Subset Spaces Modeling Knowledge-Competitive Agents

Bernhard Heinemann(✉)

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

Abstract. The bi-modal logic of subset spaces, LSS, was originally designed for revealing the intrinsic relationship between knowledge and topology. In recent years, it has been developed in several directions, not least towards a comprehensive knowledge-theoretic formalism. As to that, subset spaces have been shown to be smoothly combinable with various epistemic concepts, at least as long as attention is restricted to the single-agent case. Adjusting LSS to general multi-agent scenarios, however, has brought about few results only, presumably due to reasons inherent in the system. This is why one is led to consider more special cases. In the present paper, LSS is extended to a particular two-agent setting, where the peculiarity is given by the case that the agents are competitive in a sense; in fact, it is assumed here that one agent is always able to go ahead of another one regarding knowledge (or, the other one is possibly lagging behind in this respect), and vice versa. It turns out that such circumstances can be modeled in corresponding logical terms to a considerable extent.

Keywords: Reasoning about knowledge · Epistemic logic · Subset space semantics · Knowledge-competitive agents · Completeness · Decidability

1 Introduction

Our topic in this paper is *reasoning about knowledge*. This important foundational issue has been given a solid logical basis right from the beginning of the research into theoretical aspects of artificial intelligence, as can be seen, e.g., from the classic textbook [5]. According to this, a binary *accessibility relation* R_A connecting *possible worlds* or *conceivable states of the world,* is associated with every instance A of a given finite group G of agents. The *knowledge of A* is then defined through the set of all *valid formulas,* where validity is understood with regard to every state the agent considers possible at the actual one. This widespread and well-established view of knowledge is complemented by Moss and Parikh's bi-modal *logic of subset spaces,* LSS (see [11], [4], or Ch. 6 of [1]), of which the basic idea is reported in the following.

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

S. Zhang et al. (Eds.): KSEM 2015, LNAI 9403, pp. 3–14, 2015.
DOI: 10.1007/978-3-319-25159-2_1

as a *neighborhood* U of the actual state x of the world. Formulas are now 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 [5] again), and $\Box$ reflects a kind of *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*. Thus, it was natural to ask for the appropriate logic of 'real' topological spaces, which could be determined by Georgatos shortly afterwards; see [6]. The subsequent research into subset and topological spaces, respectively, is quoted in the handbook [1], whereas more recent developments include, among others, the papers [2] and [12].

Despite the fact that most treatises on LSS deal with the single-agent case, a corresponding multi-agent version was proposed in the paper [7]. The key idea behind that approach is to incorporate the agents in terms of additional modalities. This clearly leads to an essential modification of the logic, while the original semantics basically remains unchanged. However, avoiding such substantial add-ons to the logic in case of multiple agents, if at all possible, calls for restricting to special cases. It is the purpose of this paper to consider one of these.

The scenarios we are interested in here are (first and foremost) constituted by *two agents* which are *competitive* in the following sense. One of these can always surpass the other one with regard to knowledge.[1] Here, 'always' means formally: at every neighborhood situation referring to the latter. And what's good for the goose is good for the gander: a knowledge state of the first agent can always be beaten by one of the second. (Note, however, that such a notion of 'being better than' is assumed to be not necessarily strict everywhere, since otherwise it could be 'not converging' in some sense.) These ideas will be made precise below, with some new technical peculiarities coming along, opening up interesting new lines of research into logics of subset spaces.

Clearly, settings like this have a strong temporal flavor. Thus, it should be possible to model them by means of the common logic of knowledge with incorporated time as well (cf. [5], Sect. 4.3.), which is justified whenever one is obliged to focus on the chronological order most notably. But sometimes it is unnecessary or even undesirable to make time explicit. For example, with regard to certain teacher-student relationships, the *effort* of teaching and, respectively, learning in order to catch up with or even overtake the conveyer of knowledge might be rated as more important than just the amount of time it costs or the exact point of time it meets with success. We shall, therefore, introduce two-agent subset spaces in such a way that this kind of mutual consecutiveness of the agents is reflected. Our main concern is then dealing with the arising two-agent subset space logic.

