Multi-Agent Subset Space Logic

Yì N. Wáng
Dept. of Computing, Mathematics and Physics
Bergen University College, Norway
wonease@gmail.com

Thomas Ågotnes
Dept. of Information Science and Media Studies
University of Bergen, Norway
thomas.agotnes@infomedia.uib.no

Abstract

Subset space logics have been introduced and studied as a framework for reasoning about a notion of *effort* in epistemic logic. The seminal Subset Space Logic (SSL) by Moss and Parikh modeled a single agent, and most work in this area has focused on different extensions of the language, or different model classes resulting from restrictions on subset spaces, while still keeping the single-agent assumption. In this paper we argue that the few existing attempts at multi-agent versions of SSL are unsatisfactory, and propose a new multi-agent subset space logic which is a natural extension of single-agent SSL. The main results are a sound and complete axiomatization of this logic, as well as an alternative and equivalent relational semantics.

1 Introduction

The study of subset space logic (SSL) was initiated in [Moss and Parikh, 1992]. One of the main motivations of this logic is to characterize epistemic effort in a reasonably simple framework which, for various reasons, has been chosen to be a topological one. In SSL's *subset models* an agent is assumed to have an associated *epistemic range*, a set of points. The epistemic range is not fixed: it can shrink as the result of an *effort* by the agent. Knowledge of φ is interpreted as φ being true at all points in the current epistemic range, and thus an effort can change the agent's knowledge or ignorance. The set of possible epistemic ranges is called a *subset space*.

Consider the following example from [Parikh *et al.*, 2007]. A policeman holds a radar gun to test speeding on a motorway. The accuracy of the radar gun is ± 3 km/h, and the speed limit of the motorway is 120 km/h. Now a car is coming, and the radar gun shows it is running at 122 km/h. Is the car speeding or not? The policeman is not certain. However, if he used a different radar gun with accuracy ± 1 km/h and read 122 km/h, he would be certain. This is captured by the SSL formula $\neg K(Speeding) \wedge \neg\Box\neg K(Speeding)$. The first conjunct says that the policeman does not know, the second that it *is possible* for him to know. The language of SSL is simply propositional logic extended with the modalities K and $\Box$.

Further investigations of SSL consider, e.g., restrictions on subset spaces [Georgatos, 1994; Dabrowski *et al.*, 1996; Georgatos, 1997], or add names for points and sets to the language, making it a *hybrid logic* [Heinemann, 2003; 2004; 2008b; 2008a; Wang, 2009].

Despite the simplicity and flexibility of SSL, other dynamic epistemic approaches, such as public announcement logic [Plaza, 1989] and action model logic [Baltag and Moss, 2004], have received more attention in recent years. One reason might be that no satisfactory multi-agent version of SSL have existed. In this paper we aim to improve this situation.

Some attempts at multi-agent versions of SSL have been made [Başkent, 2007; Heinemann, 2008b; 2010; Wen *et al.*, 2011]. However, [Başkent, 2007] and [Wen *et al.*, 2011] both, first, have problems with the semantics of nested epistemic formulas (see Sections 3 and 6), and, second, have no known meta-logical results like completeness. [Heinemann, 2008b; 2010], on the other hand, have completeness results for logics that take multiple agents into account, but it can be argued that these approaches are not multi-agent extensions of SSL in the way standard multi-agent epistemic logic [Fagin *et al.*, 1995] extends single-agent epistemic logic, neither syntactically nor semantically. Syntactically, these approaches add additional operators (hybrid-like, in the case of [Heinemann, 2008b]) instead of just replacing the K-operator with a K_i-operator for each agent i as in standard epistemic logic (and possibly similarly for $\Box$). Semantically, the K_i-operators do not have the S5 properties of knowledge.

Our motivation in this paper is to study the natural and conservative extension of the original SSL language obtained by "agentizing" the K and $\Box$ modalities, without adding hybrid operators or other additional modalities, under the assumption that agents have S5 knowledge (as in the original SSL). Similarly to (single-agent) SSL, the language of our multi-agent SSL extends propositional logic with modalities K_i and $\Box_i$ for each agent i. This language is also used in, e.g., [Başkent, 2007]. We define a semantics for multi-agent SSL, and give an axiomatization that we prove is sound and complete.

We also observe and formally show that (multi-agent) SSL has a relational aspect in the sense that every subset model corresponds to a certain multi-dimensional relational model. This observation brings SSL closer to the tradition of classical modal logic based on relational semantics. As a result, we are able to make use of well-known model theory of classical modal logic, which is helpful for further research in this field.

The idea of treating a subset model as a multi-dimensional relational model is not new: a similar approach is used in the single-agent case in [Dabrowski *et al.*, 1996]. However, their result is weaker (not a one-to-one correspondence) and the technique used in this paper (for the multi-agent case) is fundamentally different. Dabrowski *et al.* [1996] also give a completeness proof for single-agent SSL, but it uses a proof technique that seems difficult to apply directly to the multi-agent case or other extensions of the logic. In this paper we present a new and detailed completeness proof for the multi-agent case, which is based on its relational aspect and standard modal logic techniques (the step-by-step method).

The paper is organized as follows. In the next section we briefly recall basic subset space logic. In Section 3 we propose a language and semantics for multi-agent SSL and present a sound axiomatization. In Section 4 we introduce relational semantics for multi-agent SSL, and show its correspondence to multi-agent subset semantics. The completeness proof is found in Section 5. We discuss related and future work and conclude in Section 6.

2 Background: Single-Agent SSL

We briefly recall the main concepts of (single-agent) SSL, see [Moss and Parikh, 1992; Dabrowski *et al.*, 1996] for details. Let PROP be a countable set of propositional variables.

Definition 1 (Single-agent language) The language $\mathcal{L}$ is given by the following grammar rule, where $p \in$ PROP:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K\varphi \mid \Box\varphi,$$

We write $\hat{K}\varphi$ for $\neg K \neg\varphi$, and $\Diamond\varphi$ for $\neg\Box\neg\varphi$.

$K\varphi$ is intended to mean that "φ is known", and $\Diamond\varphi$ that "there is a refinement of knowledge (e.g., new evidence) under which φ is true". Formally the semantics is defined as follows.

