

Contraction Semantics for Distributed Knowledge

Floris Roelofsen and Yanjing Wang
Institute for Logic, Language, and Computation
Amsterdam, Netherlands

Abstract

The standard semantics of distributed knowledge in epistemic logic does not fully correspond to its underlying intuition, and moreover, does not fit well with bisimulation as a notion of epistemic equivalence. We propose an alternative semantics which resolves these problems.

Introduction

This note concerns the semantics of distributed knowledge in epistemic logic. In this framework, the information available to a group of agents G in a situation s is typically represented by a pointed multi-**S5** model (M, s) . This model, then, is thought of as the *collective information state* of G in s . We consider only *finite* models here.¹ The basic formal language \mathcal{L}^C that is used to talk about the agents' information represented by these models is a standard propositional language with modal operators for individual and common knowledge (K_a for an agent a and C_G for a group of agents G , respectively). To explicitly talk about the distributed knowledge among a group of agents G , i.e. the knowledge that would be obtained when all members of G were to put their individual knowledge together, the language \mathcal{L}^C is enriched with modal operators D_G . The resulting language is denoted as \mathcal{L}^{CD} . Its standard semantics contains the following clause [1, 4]:

$$(M, s) \models D_G \phi \iff (M, t) \models \phi \text{ for all } t \text{ such that } s \sim_G t \quad (1)$$

where $\sim_G = \bigcap_{a \in G} \sim_a$. That is, $D_G \phi$ holds in (M, s) iff ϕ holds in all states t in M , which are indistinguishable from s for all agents a in G .

There are two problems with this semantics. We first sketch each of these problems in some detail and then propose an alternative semantics.

¹This assumption simplifies our discussion significantly. In particular, lemma 2, which is crucial for our proof of theorem 2, does not hold for infinite models.

Problem 1: Intuition

Let (M, s) be a collective information state for a group of agents G in a situation s . Then the information state $Inf_a(M, s)$ of an agent a and the information state $Inf_G(M, s)$ of a group G are defined as follows:

$$\begin{aligned} Inf_a(M, s) &= \{\phi \in \mathcal{L}^C \mid (M, s) \models K_a\phi\} \\ Inf_G(M, s) &= \bigcup_{a \in G} Inf_a(M, s) \end{aligned}$$

Intuitively, a piece of information ϕ is distributed among a group of agents G if and only if it is obtained when all agents in G put their knowledge together, that is, if and only if ϕ is a logical consequence of $Inf_G(M, s)$. This *combined information intuition* is formalized by the following condition:

$$(M, s) \models D_G\phi \iff Inf_G(M, s) \models \phi \quad (2)$$

As witnessed by the model below, the standard semantics for distributed knowledge, given by clause (1), does not respect condition (2).




Figure 1: Model (M, s) . Links are bidirectional and the accessibility relation of M is the reflexive and transitive closure of the links actually shown.

Consider the group $G = \{a, b\}$. The only world in M that is indistinguishable from s for both agent a and agent b is s itself. So, according to (1):

$$(M, s) \models D_G p$$

while p does not follow from the information available to agent a and b :

$$Inf_G(M, s) \not\models \phi$$

So the standard semantical definition does not fully correspond to the original intuition underlying the notion of distributed knowledge. This problem has been reported before by Gerbrandy [2] and van der Hoek et.al. [6].

Problem 2: Bisimulation

A bisimilarity relation between two models M and M' is a relation \simeq between their respective domains such that, whenever $s \simeq s'$ for some s in M and some s' in M' , the following hold:

1. s and s' satisfy the same proposition letters,
2. if $s \sim_a t$ in M , then for some t' in M' we have $s' \sim_a t'$ and $t \simeq t'$,
3. if $s' \sim_a t'$ in M' , then for some t in M we have $s \sim_a t$ and $t \simeq t'$.