[1] Thus, the use of the term 'competitive' is here different from the one that is nowadays common in the multi-agent systems community.

The rest of the paper is organized as follows. In the next section, we recapitulate the language and the logic of subset spaces for single agents. In Section 3, the scenarios of two knowledge-competitive agents, as sketched above, are formalized. In Section 4, the completeness of the resulting logic is proved. The subsequent Section 5 is devoted to the corresponding decidability problem. Finally, we summarize and make some additional remarks. – All relevant facts from modal logic not explicitly introduced here can be found in the standard textbook [3].

2 The Language and the Logic of Subset Spaces Revisited

The purpose of this section is threefold: to clarify the starting point of our investigation on a technical level, to set up some concepts and results to be introduced and, respectively, proved later on, and to enable a posterior validation of the thesis that the common single-agent case and the novel two-agent framework follow a closely related idea of knowledge; as to the latter, see the comments on Definition 4 below.

First in this section, the language for (single-agent) subset spaces, $\mathcal{L}$, is defined precisely. Then, the semantics of $\mathcal{L}$ is linked with the common relational semantics of modal logic. Finally, the ensuing relationship is utilized after the most important facts on the logic of subset spaces have been recalled.

To begin with, we define the syntax of $\mathcal{L}$. Let $\mathsf{Prop} = \{p, q, \dots\}$ be a denumerably infinite set of symbols called *proposition variables* (which shall represent the basic facts about the states of the world). Then, the set SF of all *subset formulas* over Prop is defined by the rule $\alpha \ ::= \ \top \mid p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \mathsf{K}\alpha \mid \Box\alpha$. The missing boolean connectives are treated as abbreviations, as needed. The operators which are dual to K and $\Box$ are denoted by L and $\Diamond$, respectively. In view of our remarks in the previous section, K is called the *knowledge operator* and $\Box$ the *effort operator*.

Second, we fix the semantics of $\mathcal{L}$. For a start, we single out 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.
2. *Let $\mathcal{S} = (X, \mathcal{O})$ be a subset frame. The set $\mathcal{N}_\mathcal{S} := \{(x, U) \mid x \in U$ and $U \in \mathcal{O}\}$ is then called the* set of neighborhood situations *of $\mathcal{S}$.*
3. *Let $\mathcal{S} = (X, \mathcal{O})$ be a subset frame. Under an $\mathcal{S}$-*valuation *we understand 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 bimodal 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 shows that values of proposition variables depend on states only. This is in accordance with the common practice in epistemic logic; see [5] 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 of $\mathcal{S}$. 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 here, too, defined as validity at all states that are indistinguishable to the agent.

Subset frames and subset spaces can be considered from a different perspective, as is known since [4] and reviewed in the following, for the reader's convenience. Let a subset frame $\mathcal{S} = (X, \mathcal{O})$ and a subset space $\mathcal{M} = (X, \mathcal{O}, V)$ based on it be given. Take $X_{\mathcal{S}} := \mathcal{N}_{\mathcal{S}}$ as a set of worlds, and define two accessibility relations $R_{\mathcal{S}}^{\mathsf{K}}$ and $R_{\mathcal{S}}^{\Box}$ on $X_{\mathcal{S}}$ by

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

for all $(x, U), (x', U') \in X_{\mathcal{S}}$. Moreover, let a valuation be defined by $V_{\mathcal{M}}(p) := \{(x, U) \in X_{\mathcal{S}} \mid x \in V(p)\}$, for all $p \in \mathsf{Prop}$. Then, bi-modal Kripke structures $S_{\mathcal{S}} := \left(X_{\mathcal{S}}, \{R_{\mathcal{S}}^{\mathsf{K}}, R_{\mathcal{S}}^{\Box}\}\right)$ and $M_{\mathcal{M}} := \left(X_{\mathcal{S}}, \{R_{\mathcal{S}}^{\mathsf{K}}, R_{\mathcal{S}}^{\Box}\}, V_{\mathcal{M}}\right)$ 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. – The proposition can easily be proved by structural induction on α. We call $S_{\mathcal{S}}$ and $M_{\mathcal{M}}$ the Kripke structures *induced* by the subset structures $\mathcal{S}$ and $\mathcal{M}$, respectively.