Definition 2 (Subset structures) A pair $(X, \mathcal{O})$ is called a *subset space*, if X is a non-empty set and $\mathcal{O} \subseteq \wp(X)$. A *subset model* is a tuple $\mathcal{X} = (X, \mathcal{O}, V)$ where $(X, \mathcal{O})$ is a subset space and V : PROP $\to \wp(X)$ is an valuation function. Elements of $\mathcal{O}$ are called *open sets*.

An *epistemic scenario* of a subset model $\mathcal{X} = (X, \mathcal{O}, V)$ is a pair (x, O) where $x \in X$ is the factual state, $O \in \mathcal{O}$ is the *epistemic range* or *evidence*, and where $x \in O$.

Definition 3 Let $\mathcal{X} = (X, \mathcal{O}, V)$ be a subset model, and (x, O) an epistemic scenario of $\mathcal{X}$. The satisfaction relation $\models$ is given as follows:

$\mathcal{X}, x, O \models p$ iff $x \in V(p)$
$\mathcal{X}, x, O \models \neg\varphi$ iff $\mathcal{X}, x, O \not\models \varphi$
$\mathcal{X}, x, O \models \varphi \wedge \psi$ iff $\mathcal{X}, x, O \models \varphi$ & $\mathcal{X}, x, O \models \psi$
$\mathcal{X}, x, O \models K\varphi$ iff $\forall y \in O. \mathcal{X}, y, O \models \varphi$
$\mathcal{X}, x, O \models \Box\varphi$ iff $\forall U \in \mathcal{O}. (x \in U \subseteq O \Rightarrow \mathcal{X}, x, U \models \varphi)$

SSL, an (equivalent variant of a) sound and complete axiomatization [Moss and Parikh, 1992] of subset space logic is given in Figure 1.

Instances of tautologies	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
$K\varphi \to \varphi$	$\neg K\varphi \to K\neg K\varphi$
$\vdash \varphi \Rightarrow \vdash K\varphi$	$\Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi)$
$\Box\varphi \to \varphi$	$\Box\varphi \to \Box\Box\varphi$
$\vdash \varphi \Rightarrow \vdash \Box\varphi$	$(p \to \Box p) \wedge (\neg p \to \Box\neg p)$
$K\Box\varphi \to \Box K\varphi$	$\vdash \varphi$ & $\vdash \varphi \to \psi \Rightarrow \vdash \psi$

Figure 1: The axiomatization **SSL** of subset space logic.

3 Multi-Agent Subset Space Logic

Things become more interesting and more complicated when more agents are involved. Consider the following example:

Example 1 Alice is driving a car on a motorway with a speed limit of 120 km/h. The speedometer displays the speed to an accuracy of ± 5 and it reads 120 km/h at the moment. She knows that the police use a radar gun with accuracy ± 1 km/h. Now she is stopped by a policeman with a radar gun which reads 122 km/h (which Alice does not see). The car was running at 121.5 km/h. The following questions are of interest: (i) Does the policeman know that Alice was speeding? (ii) Does Alice know that the policeman knows that she was speeding? (iii) Does Alice know that the policeman knows that she knows she was speeding? (iv) Alice can replace her speedometer with a new one with an accuracy of $0 \sim 5\%$ (the reading is never lower than the actual speed). Is there anything Alice can do to make sure she knows she was speeding whenever the policeman knows she was speeding (so that she will never get into trouble)?

The first question is a simple statement of the type that was mentioned in the introduction, which single-agent SSL can be used to reason about. The other questions, however, are multi-agent ones. The fourth question relates to knowledge update by *effort*. In multi-agent SSL, we must be able to reason about nested knowledge, like in questions 2–4.

A natural modeling attempt is to define a multi-agent subset model as a structure $\mathcal{X} = (X, \mathcal{O}_0, \ldots, \mathcal{O}_{n-1}, V)$ which differs from a single-agent subset model only in that it includes a set of open sets for each agent. An epistemic scenario becomes $(x, O_0, \ldots, O_{n-1})$ such that $x \in O_0 \cap \cdots \cap O_{n-1}$. Then, the definition of satisfaction can be extended in the natural way. This attempt appears in [Wen *et al.*, 2011, Section 4.2]. However, there is a problem: the truth value of a nested formula of the form $K_i K_j \varphi$ is in fact not well defined. For consider in the above example that Alice knows that the policeman knows that she was speeding, denoted by $K_0 K_1(Speeding)$. To verify $\mathcal{X}, x, O_0, \ldots, O_{n-1} \models K_0 K_1(Speeding)$, we need to check whether $\mathcal{X}, y, O_0, \ldots, O_{n-1} \models K_1(Speeding)$ for all $y \in O_0$. Now, observe that $y \in O_0$ does not necessarily mean that $y \in O_1$, which means that $(y, O_0, \ldots, O_{n-1})$ is not necessarily an epistemic scenario as defined above. But we do need the condition $y \in O_1$, to interpret $K_1(Speeding)$ (it makes no sense in SSL to have a factual state outside an epistemic range).

Başkent [2007] proposes two different semantics for the language of multi-agent SSL based on the above defined multi-agent subset models, which avoid the problem mentioned above. In the first of these, knowledge is interpreted in the following way: $\mathcal{X}, x, O_0, \ldots, O_{n-1} \models K_i\varphi$ iff $\forall y \in O_0 \cap \cdots \cap O_{n-1}.\ \mathcal{X}, y, O_0, \ldots, O_{n-1} \models \varphi$. This does avoid the problem that the factual state falls outside of an epistemic range, however, with the cost that any formula of the form $K_iK_j\varphi \leftrightarrow K_jK_i\varphi$ is valid, which is unlike in standard epistemic logic and arguably gives a much less intuitive and interesting logic. The second approach in [Başkent, 2007] has a similar problem.

