

A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases

Christoph Schwering Maurice Pagnucco

UNSW Sydney, Australia

A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases

Reasoning in multi-agent epistemic knowledge bases
reduces to
classical validity

A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases

Reasoning in multi-agent epistemic knowledge bases
reduces to
classical validity

- ▶ Logical framework: Levesque's logic of only-knowing $\mathbf{K}_A \alpha$ $\mathbf{O}_A \alpha$

A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases

Reasoning in multi-agent epistemic knowledge bases

Turing-reduces to
classical validity

- ▶ Logical framework: Levesque's logic of only-knowing $\mathbf{K}_A \alpha$ $\mathbf{O}_A \alpha$


A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases

Reasoning in multi-agent epistemic knowledge bases

Turing-reduces to

classical validity


- ▶ Logical framework: Levesque's logic of only-knowing $\mathbf{K}_A \alpha$ $\mathbf{O}_A \alpha$
- ▶ Could implement reasoning service with off-the-shelf theorem prover



■ A knows **A**


■ A knows that
but A doesn't know **B**

■ B knows **B**



KB: $\mathbf{O}_A(\mathbf{A} = 7 \wedge \forall x(\mathbf{B} = x \rightarrow \mathbf{O}_B \mathbf{B} = x))$

A only knows that $\mathbf{A} = 7$ and
that if $\mathbf{B} = x$, then B only knows that $\mathbf{B} = x$



KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (\underbrace{B = z}_{\text{de dicto}} \wedge \neg \mathbf{K}_A \underbrace{B = z}_{\text{de re}} \wedge \mathbf{K}_B \underbrace{B = z}_{\text{de re}})$

A knows that some number is equal to B , but A doesn't know what the number is, and B does know it

Reduction: Eliminate Modal Operators

KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$

Reduction: Eliminate Modal Operators

KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$

Reduction: Eliminate Modal Operators

KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$

KB': $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow P(x)))$

Reduction: Eliminate Modal Operators

$$\text{KB: } \mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$$

entails

$$\text{Query: } \mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$$

$$\text{KB': } \mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow P(x)))$$

entails

$$\text{Query': } \mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \exists x (P(x) \wedge \quad))$$

“For which x, z does $\mathbf{O}_B B = x$ entail $\mathbf{K}_B B = z$?”

Reduction: Eliminate Modal Operators

KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$

KB': $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow P(x)))$

entails

Query': $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \exists x (P(x) \wedge \quad))$

“For which x, z is $B = x \rightarrow B = z$ valid?”

Call validity oracle! [Levesque '84]

Reduction: Eliminate Modal Operators

KB: $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow \mathbf{O}_B B = x))$

entails

Query: $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \mathbf{K}_B B = z)$

KB': $\mathbf{O}_A (A = 7 \wedge \forall x (B = x \rightarrow P(x)))$

entails

Query': $\mathbf{K}_A \exists z (B = z \wedge \neg \mathbf{K}_A B = z \wedge \exists x (P(x) \wedge x = z))$

“For which x, z is $B = x \rightarrow B = z$ valid?”

Call validity oracle! [Levesque '84]

Summary

Assumption: agents always only-know something about each other.

$$\times \mathbf{O}_A(P \rightarrow \mathbf{O}_B \alpha)$$

$$\checkmark \mathbf{O}_A((P \rightarrow \mathbf{O}_B \alpha) \wedge (\neg P \rightarrow \mathbf{O}_B \beta))$$

$$\checkmark \mathbf{O}_A(\mathbf{A} = 7 \wedge \forall x(\mathbf{B} = x \rightarrow \mathbf{O}_B \mathbf{B} = x))$$

Then:

Reasoning in multi-agent epistemic knowledge bases

Turing-reduces to

classical validity

Appendix

Multi-Agent Knowledge Bases

- $O_A \alpha = A$ only-knows α [Levesque '84]
 - ▶ A considers all models of α possible