We now turn to the *logic* of subset spaces, LSS. The subsequent axiomatization from [4] 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. $\mathsf{K}\alpha \to (\alpha \wedge \mathsf{K}\mathsf{K}\alpha)$
4. $\mathsf{L}\alpha \to \mathsf{K}\mathsf{L}\alpha$
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 one, as it displays the interrelation between knowledge and effort. The members of this schema are called the *Cross Axioms* since [11]. Note that the schema involving only proposition variables is in accordance with the remark on Definition 1 above. (In other words, it is expressed by the latter schema that $\mathcal{L}$ 'only' speaks about the ongoing modification of *knowledge*.)

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 $M = (W, R, R', V)$ 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 [4]) by means of Proposition 1. This inferred completeness result can be used for proving the decidability of LSS; see [4], Sect. 2.3. We shall proceed in a similar way below, in Section 5.

3 Knowledge-Competitive Agents

The formalisms from the previous section will now be extended to the case of two knowledge-competitive agents. We again start with the logical language, which comprises two $\Box$-operators as of now. (This may appear a little surprising at first glance.) Thus, the set $\mathsf{2SF}$ of all *2-subset formulas* over Prop is defined by the rule $\alpha \ ::= \ \top \mid p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \mathsf{K}\alpha \mid \Box_1\alpha \mid \Box_2\alpha$. The above syntactic conventions apply correspondingly here. Concerning semantics, the crucial modifications follow right now.

Definition 3 (Two-Agent Subset Structures).

1. *Let X be a non-empty set and $\mathcal{O}_1, \mathcal{O}_2 \subseteq \mathcal{P}(X)$ two sets of subsets of X satisfying*
   (a) *for all $U_1 \in \mathcal{O}_1$ and every $x \in U_1$, there exists some $U_2 \in \mathcal{O}_2$ such that $x \in U_2 \subseteq U_1$, and*
   (b) *for all $U_2 \in \mathcal{O}_2$ and every $x \in U_2$, there exists some $U_1 \in \mathcal{O}_1$ such that $x \in U_1 \subseteq U_2$.*

   Then, the triple $\mathcal{S} = (X, \mathcal{O}_1, \mathcal{O}_2)$ is called a two-agent subset frame.
2. *Let $\mathcal{S} = (X, \mathcal{O}_1, \mathcal{O}_2)$ be a two-agent subset frame. The set $\mathcal{N}_\mathcal{S} := \{(x, U) \mid x \in U$ and $U \in \mathcal{O}_1 \cup \mathcal{O}_2 \cup \{X\}\}$ is then called the* set of neighborhood situations *of $\mathcal{S}$.*
3. *The notions of* $\mathcal{S}$-valuation *and* two-agent subset space *are completely analogous to those introduced in Definition 1.*

Some comments on this definition seem to be appropriate. First, the just introduced structures obviously do not correspond to the most general two-agent scenarios, but have already been adjusted to those indicated above. In fact, given that $\mathcal{O}_i$ is associated with agent i for $i \in \{1, 2\}$, condition 1.(a) says that the set of knowledge states of the first agent is 'filtered' by certain knowledge states of the second with respect to the inclusion relation; in this sense, the second agent can always increase her knowledge so that she is (at least temporarily) on par with or superior to the first. And the same applies the other way round. Thus – to say it with the example from the introduction –, the agents mutually assume the student's and the teacher's role, respectively. (This is not typical of teacher and learner in the classical understanding, but should, e.g., be kind of normal for professors and their best students.) We shall obtain simple logical counterparts to the requirements 1.(a) and 1.(b), capturing their intended meaning as just described; see below. Second, the set of all neighborhood situations of a two-agent subset frame not only is constituted of $\mathcal{O}_1$ and $\mathcal{O}_2$ but makes use of X as well. This will be advantageous for the proof of Theorem 1 below. The very fact that $\{X\}$ is written separately in Definition 3 means that the set of all states, X, is considered *indefinite,* i.e., it cannot be allocated to a particular agent.