A natural solution to these problems is to increase the order: to define a multi-agent subset space as a tuple $(X, \mathcal{O})$ for each agent, where $\mathcal{O} \subseteq \wp(\wp(X))$ (instead of $\mathcal{O} \subseteq \wp(X)$ as for single-agent subset space), such that $\mathcal{O}$ forms a set of partitions of X. An update of an individual's knowledge can then be modeled as refinement of a partition for that agent. This setting will be referred to as *partition semantics*. As we shall see, partition semantics allows us to model multi-agent nested S5 knowledge in the correct way. In the single-agent case, the relationship between partition semantics and the classical SSL semantics is just like the relationship between using equivalence classes and using universal relations in a relational model for interpreting the modal logic S5; see [Hughes and Cresswell, 1996, p. 61]. It is not hard to see that single-agent SSL is equivalently interpreted in either semantics, for in the partition semantics: i) since each agent has a set of partitions there can be multiple open sets available to model different epistemic scenarios, and ii) all the equivalence classes other than the current one in every partition are redundant. To sum up, our proposed partition semantics is equivalent to the standard SSL subset semantics for single-agent SSL but is suitable for multi-agent SSL as well.

We now formally define the language, semantics and axiomatization of multi-agent SSL. They are very natural extensions of single-agent SSL. After that, we introduce an equivalent *relational* semantics for multi-agent SSL, which we use to prove completeness.

Language and Semantics

We still let PROP be a countable set of propositional variables. Let N be the set of all agents. We assume that N is a finite non-empty set, and abuse notation by both treating it as a natural number $\mathrm{N} \geq 1$ and a set $\mathrm{N} = \{0, \ldots, \mathrm{N} - 1\}$.

Definition 4 (Multi-agent language) The language $\mathcal{L}_m$ is given by the following rule, where $p \in$ PROP and $i \in$ N:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid \Box_i\varphi,$$

As usual, we write $\hat{K}_i\varphi$ for $\neg K_i\neg\varphi$, and $\Diamond_i\varphi$ for $\neg\Box_i\neg\varphi$.

We now propose a semantics for this language in terms of multi-agent subset models, based on partition semantics.

Definition 5 (Multi-agent subset models) A *multi-agent subset model* is a tuple $(X, \mathcal{O}, V)$ where

- X is a non-empty set;
- $\mathcal{O} : \mathrm{N} \to \wp(\wp(\wp(X)))$ assigns to every agent a set of partitions of X; we write $\mathcal{O}_i$ for $\mathcal{O}(i)$;
- V : PROP $\to \wp(X)$ is an valuation function.

The set of all multi-agent subset models is denoted by $\mathfrak{Sub}$.

Given a multi-agent subset model $\mathcal{X} = (X, \mathcal{O}, V)$, a *profile* of $\mathcal{X}$ is a function $O : \mathrm{N} \to \wp(\wp(X))$ such that $O(i) \in \mathcal{O}_i$. We write O_i for $O(i)$. Here, an *epistemic scenario* is a pair (x, O) such that $x \in X$ and O is a profile. We introduce the following notation when O and U are profiles:

$[x]_{O_i}$	$=_{\mathrm{def}}$	the equivalence class in O_i that contains x
$U_i \preceq O_i$	$\Leftrightarrow_{\mathrm{def}}$	U_i is a finer partition than O_i, i.e., a member of O_i is the union of some members of U_i
$U \preceq_i O$	$\Leftrightarrow_{\mathrm{def}}$	$U_i \preceq O_i$ & $\forall j \in \mathrm{N}.(j \neq i \Rightarrow U_j = O_j)$.

The notation $[x]_{O_i}$ will be used in a more general form: $[x]_P$, where P is any partition. Moreover, the relation $\preceq$ is a partial order as the notation already suggests.

Definition 6 (Multi-agent subset semantics) Let $\mathcal{X}$ be a multi-agent subset model, and (x, O) an epistemic scenario of $\mathcal{X}$:

$\mathcal{X}, x, O \models p$	iff	$x \in V(p)$
$\mathcal{X}, x, O \models \neg\varphi$	iff	$\mathcal{X}, x, O \not\models \varphi$
$\mathcal{X}, x, O \models \varphi \wedge \psi$	iff	$\mathcal{X}, x, O \models \varphi$ & $\mathcal{X}, x, O \models \psi$
$\mathcal{X}, x, O \models K_i\varphi$	iff	$\forall y \in [x]_{O_i}.\ \mathcal{X}, y, O \models \varphi$
$\mathcal{X}, x, O \models \Box_i\varphi$	iff	$\forall U \preceq_i O.\ \mathcal{X}, x, U \models \varphi.$

We write $\mathcal{X} \models \varphi$, if $\mathcal{X}, (x, O) \models \varphi$ holds for all epistemic scenarios (x, O). In a similar fashion, we can define the validity of φ in a multi-agent subset space $(X, \mathcal{O})$ (notation: $X, \mathcal{O} \models \varphi$), and so on.

As we have already argued, when $\mathrm{N} = 1$, the multi-agent subset semantics is equivalent to the original subset semantics.

Let us formalize Example 1. Let 0 and 1 be Alice and the policeman respectively. Let $\mathcal{X} = (\mathbb{R}^+, \mathcal{O}, V)$, with $\mathcal{O}_0 = \{O_0, U_0\}$ and $\mathcal{O}_1 = \{O_1\}$, where $O_0 = \{(114, 126), \mathbb{R}^+ \setminus (114, 126)\}$, $O_1 = \{(121, 123), \mathbb{R}^+ \setminus (121, 123)\}$, and $U_0 = \{(114, 120), [120, 126), \mathbb{R}^+ \setminus (114, 126)\}$; and V such that $V(\mathit{Speeding}) = (120, +\infty)$. The current epistemic scenario is $(121.5, \{\langle 0, O_0\rangle, \langle 1, O_1\rangle\})$. Whether the formulas $K_1(\mathit{Speeding})$, $K_0K_1(\mathit{Speeding})$, $K_0K_1K_0(\mathit{Speeding})$ and $\Diamond_0(K_1(\mathit{Speeding}) \to K_0(\mathit{Speeding}))$ are true in this situation represents the questions in Example 1. It is an easy exercise to check that these are true, false, false, and true, respectively. For the truth of the last formula, U_0 represents the result of an effort (replacing the speedometer) by which $K_0(\mathit{Speeding})$ becomes true.

