial order $\leq_m^i$, with respect to which there is a least element $\bot \in X_m$,
3. $j_m^i : P_m^i \to \mathcal{P}(X_m)$ is a function such that $p \leq_m^i q \iff j_m^i(p) \supseteq j_m^i(q)$, for all $p, q \in P_m^i$,
4. $t_m^i : X_m \times P_m^i \to \mathcal{C}$ is a partial function such that, for all $x, y \in X_m$ and $p, q \in P_m^i$,
   (a) $t_m^i(x, p)$ is defined iff $x \in j_m^i(p)$; in this case it holds that
      i. if $y \in j_m^i(p)$, then $t_m^i(x, p) \stackrel{\mathsf{I}}{\longrightarrow} t_m^i(y, p)$,
      ii. if $p \leq_m^i q$, then $t_m^i(x, p) \stackrel{\Box}{\longrightarrow} t_m^i(x, q)$,
   (b) $t_m^i(x_0, \bot) = \Gamma$.

The next four conditions say to what extent the final model is approximated by the structures $\big(X_m, (P_m^1, \ldots, P_m^n), (j_m^1, \ldots, j_m^n), (t_m^1, \ldots, t_m^n)\big)$. Actually, it will be guaranteed that, for all $m \in \mathbb{N}$ and $i \in \{1, \ldots, n\}$,

5. $X_m \subseteq X_{m+1}$,
6. P_{m+1}^i is an *end extension* of P_m^i (i.e., a superstructure of P_m^i such that no element of $P_{m+1}^i \setminus P_m^i$ is strictly smaller than any element of P_m^i),
7. $j_{m+1}^i(p) \cap X_m = j_m^i(p)$ for all $p \in P_m^i$,
8. $t_{m+1}^i \mid_{X_m \times P_m^i} = t_m^i$.

[4] Due to the lack of space, we can only give a proof sketch here; however, some of the technical differences will be highlighted below.

Finally, the construction complies with the following requirements on existential formulas: for all $n \in \mathbb{N}$ and $i \in \{1, \ldots, n\}$,

9. if $\mathsf{J}\,\beta \in t^i_m(x,p)$, then there are $m < k \in \mathbb{N}$ and $y \in j^i_k(p)$ such that $\beta \in t^i_k(y,p)$,
10. if $\Diamond\beta \in t^i_m(x,p)$, then there are $m < k \in \mathbb{N}$ and $p \leq^i_k q \in P^i_k$ such that $\beta \in t^i_k(x,q)$.

With that, the final model refuting α can be defined easily. Furthermore, a relevant *Truth Lemma* (see [4], 4.21) can be proved for it, from which the completeness of LSS with respect to the multi-agent semantics follows immediately. Thus, it remains to specify, for all $m \in \mathbb{N}$, the approximating structures $\big(X_m, (P^1_m, \ldots, P^n_m), (j^1_m, \ldots, j^n_m), (t^1_m, \ldots, t^n_m)\big)$ in a way that all the above requirements are met. This makes up one of the crucial parts of the proof.

Since the case $m = 0$ is rather obvious, we focus on the induction step. Here, some existential formula contained in some maximal LSS-consistent set $t^i_m(x,p)$ must be made true, where $x \in X_m$ and $p \in P^i_m$; see item 9 and item 10 above. We confine ourselves to the case of the implicit knowledge operator. So let $\mathsf{J}\,\beta \in t^i_m(x,p)$. We choose a new point $y \in Y$ and let $X_{m+1} := X_m \cup \{y\}$. The sets $P^1_m, \ldots, P^n_m$ remain unchanged (and the associated partial orders therefore as well), i.e., we define $P^i_{m+1} := P^i_m$ for $i = 1, \ldots, n$. However, the mappings $j^1_m, \ldots, j^n_m$ are modified as follows. We let $j^i_{m+1}(q) := j^i_m(q) \cup \{y\}$, for all $q \in P^i_m$ satisfying $q \leq^i_m p$ *and all* $i \in \{1, \ldots, n\}$. The latter requirement obviously guarantees that the new point is really in the intersection of the local knowledge states. Finally, the mappings $t^1_m, \ldots, t^n_m$ are adjusted. From the *Existence Lemma* of modal logic (see [4], 4.20) we know that, for every $i \in \{1, \ldots, n\}$, there is some point Γ_i of $\mathcal{C}$ such that $t^i_m(x,p) \xrightarrow{\mathsf{I}} \Gamma_i$ and $\beta \in \Gamma_i$. Thus, we define $t^i_{m+1}(y,p) := \Gamma_i$. Moreover, the maximal consistent sets which are to be assigned to the pairs (y,q) where $q \leq^i_m p$ and $q \neq p$, are obtained by means of the cross property (which in fact holds on the canonical model); for all other pairs $(z,r) \in X_{m+1} \times P^i_{m+1}$, we let $t^i_{m+1}(z,r) := t^i_m(z,r)$. This completes the definition of $\big(X_{m+1}, (P^1_{m+1}, \ldots, P^n_{m+1}), (j^1_{m+1}, \ldots, j^n_{m+1}), (t^1_{m+1}, \ldots, t^n_{m+1})\big)$ in the case under consideration.

We must now check that the validity of the properties stated in items 1 – 8 above is transferred from m to $m+1$. Doing so, several items prove to be evident from the construction. In some cases, however, the particularities of the accessibility relations on $\mathcal{C}$ (like the cross property) have to be applied. Further details regarding this must be omitted here.