With regard to satisfaction and validity, we need not completely present the analogue of Definition 2 here, but may confine ourselves to the clauses for the new operators.

Definition 4 (Satisfaction). *Let $\mathcal{S} = (X, \mathcal{O}_1, \mathcal{O}_2)$ be a two-agent subset frame, $\mathcal{M} = (X, \mathcal{O}_1, \mathcal{O}_2, V)$ a two-agent subset space based on $\mathcal{S}$, and $x, U \in \mathcal{N}_\mathcal{S}$ a neighborhood situation of $\mathcal{S}$ (i.e., $U \in \mathcal{O}_1 \cup \mathcal{O}_2 \cup \{X\}$). Then, for every $\alpha \in \mathsf{2SF}$,*

$$x, U \models_\mathcal{M} \Box_1\alpha :\Longleftrightarrow \forall\, U_1 \in \mathcal{O}_1 : [x \in U_1 \subseteq U \Rightarrow x, U_1 \models_\mathcal{M} \alpha]$$
$$x, U \models_\mathcal{M} \Box_2\alpha :\Longleftrightarrow \forall\, U_2 \in \mathcal{O}_2 : [x \in U_2 \subseteq U \Rightarrow x, U_2 \models_\mathcal{M} \alpha]\,.$$

Now, the knowledge of the two involved agents can be defined through the validity of knowledge formulas at the *respective* neighborhood situations; in other words, *agent 1 knows* α *at* x, U by definition, iff $x, U \models_{\mathcal{M}} \mathsf{K}\alpha$ *and* $U \in \mathcal{O}_1$, and *agent 2 knows* α *at* x, U, iff $x, U \models_{\mathcal{M}} \mathsf{K}\alpha$ *and* $U \in \mathcal{O}_2$.

These fixings clearly require justification. To this end, note that the knowledge operator K can no longer be assigned to a particular agent unambiguously. This instead happens 'externally', i.e., by means of an additional semantic condition having no direct counterpart in the object language, namely the requirement that the subset component U of the actual neighborhood situation be contained in the set of all knowledge states of the agent in question; i.e., U must have been 'enabled' for the usage of K by a preceding application of the corresponding $\Box$. The modality $\Box_i$ might therefore be called the *knowledge-enabling operator of agent* i, whereas $\mathsf{K}\alpha$ expresses that knowledge of α by agent i is actually present ($i = 1, 2$). Relating to this, it should be mentioned that all the knowledge of agents we talk about in this paper is an 'ascribed' one (cf. [5], p. 8), in fact, by the system designer utilizing epistemic logic as a formal tool for specifying multi-agent scenarios. This gives us a kind of freedom regarding the choice of the relevant system properties, which is only limited by the suitability of the approach for the intended applications. Here, the expressive power of formulas has to be restricted to some extent, on the other hand, the appearing relaxation makes it possible to describe the competitive knowledge development of the two agents under discussion; see below for some examples.

The final semantic issue to be mentioned is that *induced* Kripke structures are formed in the same way as in Section 2 here so that the two-agent analogue of Proposition 1 is obviously valid.

The *subset space logic of two knowledge-competitive agents,* $\mathsf{2LSS}$, is given by the following list of axioms, where $i, j \in \{1, 2\}$, $p \in \mathsf{Prop}$, and $\alpha, \beta \in \mathsf{2SF}$.