Axiomatization

The axiomatization $\mathbf{SSL_m}$ for multi-agent subset space logic is given in Figure 2. $\mathbf{SSL_m}$ contains multi-agent versions of all axioms and rules of **SSL**, in addition to three axiom schemata for characterizing (non-)interaction between dimensions. They are in two groups: COM for *commutativity*, and ChR for *Church-Rosser*. We do not need $\mathrm{ChR}^{\bullet}_{\circ}$, i.e. $\hat{K}_i\Box_j\varphi \to \Box_j\hat{K}_i\varphi$ where $j \neq i$, as an axiom, because it can be derived in $\mathbf{SSL_m}$. It is clear that $\mathbf{SSL_m}$ coincides with **SSL** when restricted to a single agent.

Theorem 1 (Soundness) *Every $\mathbf{SSL_m}$-theorem is valid in the class of all multi-agent subset models.*

4 Relational Correspondence

We now define an equivalent relational semantics for multi-agent SSL. We first define a class of relational models called

(PC)	All instances of tautologies
(AP)	$(p \to \Box_i p) \land (\neg p \to \Box_i \neg p)$
(K•)	$K_i(\varphi \to \psi) \to (K_i\varphi \to K_i\psi)$
(K∘)	$\Box_i(\varphi \to \psi) \to (\Box_i\varphi \to \Box_i\psi)$
(T•)	$K_i\varphi \to \varphi$
(T∘)	$\Box_i\varphi \to \varphi$
(5•)	$\neg K_i\varphi \to K_i\neg K_i\varphi$
(4∘)	$\Box_i\varphi \to \Box_i\Box_i\varphi$
(CR)	$K_i\Box_i\varphi \to \Box_i K_i\varphi$
($\mathrm{ChR}^{\circ}_{\circ}$)	$\Diamond_i\Box_j\varphi \to \Box_j\Diamond_i\varphi$, if $i \neq j$
($\mathrm{COM}^{\bullet}_{\circ}$)	$K_i\Box_j\varphi \leftrightarrow \Box_j K_i\varphi$, if $i \neq j$
($\mathrm{COM}^{\circ}_{\circ}$)	$\Box_i\Box_j\varphi \to \Box_j\Box_i\varphi$
(N•)	$\vdash \varphi \Rightarrow \vdash K_i\varphi$
(N∘)	$\vdash \varphi \Rightarrow \vdash \Box_i\varphi$
(MP)	$\vdash \varphi \;\&\; \vdash \varphi \to \psi \Rightarrow \vdash \psi$

Figure 2: The axiomatization $\mathbf{SSL_m}$of multi-agent SSL.

R-models and their interpretation of $\mathcal{L}_m$, define a strict subset of them called R$^+$-models, and then show that R$^+$-models one-to-one correspond to multi-agent subset models. Thus, we are capable of investigating subset space logic in terms of relational structures, and by doing so, conditions lying in the "point-set" structure become frame conditions over relational structures. The correspondence result will be fundamental in the completeness proof of $\mathbf{SSL_m}$ in Section 5.

The idea of treating SSL relationally can be traced back to [Dabrowski *et al.*, 1996]. Section 2.3 of [Dabrowski *et al.*, 1996] introduces the notion of a "cross axiom model", which is based on the relational semantics. But this notion is defined only for the purpose of achieving decidability. Technically, every subset model can be truth-preservingly translated into a cross axiom model, but not necessarily vice versa, the direction which we are more interested in, for the sake of the completeness proof. With a one-to-one correspondence, we will get an alternative, relational, semantics for SSL.

Definition 7 (R-structures) An *R-frame* is a tuple $(M, \sim, \rhd)$ where

- $M = Y \times \mathcal{U}_0 \times \cdots \times \mathcal{U}_{\mathrm{N}-1}$ is an $(\mathrm{N}+1)$-dimensional non-empty domain such that Y is an arbitrary non-empty set and every $\mathcal{U}_i$ $(0 \le i \le \mathrm{N}-1)$ is a set of partitions of Y; an element of M is called a *state* and is denoted by a pair (x, O) where $x \in Y$ and $O \in \mathcal{U}_0 \times \cdots \times \mathcal{U}_{\mathrm{N}-1}$; we write O_i for the projection of O on $\mathcal{U}_i$.
- $\sim$: $\mathrm{N} \to \wp(M^2)$ assigns to every agent i an equivalence relation, $\sim_i$, on M such that $\{\{(y, O) \mid y \in [x]_{O_i}\} \mid (x, O) \in M\}$ is the set of all $\sim_i$-equivalence classes, or in other words, $(x, O) \sim_i (y, U)$ iff $O = U$ and $y \in [x]_{O_i}$;
- $\rhd$: $\mathrm{N} \to \wp(M^2)$ assigns to every agent i a reflexive transitive relation, $\rhd_i$, on M such that:
  - $\rhd_i$ links only states variant in $\mathcal{U}_i$; i.e., $(x, O) \rhd_i (y, U)$ implies $x = y$ and $\forall j.(j \neq i \Rightarrow O_j = U_j)$;
  - $\rhd_i$ is indifferent along Y; i.e., for all x, y, O, U, $(x, O) \rhd_i (x, U)$ implies $(y, O) \rhd_i (y, U)$;
  - $\rhd_i$ is indifferent along $\mathcal{U}_j$ $(j \neq i)$; i.e., for all x, y, O, U, j, if $(x, (O_j, -)) \rhd_i (y, (O_j, -'))$ then $(x, (U_j, -)) \rhd_i (y, (U_j, -'))$, where the symbols $-$ and $-'$ stand for other components of the vectors;

$\mathfrak{M} = (\mathfrak{F}, V)$ is called an *R-model (based on $\mathfrak{F}$)* if $\mathfrak{F}$ is an R-frame and $V : \mathrm{PROP} \to \wp(M)$ is an evaluation function.

By the definition of $\rhd$, it makes sense to use the notation $O \rhd_i U$ where O and U are elements of $\mathcal{U}_0 \times \cdots \times \mathcal{U}_{\mathrm{N}-1}$.

Definition 8 (R$^+$-models) An *R$^+$-model* $(M, \sim, \rhd, V)$ is an R-model that is *first-component evaluated (1CE)* and *fully designated (FD)*, as defined below:

(1CE) Evaluation is determined by the first component of a state, i.e, $(x, O) \in V(p) \Leftrightarrow (x, U) \in V(p)$;

(FD) $O \rhd_i U$ iff $O \succeq_i U$.