As to the validity of item 9 and item 10, it has to be ensured that *all* possible cases are eventually exhausted. To this end, processing must suitably be scheduled with regard to both modalities. This can be done with the aid of appropriate enumerations. The reader is referred to the paper [5] to see how this works in the single-agent case. – In the following theorem, the above achievements are summarized.

Theorem 2 (Completeness). *If the formula* $\alpha \in \mathsf{SF}$ *is valid in all multi-agent subset spaces, then* α *is* LSS*-derivable.*

Proposition 2 and Theorem 2 together constitute the main result of this paper, saying that LSS is sound and complete with respect to the class of all multi-agent subset spaces.

6 Remarks on Decidability and Complexity

While the soundness and the completeness of a logic depend on the underlying semantics by definition, *decidability* is a property of the logic (as a set of formulas) by itself. Thus, this property could be established by using a different semantics, and this is actually the case with LSS.

At the end of Section 2, it was stated that LSS is sound and complete for cross axiom models. In addition, LSS satisfies the *finite model property* with respect to this class of models, as was shown in [5], Sect. 2.3. Now, it is known from modal logic that both properties together imply decidability (see [4], Sect. 6.2). Thus, we can readily adopt this fact with reference to the present context.

Theorem 3 (Decidability). LSS *is a decidable set of formulas.*

And what is true of decidability is just as true of *complexity:* being a property of the logic alone. Unfortunately, the precise complexity of LSS has not yet been determined. Quite recently, the *weak logic of subset spaces,* which results from LSS by forgetting the Cross Axioms, was proved to be PSPACE-complete; see [2]. Thus, we have a partial corollary at least. The general case, however, still awaits a solution.

7 Conclusion

In this paper, we have introduced a new description of the implicit knowledge of a group of agents on the basis of subset spaces. Actually, the usual logic of implicit knowledge and Moss and Parikh's logic of subset spaces have been synthesized. The result is a novel semantics for implicit knowledge first, where the actual knowledge states of the individual agents are represented by the semantic atoms. We have argued that, relating to this framework, the implicit knowledge operator I takes over the role of the knowledge operator K from the language $\mathcal{L}$ for (single-agent) subset spaces. Thus, subset space formulas can speak about implicit knowledge and its dynamic change when interpreting them in multi-agent subset spaces.

The second outcome of this paper is a meta-theorem on the logic accompanying this non-standard semantics of I. We have proved that the logic of subset spaces, LSS (see Section 3), is sound and complete with respect to multi-agent subset spaces, too. Moreover, this logic is even decidable, which is obtained as a consequence of earlier results.

It has been argued here and there that subset spaces provide an alternative basis for reasoning about knowledge, complementing the most common and well-established epistemic logic as proposed, e.g., in [13]. It appears to us that the present paper as well makes a contribution underpinning this thesis.

An important open problem regarding LSS was addressed near the end of Section 6. Yet a lot remains to be done beyond answering that question. In particular, one should try to bring subset spaces into line with as many theoretical or practice-oriented epistemic concepts as possible, according to the thesis which has just been mentioned. (An additional justification for such a project was already indicated above, right before the final section of the introduction.)

References

1. Aiello, M., Pratt-Hartmann, I.E., van Benthem, J.F.A.K.: Handbook of Spatial Logics. Springer, Dordrecht (2007)
2. Balbiani, P., van Ditmarsch, H., Kudinov, A.: Subset space logic with arbitrary announcements. In: Lodaya, K. (ed.) ICLA 2013. LNCS, vol. 7750, pp. 233–244. Springer, Heidelberg (2013)
3. van Benthem, J.: Logical Dynamics of Information and Interaction. Cambridge University Press, Cambridge (2011)
4. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge Tracts in Theoretical Computer Science, vol. 53. Cambridge University Press, Cambridge (2001)
5. Dabrowski, A., Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. Annals of Pure and Applied Logic 78, 73–110 (1996)
6. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Synthese Library, vol. 337. Springer, Dordrecht (2007)
7. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about Knowledge. MIT Press, Cambridge (1995)
8. Georgatos, K.: Knowledge theoretic properties of topological spaces. In: Masuch, M., Polos, L. (eds.) Logic at Work 1992. LNCS (LNAI), vol. 808, pp. 147–159. Springer, Heidelberg (1994)
9. Georgatos, K.: Knowledge on treelike spaces. Studia Logica 59, 271–301 (1997)
10. Heinemann, B.: Topology and knowledge of multiple agents. In: Geffner, H., Prada, R., Machado Alexandre, I., David, N. (eds.) IBERAMIA 2008. LNCS (LNAI), vol. 5290, pp. 1–10. Springer, Heidelberg (2008)
11. Heinemann, B.: Logics for multi-subset spaces. Journal of Applied Non-Classical Logics 20(3), 219–240 (2010)
12. Heinemann, B.: Characterizing subset spaces as bi-topological structures. In: McMillan, K., Middeldorp, A., Voronkov, A. (eds.) LPAR 2013. LNCS, vol. 8312, pp. 373–388. Springer, Heidelberg (2013)
13. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge Tracts in Theoretical Computer Science, vol. 41. Cambridge University Press, Cambridge (1995)
14. Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. In: Moses, Y. (ed.) Theoretical Aspects of Reasoning about Knowledge (TARK 1992), pp. 95–105. Morgan Kaufmann, Los Altos (1992)
15. Wáng, Y.N., Ågotnes, T.: Subset space public announcement logic. In: Lodaya, K. (ed.) ICLA 2013. LNCS, vol. 7750, pp. 245–257. Springer, Heidelberg (2013)