1. All instances of propositional tautologies
2. $\mathsf{K}(\alpha \to \beta) \to (\mathsf{K}\alpha \to \mathsf{K}\beta)$
3. $\mathsf{K}\alpha \to (\alpha \wedge \mathsf{K}\mathsf{K}\alpha)$
4. $\mathsf{L}\alpha \to \mathsf{K}\mathsf{L}\alpha$
5. $(p \to \Box_i p) \wedge (\Diamond_i p \to p)$
6. $\Box_i (\alpha \to \beta) \to (\Box_i \alpha \to \Box_i \beta)$
7. $\Box_i \alpha \to \Box_j \Box_i \alpha$
8. $\mathsf{K}\Box_i \alpha \to \Box_i \mathsf{K}\alpha$
9. $\Box_i \alpha \to \Diamond_i \alpha$

At first glance, this list is like a doubling of that for LSS; cf. Section 2. Differences arise, in particular, in the seventh schema. The original one obviously consists of two parts, $\Box\alpha \to \alpha$ and $\Box\alpha \to \Box\Box\alpha$, by means of which, regarding the relational semantics, the reflexivity and the transitivity, respectively, of the associated accessibility relation are expressed. Now, the reflexivity axiom has been separated off and weakened to *seriality* with respect to both agents; this yields the new schema 9.[2] This schema is responsible for the above mentioned

[2] A binary relation R is called *serial* iff $\forall x \exists y\, xRy$; see [5], p. 57.

filtering of $\mathcal{O}_1$ by $\mathcal{O}_2$ and vice versa (which may, therefore, be also called the *interleaving* of $\mathcal{O}_1$ and $\mathcal{O}_2$). Thus, this is the point where, compared to LSS, one of the crucial changes appears: that weakening of $\Box$ allows for interpreting the $\Box_i$'s in knowledge-competitive scenarios as described above; and Theorem 1 below can, in fact, be proved with this.

On the other hand, the new schema 7 comprises, in particular, two-agent transitivity, by letting $i = j$. In case $i \neq j$, however, additional new axioms appear, mirroring the fact that the interleaving of $\mathcal{O}_1$ and $\mathcal{O}_2$ is compatible with the subset space structure of the respective collections of knowledge states; see the proof of Proposition 2 (and also the comments on Definition 5) below.

As to examples of derived 2LSS-sentences, let us call a formula α *i-stable* in a two-agent subset space, iff $\mathsf{K}\Box_i\alpha$ is valid there. Since $\mathsf{K}\Box_i\alpha$ implies $\Box_i\mathsf{K}\alpha$, i-stability means, in particular, that agent i knows α with a kind of future certainty. Moreover, it can be asserted that i-stable formulas will in actual fact be known, because $\Diamond_i\mathsf{K}\alpha$ can be deduced from $\mathsf{K}\Box_i\alpha$. Going beyond that, it can easily be shown that i-stability itself is stable knowledge of each of the two agents, or, to put it another way, $\mathsf{K}\Box_i\alpha \rightarrow \mathsf{K}\Box_j\mathsf{K}\Box_i\alpha$ belongs to 2LSS for $i, j \in \{1, 2\}$. All this can easily be obtained with the aid of the above axioms.

Concluding our discussion on the arising logic, we would like to draw the reader's attention to the multi-method logics of subset spaces examined in [8], Sect. 4. Despite all the differences in details, those are formally more similar to the present approach than the multi-agent version of LSS quoted in the introduction.

Finally in this section, it is proved that the logic 2LSS is *sound* with respect to the class of all two-agent subset spaces.

Proposition 2. *Let $\mathcal{M} = (X, \mathcal{O}_1, \mathcal{O}_2, V)$ be a two-agent subset space. Then, every axiom from the above list is valid in $\mathcal{M}$.*

Proof. We confine ourselves to the instance $\Box_1\alpha \rightarrow \Box_2\Box_1\alpha$ of the seventh schema. Let $x, U \models_{\mathcal{M}} \Box_1\alpha$ be satisfied. This means that, for all $U_1 \in \mathcal{O}_1$ such that $x \in U_1 \subseteq U$, we have $x, U_1 \models_{\mathcal{M}} \alpha$. Now, let $U_2 \in \mathcal{O}_2$ be any subset of U containing x. Furthermore, let $U' \in \mathcal{O}_1$ be an arbitrary element satisfying $x \in U' \subseteq U_2$. Then, in particular, $U' \subseteq U$. Thus, $x, U' \models_{\mathcal{M}} \alpha$. It follows that $x, U_2 \models_{\mathcal{M}} \Box_1\alpha$. Consequently, $x, U \models_{\mathcal{M}} \Box_2\Box_1\alpha$, since U_2 was chosen arbitrarily as well. This proves (the particular case of) the proposition.