The set of all R$^+$-models is denoted by $\mathfrak{R}^+$.

An example of an R$^+$-model can be found in Figure 3.

Definition 9 ($\mathfrak{R}^+$-$\mathfrak{Sub}$-translation) We define a translation $\rho : \mathfrak{R}^+ \to \mathfrak{Sub}$ as follows. Let $\mathfrak{M} = (Y \times \mathcal{U}_0 \times \cdots \times \mathcal{U}_{\mathrm{N}-1}, \sim, \rhd, V)$ be an R$^+$-model. Its translation $\rho(\mathfrak{M})$ is a multi-agent subset model $(X, \mathcal{O}, \nu)$ such that $X = Y$, $\mathcal{O}_i = \mathcal{U}_i$ for all $i \in \mathrm{N}$, and $\nu(p) = \{x \mid (x, O) \in V(p) \;\&\; O \in \mathcal{U}_0 \times \cdots \times \mathcal{U}_{\mathrm{N}-1}\}$ for all $p \in \mathrm{PROP}$.

Definition 10 ($\mathfrak{Sub}$-$\mathfrak{R}^+$-translation) We define a translation $\sigma : \mathfrak{Sub} \to \mathfrak{R}^+$ as follows. Let $\mathcal{X} = (X, \mathcal{O}, V)$ be a multi-agent subset model. Its translation $\sigma(\mathcal{X})$ is an R$^+$-model $(M, \sim, \rhd, \mu)$ defined below:

- $M = X \times \mathcal{O}_0 \times \cdots \times \mathcal{O}_{\mathrm{N}-1}$;
- $\forall i \in \mathrm{N}, \sim_i = \{((x, O), (y, O)) \in M \times M \mid \{x, y\} \subseteq O_i\}$;
- $\forall i \in \mathrm{N}, \rhd_i = \{((x, O), (x, U)) \in M \times M \mid O \succeq_i U\}$, i.e., for any profiles O and U, $O \rhd_i U$ iff $O \succeq_i U$;
- $\forall p \in \mathrm{PROP}, \mu(p) = \{(x, O) \mid x \in V(p) \;\&\; O \text{ is a profile}\}$.

Definitions 9 and 10 are well defined. Namely, given $\mathfrak{M}$ an R$^+$-model, $\rho(\mathfrak{M})$ is a multi-agent subset model; and given $\mathcal{X}$ a multi-agent subset model, $\sigma(\mathcal{X})$ is an R$^+$-model. It is easy to see that a state of an R$^+$-model is itself an epistemic scenario in a corresponding multi-agent subset model.

Theorem 2 ($\mathfrak{R}^+$-$\mathfrak{Sub}$-correspondence) *For any multi-agent subset model $\mathcal{X}$ and any R$^+$-model $\mathfrak{M}$:*

$$\rho(\sigma(\mathcal{X})) = \mathcal{X}, \quad \textit{and} \quad \sigma(\rho(\mathfrak{M})) = \mathfrak{M}.$$

Thus, ρ and σ are both bijective functions, and they are the inverse functions of each other. An illustration of the $\mathfrak{R}^+$-$\mathfrak{Sub}$-correspondence is given in Figure 3.

Definition 11 (R-satisfaction) Let $\mathfrak{M} = (M, \sim, \rhd, V)$ be an R-model and a be a state of $\mathfrak{M}$.

$\mathfrak{M}, \mathsf{a} \Vdash p$	iff	$\mathsf{a} \in V(p)$
$\mathfrak{M}, \mathsf{a} \Vdash \neg\varphi$	iff	$\mathfrak{M}, \mathsf{a} \nVdash \varphi$
$\mathfrak{M}, \mathsf{a} \Vdash \varphi \land \psi$	iff	$\mathfrak{M}, \mathsf{a} \Vdash \varphi \;\&\; \mathfrak{M}, \mathsf{a} \Vdash \psi$
$\mathfrak{M}, \mathsf{a} \Vdash K_i\varphi$	iff	$\forall \mathsf{b} \in M.(\mathsf{a} \sim_i \mathsf{b} \Rightarrow \mathfrak{M}, \mathsf{b} \Vdash \varphi)$
$\mathfrak{M}, \mathsf{a} \Vdash \Box_i\varphi$	iff	$\forall \mathsf{b} \in M.(\mathsf{a} \rhd_i \mathsf{b} \Rightarrow \mathfrak{M}, \mathsf{b} \Vdash \varphi)$.

Theorem 3 (Semantic correspondence) *Let $\mathfrak{M}$ be an R$^+$-model and $\mathcal{X}$ be a multi-agent subset model. Then, for any formula φ,*

1. *For any state (x, O), $\mathfrak{M}, (x, O) \Vdash \varphi$ iff $\rho(\mathfrak{M}), x, O \models \varphi$;*
2. *For any epistemic scenario (x, O) of $\mathcal{X}$, $\mathcal{X}, x, O \models \varphi$ iff $\sigma(\mathcal{X}), (x, O) \Vdash \varphi$.*

Corollary 1 *If a set of formulas is satisfiable in an R$^+$-model, then it is also satisfiable in a multi-agent subset model.*

Let $N = 2$ and consider $\mathcal{S} = (X, \mathcal{O})$, where: $X = \{a, b, c, d\}$ and $\mathcal{O} = \{\langle 0, \{O_0^0, O_0^1, O_0^2\}\rangle, \langle 1, \{O_1^0, O_1^1\}\rangle\}$ such that $O_0^0 = \{\{a\}, \{b, c, d\}\}$, $O_0^1 = \{\{a, b, c\}, \{d\}\}$, $O_0^2 = \{\{a, b\}, \{c\}, \{d\}\}$, $O_1^0 = \{\{a, b\}, \{c, d\}\}$, and $O_1^1 = \{\{a, b\}, \{c\}, \{d\}\}$.

The following illustrates the R^+-frame $(X \times \mathcal{O}_0 \times \mathcal{O}_1, \sim, \rhd)$.