Two pointed models (M, s) and (M', s') are bisimilar if there is a bisimilarity relation \simeq between M and M' such that $s \simeq s'$. Bisimulation is considered to be a suitable notion of epistemic equivalence: if two finite pointed models (M, s) and (N, t) are bisimilar, then they satisfy exactly the same formulas in \mathcal{L}^C (we say that \mathcal{L}^C is *invariant under bisimulation*), and vice versa. However, things are different when the basic epistemic language is enriched with distributed knowledge operators and interpreted as in (1). To see this compare the following model with the one in figure 1:




Figure 2: Model (N, t) .

Although (M, s) and (N, t) are bisimilar, we have:

$$\begin{aligned} (M, s) &\not\models D_G p \\ (N, t) &\models D_G p \end{aligned}$$

So adopting the standard semantics for distributed knowledge undesirably implies that \mathcal{L}^{CD} is not invariant under bisimulation. This problem has been mentioned before by Gerbrandy [3] and van Benthem [5].

Solution: Bisimulation Contraction

Our solution to the problems sketched above relativizes the given semantics of distributed knowledge for a model M to the *bisimulation contraction* M' of M , which is obtained from M by identifying all bisimilar states [5].

Any model M is bisimilar to its own bisimulation contraction M' . Therefore M' can be thought of as the simplest representation of the collective information state represented by M . If $M = M'$, then we say that M is *maximally contracted under bisimulation*. An important property of maximally contracted models is that they are *distinguishing* w.r.t. \mathcal{L}^C .

Definition 1 (Distinguishing Models) *Two states s and t in a model M are distinguishable w.r.t. a language \mathcal{L} , if and only if there is a formula $\phi_{s,t} \in \mathcal{L}$ such that $(M, s) \models \phi_{s,t}$ and $(M, t) \not\models \phi_{s,t}$. M is called distinguishing w.r.t. \mathcal{L} , if every two states in M are distinguishable w.r.t. \mathcal{L} .*

Lemma 1 *If a model M is maximally contracted under bisimulation, then for every two states s and t in M , there is a formula $\beta_{s,t}$ in \mathcal{L}^C such that:*

$$\begin{aligned} (M, s) &\models \beta_{s,t} \\ (M, t) &\not\models \beta_{s,t} \end{aligned}$$

This lemma follows directly from the following result, which is attributed to Alexandru Baltag (see [5] for a proof).

Lemma 2 *For every model (M, s) there is a formula $\delta(M, s)$ in \mathcal{L}^C , such that for every model (N, t) , the following are equivalent:*

1. $(N, t) \models \delta(M, s)$
2. (M, s) and (N, t) are bisimilar.

In fact, from lemma 1 and our earlier remark that \mathcal{L}^C is invariant under bisimulation, it follows that the class of maximally contracted models is exactly the class of distinguishing models w.r.t. \mathcal{L}^C .

Lemma 3 *A model is maximally contracted under bisimulation if and only if it is distinguishing w.r.t. \mathcal{L}^C .*

Our proposal is to replace clause (1) with the following:

$$(M, s) \models D_G \phi \iff (M', t') \models \phi \text{ for all } t' \text{ such that } s' \sim_G t' \quad (3)$$

where s' is the state in M' that corresponds to s in M . That is, we suggest to impose the same condition as in (1), but applied to the bisimulation contraction (M', s') of (M, s) rather than to (M, s) itself. Doing so resolves both problems sketched above. Distributed knowledge becomes invariant under bisimulation, and compliant with the combined information intuition.

Theorem 1 *Two finite pointed models (M, s) and (N, t) are bisimilar if and only if they satisfy exactly the same formulas in \mathcal{L}^{CD} according to the contraction semantics for distributed knowledge given by clause (3).*

Proof.

(\Rightarrow) If (M, s) and (N, t) are bisimilar, then their bisimulation contractions are isomorphic. It follows that $(M, s) \models \phi$ if and only if $(N, t) \models \phi$ for any formula $\phi \in \mathcal{L}^{CD}$.