As the transitivity of the subset relation is crucially used for the preceding proof, one may call the just treated schemata the *quasi-transitivity axioms.*

4 Completeness

In this section, we primarily present the new concepts required for proving the semantic completeness of 2LSS on the class of all two-agent subset spaces. As it is mostly the case with subset space logics, the overall structure of such a proof consists of an infinite step-by-step model construction.[3] Utilizing a procedure of that kind seems to be necessary, since subset spaces in a sense do not harmonize with the main modal means supporting completeness, viz *canonical models*.

The canonical model of 2LSS will come into play nevertheless. So let us fix some notations concerning that model first. Let $\mathcal{C}$ be the set of all maximal 2LSS-consistent sets of formulas. Furthermore, let $\xrightarrow{\mathsf{K}}$ and $\xrightarrow{\Box_i}$ be the accessibility relations induced on $\mathcal{C}$ by the modalities K and $\Box_i$, respectively, where $i \in \{1, 2\}$. And finally, let $\alpha \in 2\mathsf{SF}$ be a formula which is *not* contained in 2LSS. Then, we have to find a model for $\neg\alpha$.

This model is constructed stepwise and incrementally in such a way that better and better intermediary structures are obtained (which means that more and more existential formulas are realized). In order to ensure that the finally resulting limit structure behaves as desired, several requirements on those approximations have to be met at every stage. This makes up the technical core of the proof, of which only the outset is specified in detail below. With this aim in view, we need a definition.

Definition 5 (Almost Partially Ordered Sets). *Let P be a non-empty set.*

1. *A binary relation $\sqsubset$ on P is called* weakly trichotomous *iff, for all $\pi, \rho \in P$, at most one out of $(\pi \sqsubset \rho, \rho \sqsubset \pi, \pi = \rho)$ is true. Now, $(P, \sqsubset)$ is called an* almost partially ordered (apo) set *and $\sqsubset$ an* almost partial order on P, *iff $\sqsubset$ is transitive and weakly trichotomous.*
2. *Let $\sqsubset_1, \sqsubset_2$ be almost partial orders on P. Then, $(P, \sqsubset_1, \sqsubset_2)$ is called a* twofold almost partially ordered (2apo) set *iff, for all $\pi, \rho, \sigma \in P$ and $i, j \in \{1, 2\}$, it ensues from $\pi \sqsubset_i \rho \sqsubset_j \sigma$ that $\pi \sqsubset_j \sigma$.*

We comment on 5 first. Compared against partial orders, reflexivity is obviously missing, whereas antisymmetry ensues from weak trichotomy. – Concerning 5, it should be remarked that not only are 2apo-sets those equipped with two almost partial orders, but these satisfy an additional requirement corresponding, among other things, to the fact that the interleaving of the two distinguished sets of subsets of a two-agent subset frame is 'good-natured' with respect to inclusion. Actually, this condition is the semantic equivalent of the quasi-transitivity axioms in case $i \neq j$, hence itself is called *quasi-transitivity (of* $(\sqsubset_1, \sqsubset_2)$*)*.

We now describe the ingredients of the above mentioned approximation structures. Their possible worlds are successively taken from a denumerably infinite set of points, Y, chosen in advance. Also, another denumerably infinite set, Q, is chosen such that $Y \cap Q = \emptyset$. The latter set shall gradually contribute to a 2apo-set representing the subset space structure of the desired limit model. Finally, we fix particular 'starting elements' $x_0 \in Y$, $\bot \in Q$, and $\Gamma \in \mathcal{C}$ containing the formula $\neg\alpha$ from above. Then, a sequence of quadruples (X_m, P_m, j_m, t_m) has to be defined inductively such that, for all $m \in \mathbb{N}$,

[3] See [4], Sect. 2.2, for a fully completed proof regarding LSS, and [9], Sect. 5, for an outline of a particular variation.