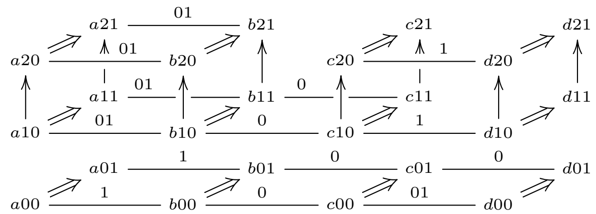

Figure 3: Illustration of a multi-agent subset space and its corresponding R^+-frame. We use labeled lines to denote the $\sim_i$-relation (where $i = 0$ or 1 is the label), and use single arrows (resp. double arrows) to denote the $\rhd_0$- (resp. $\rhd_1$-) relation. A state αxy of the R^+-frame is a shorthand for (α, O_0^x, O_1^y); e.g., $b10$ stands for (b, O_0^1, O_1^0).

5 Completeness

Completeness proofs for subset space logics turn out to be non-trivial [Dabrowski *et al.*, 1996]. In short, the canonical model method fails (already in the single-agent case it turns out to be difficult to define the set of open sets in the canonical subset model). A well known general technique for such cases is the *step-by-step method* (see, e.g., [Blackburn *et al.*, 2001, Chapter 4]). Compared with the completeness proof in [Dabrowski *et al.*, 1996] for single-agent SSL, which seems difficult to extend directly, we here instead show a completeness result with respect to $\mathfrak{R}^+$, which is equivalent to the completeness of $\mathbf{SSL_m}$ with respect to $\mathfrak{Sub}$ by the correspondence results in Section 4. The proof makes use of the standard step-by-step method for the relational semantics, which makes it clearer and more adaptable than a more ad-hoc proof. The method constructs a "canonical" model via a "network" construction, which is not a proper canonical model but shares some similarities.

We say Φ is *consistent* if $\Phi \nvdash \bot$, and *inconsistent* otherwise. Furthermore, Φ is *maximal consistent* if Φ is consistent and there is no consistent Ψ such that $\Phi \subset \Psi$. The set of all maximal consistent sets of formulas is denoted by MCS, and if $\Phi \in$ MCS, then we also say Φ *is an* MCS. For any agent i, we define $\sim_i^c$ as a binary relation on MCS such that $\Phi \sim_i^c \Psi$ iff $\forall \varphi \in \mathcal{L}_m.(K_i\varphi \in \Phi \Rightarrow \varphi \in \Psi)$. Similarly, $\rhd_i^c$ is defined to be a binary relation on MCS such that $\Phi \rhd_i^c \Psi$ iff $\forall \varphi \in \mathcal{L}_m.(\Box_i\varphi \in \Phi \Rightarrow \varphi \in \Psi)$. These two relations will serve as "canonical relations" for our "canonical" R^+-model.

Proposition 1 *Given $i \in N$ and $\Phi, \Psi \in$ MCS:*

1. *$\sim_i^c$ is an equivalence relation, and $\rhd_i^c$ is a reflexive transitive relation;*
2. *$\Phi \rhd_i^c \Psi$ implies $\forall p \in$ PROP.$(p \in \Phi \Leftrightarrow p \in \Psi)$.*

What normally follows now is to show an *existence lemma*. However, in our multi-dimensional setting with extra conditions on the frames, we need a quite sophisticated existence lemma. Intuitively, a "canonical" R^+-model is like a hypercube whose vertices are labeled by MCSs and edges by canonical relations. The question is whether we are able to label all the vertices and edges coherently. While there are general methods in the literature which are of help, we face the new problem of labeling the last vertex of the hypercube. To handle this problem, we now define the concepts of MCS-*s-cubes* and *completing vertices* before stating the existence lemma. The vertices and edges of an s-dimensional hypercube will be called an *s-cube*.

Definition 12 (MCS-*s*-cubes) For an $s \in \mathbb{N}$ with $s \geq 2$, an MCS-*s-cube* is an s-cube that forms an *R^+-frame* such that:

1. Each of its vertices is an MCS;
2. (Undirected) edges in the first dimension are labeled by a set of agents and reflect $\sim_i^c$-relations for all agents i;
3. (Directed) edges in the j-th dimension, with $2 \leq j \leq s$, are all labeled by j and reflect the $\rhd_{j-2}^c$-relation;
4. The property FD is satisfied;
5. All the $\rhd_x^c$-relations indeed link related vertices.

In the above definition we choose to say that a vertex is an MCS instead of that an MCS labels a vertex. This is not an essential difference. An MCS-s-cube is an R^+-frame in the usual sense, and it is of a special type where all edges are indeed connected by a $\sim_A^c$- or $\rhd_i^c$-relation. We mention here that the "canonical" R^+-model that we will build meets this requirement.

Definition 13 (Completing vertices) A *completing vertex* of an MCS-s-cube is a vertex such that:

1. If it is a part of a face[1] with two $\sim_i^c$-lines and two $\rhd_i^c$-arrows as edges, then it has an outgoing $\rhd_i^c$-arrow; and
2. If it is a part of a face with two $\rhd_i^c$- and two $\rhd_j^c$-arrows as edges, then it is not the *input*, i.e., the vertex where both arrows are outgoing.

Completing vertices always exist in a cube.

Lemma 1 (Existence) *Let $i, j \in N$ such that $i \neq j$, and $\Phi, \Psi, \Xi \in$ MCS, then*

1. *$\hat{K}_i\varphi \in \Phi \Rightarrow \exists\Psi.(\varphi \in \Psi \;\&\; \Phi \sim_i^c \Psi)$;*
2. *$\Diamond_i\varphi \in \Phi$ implies $\exists\Psi.(\varphi \in \Psi \;\&\; \Phi \rhd_i^c \Psi)$;*
3. *$\exists\Gamma.(\Phi \rhd_i^c \Gamma \sim_i^c \Psi)$ implies $\exists\Delta.(\Phi \sim_i^c \Delta \rhd_i^c \Psi)$;*
4. *$\exists\Gamma.(\Phi \sim_i^c \Gamma \rhd_j^c \Psi)$ iff $\exists\Delta.(\Phi \rhd_j^c \Delta \sim_i^c \Psi)$;*
5. *$\exists\Gamma.(\Phi \rhd_i^c \Gamma \rhd_j^c \Psi)$ iff $\exists\Delta.(\Phi \rhd_j^c \Delta \rhd_i^c \Psi)$;*
6. *$\Phi \sim_i^c \Psi \;\&\; \Phi \rhd_j^c \Xi$ implies $\exists\Gamma.(\Psi \rhd_j^c \Gamma \;\&\; \Xi \sim_i^c \Gamma)$;*
7. *$\Phi \rhd_i^c \Psi \;\&\; \Phi \rhd_j^c \Xi$ implies $\exists\Gamma.(\Psi \rhd_j^c \Gamma \;\&\; \Xi \rhd_i^c \Gamma)$;*
8. *If there is a possible MCS-s-cube with only one completing vertex unfinished, then there exists an MCS which can be filled in as the last vertex so as to form an MCS-s-cube.*