(\Leftarrow) If (M, s) and (N, t) are not bisimilar, then, by lemma 2, there is a formula $\phi \in \mathcal{L}^{CD}$ such that $(M, s) \models \phi$ and $(N, t) \not\models \phi$. \square

Theorem 2 *The contraction semantics for distributed knowledge given by clause (3) matches the combined information intuition expressed by condition (2).*

Proof. We should prove that for every (M, s) and every ϕ in \mathcal{L}^{CD} :

$$Inf_G(M, s) \models \phi \Leftrightarrow (M', t') \models \phi \text{ for all } t' \text{ such that } s' \sim_G t' \quad (4)$$

Theorem 1 implies that for every (M, s) and every ϕ :

$$Inf_G(M, s) \models \phi \Leftrightarrow Inf_G(M', s') \models \phi \quad (5)$$

Therefore it is sufficient to prove that for every formula ϕ in \mathcal{L}^{CD} and every *maximally contracted* model (M, s) :

$$Inf_G(M, s) \models \phi \Leftrightarrow (M, t) \models \phi \text{ for all } t \text{ such that } s \sim_G t \quad (6)$$

Suppose (M, s) is maximally contracted under bisimulation.

(\Rightarrow) Suppose $Inf_G(M, s) \models \phi$ and let t be any state in M such that $s \sim_G t$. For every $\varphi \in Inf_G(M, s)$ we have $(M, t) \models \varphi$. This implies $(M, t) \models \phi$.

(\Leftarrow) Lemma 1 yields that for every two states s and t in M , there is a formula $\beta_{s,t}$ in \mathcal{L}^C such that $(M, s) \models \beta_{s,t}$ and $(M, t) \not\models \beta_{s,t}$. It follows that for every state s in M the formula $\beta_s = \bigwedge_{t \in M} \beta_{s,t}$ is such that $(M, s) \models \beta_s$ if and only if $t = s$. More generally, for every set of states X in M the formula $\beta_X = \bigvee_{s \in X} \beta_s$ is such that $(M, s) \models \beta_X$ if and only if $s \in X$.

Now, for every agent $a \in G$, let $X_a = \{t \in M \mid s \sim_a t\}$ be the set of states in M that a considers indistinguishable from s . Suppose that ϕ holds in any state in M that is indistinguishable from s for all agents $a \in G$. Then, for every state w in M , we have $(M, w) \models (\bigwedge_{a \in G} \beta_{X_a}) \rightarrow \phi$. So every agent's individual information state comprises: $(\bigwedge_{a \in G} \beta_{X_a}) \rightarrow \phi$. The individual information state of agent a moreover contains β_{X_a} , so:

$$\{\beta_{X_a}\}_{a \in G} \cup \{(\bigwedge_{a \in G} \beta_{X_a}) \rightarrow \phi\} \subseteq Inf_G(M, s)$$

from which we may conclude that $Inf_G(M, s) \models \phi$. \square

Discussion

Our proposal is based on the following three observations:

1. Every collective information state can be represented by a model that is maximally contracted under bisimulation.
2. A model is maximally contracted under bisimulation if and only if it is distinguishing w.r.t. \mathcal{L}^C (see lemma 3).
3. For models that are distinguishing w.r.t. \mathcal{L}^C , condition (2) and (3) coincide.

The first observation is folklore in epistemic logic. The second is new, to the best of our knowledge, but of course not very deep by itself. The third is a somewhat generalized version of a previous result by van der Hoek et.al.². Together, these observations naturally lead to the revised semantical account we have presented in this paper, thereby resolving both problems associated with the generally accepted semantics for distributed knowledge.

References

- [1] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [2] J. Gerbrandy. Distributed knowledge. In *Formal Semantics and Pragmatics of Dialogue (TWENDIAL 98)*, pages 111–124, 1998.
- [3] J. Gerbrandy. Epistemic equivalence and bisimulation. Manuscript, submitted for publication. 2005.
- [4] J.-J.Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press, 1995.
- [5] J. van Benthem. One is a lonely number: on the logic of communication. *Technical Report PP-2002-27, ILLC Amsterdam*, 2002.
- [6] W. van der Hoek, B. van Linder, and J. Meyer. Group knowledge is not always distributed (neither is it always implicit). *Mathematical social sciences*, 38:215–240, 1999.

²Van der Hoek et.al. [6] do not consider common knowledge.