- X_m is a finite subset of Y containing x_0,
- P_m is a finite subset of Q containing $\bot$ and carrying two almost partial orders $\sqsubset_1, \sqsubset_2$ such that
  - $(P_m, \sqsubset_1, \sqsubset_2)$ is a 2apo-set and
  - $\bot$ is the *least element* in P_m (i.e., $\bot \sqsubset_1 \pi$ or $\bot \sqsubset_2 \pi$ for all $\pi \in P_m$),
- $j_m : P_m \to \mathcal{P}(X_m)$ is a function satisfying $(\pi \sqsubset_i \rho \iff j_m(\pi) \supset j_m(\rho))$, for all $\pi, \rho \in P_m$ and $i \in \{1, 2\}$, and
- $t_m : X_m \times P_m \to \mathcal{C}$ is a *partial* function such that, for all $x, y \in X_m$ and $\pi, \rho \in P_m$,
  - $t_m(x, \pi)$ is defined iff $x \in j_m(\pi)$; in this case it holds that
    - if $y \in j_m(\pi)$, then $t_m(x, \pi) \xrightarrow{\mathsf{K}} t_m(y, \pi)$,
    - if $\pi \sqsubset_i \rho$, then $t_m(x, \pi) \xrightarrow{\Box_i} t_m(x, \rho)$, where $i \in \{1, 2\}$,
  - $t_m(x_0, \bot) = \Gamma$.

By the way, the intermediary sets $\mathcal{O}_1^m$ and $\mathcal{O}_2^m$ of subsets of X_m are obtained from that as follows: $\mathcal{O}_i^m := \{j_m(\pi) \mid \pi \in P_m \text{ and } \pi \text{ has a } \sqsubset_i\text{-predecessor}\}$, for $i = 1, 2$.

During the construction indicated above, the sets X_m and P_m must be enlarged with new elements and the mappings j_m and t_m correspondingly be extended in each step. It turns out that this plan can indeed be followed faithfully. All this finally yields the subsequent theorem.

Theorem 1 (Completeness). *Let $\alpha \in \mathsf{2SF}$ be a formula which is valid in all two-agent subset spaces. Then α belongs to the logic $\mathsf{2LSS}$.*

5 Decidability

The standard method for proving the decidability of a given modal logic is *filtration,* which restricts inspection of the relevant models to the finite ones among them and enables a decision procedure thus. However, just as subset spaces do not harmonize with canonical models, they are incompatible with filtration. Thus, a detour is required, which takes us back into the relational semantics. In the following, we shall single out a class of tri-modal Kripke structures for which $\mathsf{2LSS}$ is as well sound and complete, and which is closed under filtration in a suitable manner. This will give us the desired decidability result. Subsequently, K is supposed to correspond to R, and $\Box_i$ to R_i' for $i = 1, 2$.

Definition 6 (Two-Agent Model). *Let $M := (W, R, R_1', R_2', V)$ be a tri-modal Kripke model, where $R, R_1', R_2' \subseteq W \times W$ are binary relations and V is a valuation. Then M is called a* two-agent model, *iff the following conditions are satisfied.*

1. *R is an equivalence relation,*
2. *both R_1' and R_2' are serial and transitive,*
3. *both pairs (R_1', R_2') and (R_2', R_1') satisfy the quasi-transitivity condition,*
4. *both pairs (R, R_1') and (R, R_2') satisfy the cross property, and*

5. *the valuation V of M is constant along every R'_i-path, for $i \in \{1, 2\}$ and all proposition variables.*

The class of all Kripke models induced by a two-agent subset space is contained in the class of all two-agent models, as can be seen easily. It follows that 2LSS is (sound and) complete with respect to the latter class; see the final part of Section 2 above. Therefore, it remains to be proved that this class is closed under filtration.

For this purpose, let a 2LSS-consistent formula $\alpha \in$ 2SF be given. Then, a *filter set* of formulas, involving the set $\mathrm{sf}(\alpha)$ of all subformulas of α, is defined as follows. We start off with $\Sigma_0 := \mathrm{sf}(\alpha) \cup \{\neg\beta \mid \beta \in \mathrm{sf}(\alpha)\}$. In the next step, we take the closure of Σ_0 under finite conjunctions of pairwise distinct elements of Σ_0. After that, we close under single applications of the operator L. And finally, we join the sets of subformulas of all the elements of the set obtained last. (This final step is necessary because L was introduced as an abbreviation.) The resulting set of formulas, denoted by Σ, is the one that meets the current requirements. Note that Σ is *finite*.