We now define the notion of a *perfect network* for a given consistent set Φ of formulas. It is then used to induce a satisfying "canonical" model for Φ, by the use of the existence lemma, which gives us completeness.

[1] A face (a.k.a. a 2-dimensional face) of a hypercube is a standard concept in elementary geometry. It is different from a 2-cube, in the sense that a 2-cube is restricted to a single agent whereas a face of a hypercube can have relations for multiple agents.

Definition 14 (Network) A *network* for multi-agent SSL is a tuple $(M, \sim, \rhd, l)$, where $(M, \sim, \rhd)$ is an R-frame, and $l : M \to \text{MCS}$ a *labeling function*.

A network $(M, \sim, \rhd, l)$ is called *coherent*, if satisfies:

(C1) For all $\mathsf{a}, \mathsf{b} \in M$, $\mathsf{a} \sim_i \mathsf{b} \Rightarrow l(\mathsf{a}) \sim_i^c l(\mathsf{b})$ and $\mathsf{a} \rhd_i \mathsf{b} \Rightarrow l(\mathsf{a}) \rhd_i^c l(\mathsf{b})$;

(C2) The R-frame $(M, \sim, \rhd)$ satisfies FD (see Def. 8);

(C3) For any x and i, the structure $(\{(x, (O_i, -)) \mid o_i \in M_{i+2}\}, \rhd_i)$ is a weakly-connected directed graph, i.e., it becomes a connected undirected graph when all of its $\rhd_i$-edges are replaced with undirected edges.

A network $(M, \sim, \rhd, l)$ is called *saturated*, if for all $\mathsf{a} \in M$ and all $\varphi \in \mathcal{L}_m$:

(S1) $\hat{K}_i\varphi \in l(\mathsf{a})$ implies $\exists \mathsf{b} \in M. (\mathsf{a} \sim_i \mathsf{b} \,\&\, \varphi \in l(\mathsf{b}))$, and

(S2) $\diamond_i\varphi \in l(\mathsf{a})$ implies $\exists \mathsf{b} \in M. (\mathsf{a} \rhd_i \mathsf{b} \,\&\, \varphi \in l(\mathsf{b}))$.

A network is called *perfect*, if it is coherent and saturated. $\mu = (M, \sim, \rhd, l)$ is called a *perfect network* for a set Φ of formulas, if it is perfect and there is an $\mathsf{a} \in M$ with $\Phi \subseteq l(\mathsf{a})$.

Definition 15 (Defects) Let $\mu = (M, \sim, \rhd, l)$ be a network. An *S1-defect* of μ consists of a designated point $\mathsf{a} \in M$ and a formula $\hat{K}_i\varphi \in l(\mathsf{a})$ such that $\neg\exists \mathsf{b}.(\mathsf{a} \sim_i \mathsf{b} \,\&\, \varphi \in l(\mathsf{b}))$. An *S2-defect* of μ consists of a designated point $\mathsf{a} \in M$ and a formula $\diamond_i\varphi \in l(\mathsf{a})$ such that $\neg\exists \mathsf{b}.(\mathsf{a} \rhd_i \mathsf{b} \,\&\, \varphi \in l(\mathsf{b}))$.

Let $\mu = (M, \sim, \rhd, l)$ be a network. The R-model $\mathfrak{M}_\mu = (M, \sim, \rhd, V_\mu)$ with $V_\mu(p) = \{\mathsf{a} \in M \mid p \in l(\mathsf{a})\}$ for all $p \in \text{PROP}$, is called the *μ-induced R-model*. The following lemma states the key steps of the step-by-step method. In particular, it defines our "canonical" R$^+$-model.

Lemma 2 (Step-by-step method) *Let μ be a network and $\mathfrak{M}$ the μ-induced R-model. Let Φ be a set of formulas.*

1. *If μ is coherent, then $\mathfrak{M}$ is an R$^+$-model;*
2. *If μ is perfect, then $\forall\varphi\forall\mathsf{a}.(\mathfrak{M}, \mathsf{a} \Vdash \varphi \Leftrightarrow \varphi \in l(\mathsf{a}))$;*
3. *If there is a perfect network for Φ, then Φ is satisfiable in an R$^+$-model;*
4. *(Repair lemma) If μ is finite and coherent, and let D be a defect of it, then there is a finite coherent network μ' extending μ, such that D is not a defect of μ'.*
5. *Φ is consistent iff there is a perfect network for Φ;*
6. ***SSL$_m$** is strongly complete wrt. the class of all R$^+$-models.*

We get the following from Lemma 2(6) and Corollary 1.

Theorem 4 (Completeness) ***SSL$_m$** is strongly complete with respect to the class of all multi-agent subset models.*

6 Discussion

In this paper we have proposed a multi-agent version of subset space logic (SSL), which is a natural extension of single-agent SSL. As discussed earlier, the few existing attempts at multi-agent SSL are, in our view, unsatisfactory solutions to the problem we solve. In [Başkent, 2007] any formula of the form $K_iK_j\varphi \leftrightarrow K_jK_i\varphi$ is valid, unlike in standard epistemic logic. In [Wen *et al.*, 2011], the truth value of a formula of the form $K_iK_j\varphi$ is in fact not even defined. The semantics we have proposed in this paper is well defined, it gives the standard meaning to nested knowledge, and its single-agent restriction is equivalent to the original SSL. Heinemann [2008b; 2010] obtain completeness results by extending the language with other operators (unlike in this paper), where agents do not have S5 knowledge, which are conceptually significant differences compared to the original SSL. We also introduced a relational semantics, which we used to show a correspondence between (multi-agent) SSL and a classical multi-dimensional modal logic. This brings SSL closer to classical modal logic, so that many existing techniques are now at hand. In particular, we are able to use the step-by-step method to show completeness.