Now, the canonical model of 2LSS is filtered through Σ. As a filtration of the corresponding accessibility relations, we take the *smallest* one in each of the three cases. Let $M := (W, R, S_1, S_2, V)$ be the resulting model, where the valuation V shall be in accordance with Definition 6 for the proposition variables outside of Σ. Then, the following lemma is crucial.

Lemma 1. *The structure M is a finite two-agent model. Furthermore, the size of M can be computed from the length of α.*

Proof. The finiteness of W follows from that of Σ, and we must now show that the five conditions from Definition 6 are satisfied. According to the way the filter set Σ was formed, the verification of 1 and 4 is not difficult. Next, both the validity of 5 for the proposition variables occurring in Σ and the seriality of the R'_is can easily be concluded from the fact that M is the result of a filtration. Moreover, establishing the transitivity of the R'_is is covered by the proof of Lemma 2.10 from [4]. Thus only the verification of the third condition requires a separate argument. Fortunately, it turns out that one can proceed for it in a way similar to the one taken for transitivity (but this is the most sophisticated portion of the whole proof). In this manner, the lemma is proved.

The desired decidability result is now an immediate consequence of Lemma 1 and the facts stated above.

Theorem 2 (Decidability). *The logic* 2LSS *is a decidable set of formulas.*

6 Conclusion

In this paper, a subset space logic of two knowledge-competitive agents, denoted by 2LSS, has been introduced. A corresponding axiomatization was proposed, which turned out to be *sound and complete* with respect to the intended class of models. This constitutes the first of our main results. The second assures the *decidability* of the new logic.

It is to be expected that the *complexity* of 2LSS can be determined not until solving this problem for the usual logic of subset spaces. As to that, only partial results are known; see [2].

Generalizing our approach to the case of *more than two agents* does not pose difficulties on the formal side. However, the interpretation of some of the competitively relevant formulas is different then; for example, seriality in the more general context means that all agents different from a particular one can do better than the latter. Finally, the question for *other interesting agent interrelationships* and the effects of them on knowledge comes up and should be answered by future research; concerning this, the paper [10] may serve as a starting point.

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.) Logic and Its Applications. LNCS, vol. 7750, pp. 233–244. Springer, Heidelberg (2013)
3. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic, Cambridge Tracts in Theoretical Computer Science, vol. 53. Cambridge University Press, Cambridge (2001)
4. Dabrowski, A., Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. Annals of Pure and Applied Logic **78**, 73–110 (1996)
5. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about Knowledge. MIT Press, Cambridge (1995)
6. Georgatos, K.: Knowledge theoretic properties of topological spaces. In: Masuch, M., Polos, L. (eds.) Logic at Work 1992. LNCS, vol. 808, pp. 147–159. Springer, Heidelberg (1994)
7. 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)
8. Heinemann, B.: Logics for multi-subset spaces. Journal of Applied Non-Classical Logics **20**(3), 219–240 (2010)
9. Heinemann, B.: Coming upon the classic notion of implicit knowledge again. In: Buchmann, R., Kifor, C.V., Yu, J. (eds.) KSEM 2014. LNCS, vol. 8793, pp. 1–12. Springer, Heidelberg (2014)
10. Lomuscio, A., Ryan, M.: Ideal agents sharing (some!) knowledge. In: Prade, H. (ed.) 13th European Conference on Artificial Intelligence, ECAI 1998, pp. 557–561. John Wiley & Sons Ltd, Chichester (1998)
11. 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)
12. Wáng, Y.N., Ågotnes, T.: Subset space public announcement logic. In: Lodaya, K. (ed.) Logic and Its Applications. LNCS, vol. 7750, pp. 245–257. Springer, Heidelberg (2013)