There are several recent works on the relationship between public announcement logic (PAL) [Plaza, 1989] and SSL. Wáng and Ågotnes [2013] give a completeness proof for (single-agent) PAL interpreted in (single-agent) subset models. They make use of a corresponding relational semantics, but only for the single-agent case, which does not seem to be easily extended to the multi-agent case. The (single-agent) completeness result and technique in [Wáng and Ågotnes, 2013] share some similarities but is fundamentally different from the ones in the current paper. Balbiani *et al.* [2013] add public announcement operators to a (non-standard) variant of SSL, and completeness of an axiomatization is shown, but again only in the single-agent case. There is no clear direct relationship to the results or the techniques in the current paper. However, there is an interesting conceptual relationship here, as follows. The logic we have studied in this paper uses multi-agent versions of operators both for knowledge and for refinement, i.e., K_i and $\Box_i$. An agent-independent operator for refinement would also be interesting. For example, instead of talking about "agent 1 can come up with evidence so that agent 2 gets to know p" ($\diamond_1 K_2 p$), it also makes sense to talk about "there is new evidence so that agent 2 gets to know p", which we can denote by $\diamond K_2 p$. This $\diamond$ operator is more-or-less related to the $\diamond$ operator in arbitrary public announcement logic [Balbiani *et al.*, 2008] which can be read as "there is a public announcement such that ...". Balbiani *et al.* [2013] study such operators in their subset space setting. We consider using the framework developed in this paper to extend the work on public announcements in the SSL setting in [Wáng and Ågotnes, 2013; Balbiani *et al.*, 2013] to the multi-agent case a very interesting direction for future work. Also interesting for future work is incorporating group notions of knowledge into multi-agent SSL, and we believe the framework developed in this paper will help doing that. Another direction is to further extend other variants of SSL, e.g., with restrictions on subset spaces [Georgatos, 1994; Dabrowski *et al.*, 1996; Georgatos, 1997], to the multi-agent case.

Acknowledgments

Yì Wáng gratefully acknowledges funding support from the Major Project of National Social Science Foundation of China (No. 11&ZD088). Thomas Ågotnes is also affiliated with Southwest University, China, and Bergen University College, Norway.

References

[Başkent, 2007] Can Başkent. Topics in subset space logic. Master's thesis, ILLC, UvA, 2007.

[Balbiani *et al.*, 2008] Philippe Balbiani, Alexandru Baltag, Hans van Ditmarsch, Andreas Herzig, Tomohiro Hoshi, and Tiago de Lima. 'knowable' as 'known after an announcement'. *Review of Symbolic Logic*, 1(3):305–334, 2008.

[Balbiani *et al.*, 2013] Philippe Balbiani, Hans van Ditmarsch, and Andrey Kudinov. Subset space logic with arbitrary announcements. In *Proceedings of ICLA 2013*, 2013.

[Baltag and Moss, 2004] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139:165–224, 2004.

[Blackburn *et al.*, 2001] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53. Cambridge University Press, 2001.

[Dabrowski *et al.*, 1996] Andrew Dabrowski, Lawrence S. Moss, and Rohit Parikh. Topological reasoning and the logic of knowledge. *Annals of Pure and Applied Logic*, 78(1-3):73–110, 1996.

[Fagin *et al.*, 1995] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. The MIT Press: Cambridge, MA, 1995.

[Georgatos, 1994] Konstantinos Georgatos. Knowledge theoretic properties of topological spaces. In *Knowledge Representation and Reasoning Under Uncertainty*, volume 808 of *LNCS*, pages 147–159. Springer, 1994.

[Georgatos, 1997] Konstantinos Georgatos. Knowledge on treelike spaces. *Studia Logica*, 59(2):271–301, 1 1997.

[Heinemann, 2003] Bernhard Heinemann. Extended canonicity of certain topological properties of set spaces. In *Proc. of LPAR*, pages 137–151, 2003.

[Heinemann, 2004] Bernhard Heinemann. A hybrid logic of knowledge supporting topological reasoning. In Charles Rattray et al., editor, *Proc. of AMAST*, pages 181–195, 2004.

[Heinemann, 2008a] Bernhard Heinemann. A hybrid logic for reasoning about knowledge and topology. *JoLLI*, 17(1):19–41, 2008.

[Heinemann, 2008b] Bernhard Heinemann. Topology and knowledge of multiple agents. In *Proc. of IBERAMIA*, pages 1–10, 2008.

[Heinemann, 2010] Bernhard Heinemann. Logics for multisubset spaces. *JANCL*, 20(3):219–240, 2010.

[Hughes and Cresswell, 1996] George Edward Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.

[Moss and Parikh, 1992] Lawrence S. Moss and Rohit Parikh. Topological reasoning and the logic of knowledge. In Yoram Moses, editor, *Proc. of TARK*, pages 95–105, 1992.

[Parikh *et al.*, 2007] Rohit Parikh, Lawrence S. Moss, and Chris Steinsvold. Topology and epistemic logic. In *Handbook of Spatial Logics*, pages 299–341, 2007.

[Plaza, 1989] Jan A. Plaza. Logics of public communications. In *Proc. of ISMIS*, pages 201–216, 1989.

[Wáng and Ågotnes, 2013] Yì N. Wáng and Thomas Ågotnes. Subset space public announcement logic. In Kamal Lodaya, editor, *Proceedings of ICLA*, volume 7750, pages 245–257, Chennai, India, January 10-12 2013. Springer Berlin Heidelberg.

[Wang, 2009] Yi Nicholas Wang. A two-dimensional hybrid logic of subset spaces. In *Proc. of ICLA*, volume 5378 of *LNAI*, pages 198–211. Springer, 2009.

[Wen *et al.*, 2011] Xuefeng Wen, Hu Liu, and Fan Huang. An alternative logic for knowability. In *Proc. of LORI*, pages 342–355, 2011.