Multi-Agent Knowledge Bases

- $O_A \alpha = A$ only-knows α [Levesque '84]
 - ▶ A considers all models of α possible
- $O_A \phi$ entails $K_A \psi \iff \phi \rightarrow \psi$ is valid
 - ▶ provided that ϕ, ψ are **objective!**

Multi-Agent Knowledge Bases

- $O_A \alpha = A \text{ only-knows } \alpha$ [Levesque '84]
 - ▶ A considers all models of α possible
- $O_A \phi$ entails $K_A \psi \iff \phi \rightarrow \psi$ is valid
 - ▶ provided that ϕ, ψ are **objective!**
- $O_A \alpha$ implies $O_A \beta \iff \alpha$ and β are equivalent
 - ▶ A can only-know at most one formula

Multi-Agent Knowledge Bases

- $\mathbf{O}_A \alpha$ = A only-knows α [Levesque '84]
 - ▶ A considers all models of α possible
- $\mathbf{O}_A \phi$ entails $\mathbf{K}_A \psi \iff \phi \rightarrow \psi$ is valid
 - ▶ provided that ϕ, ψ are **objective!**
- $\mathbf{O}_A \alpha$ implies $\mathbf{O}_A \beta \iff \alpha$ and β are equivalent
 - ▶ A can only-know at most one formula
- $\mathbf{O}_A \alpha$ is a multi-agent KB \iff every model of α satisfies some $\mathbf{O}_B \beta$
 - ✗ $\mathbf{O}_A (P \rightarrow \mathbf{O}_B \alpha)$
 - ✓ $\mathbf{O}_A ((P \rightarrow \mathbf{O}_B \alpha) \wedge (\neg P \rightarrow \mathbf{O}_B \beta))$
 - ✓ $\mathbf{O}_A \forall x (f = x \rightarrow \mathbf{O}_B \alpha(x))$

Reduction

- Replace each $\mathbf{O}_A \alpha(\vec{x})$ with a fresh atom $P_\alpha(\vec{x})$

Reduction

- Replace each $\mathbf{O}_A \alpha(\vec{x})$ with a fresh atom $P_\alpha(\vec{x})$
- Replace each $\mathbf{K}_A \gamma(\vec{z})$ with a disjunction of

$$\exists \vec{x} (P_\alpha(\vec{x}) \wedge \text{“for which } \vec{x}, \vec{z} \text{ is } \alpha(\vec{x}) \rightarrow \gamma(\vec{z}) \text{ is valid?”})$$

over all $\mathbf{O}_A \alpha(\vec{x})$ at the same modal nesting level

Reduction

- Replace each $\mathbf{O}_A \alpha(\vec{x})$ with a fresh atom $P_\alpha(\vec{x})$

- Replace each $\mathbf{K}_A \gamma(\vec{z})$ with a disjunction of

$$\exists \vec{x} (P_\alpha(\vec{x}) \wedge \text{"for which } \vec{x}, \vec{z} \text{ is } \alpha(\vec{x}) \rightarrow \gamma(\vec{z}) \text{ is valid?"})$$

over all $\mathbf{O}_A \alpha(\vec{x})$ at the same modal nesting level

- Axiomatise that $P_\alpha(\vec{x}), P_\beta(\vec{y})$ introduced for $\mathbf{O}_A \alpha(\vec{x}), \mathbf{O}_A \beta(\vec{y})$

$$P_\alpha(\vec{x}) \rightarrow (P_\beta(\vec{y}) \leftrightarrow \text{"for which } \vec{x}, \vec{y} \text{ is } \alpha(\vec{x}) \rightarrow \beta(\vec{y}) \text{ is valid?"})$$

Summary

Multi-agent KB:

- Based on Levesque's **logic of only-knowing**
- Every model of a multi-agent KB must satisfy some $\mathbf{O}_B \beta$
- Allows for **incomplete knowledge** about other agent's knowledge

Summary

Multi-agent KB:

- Based on Levesque's **logic of only-knowing**
- Every model of a multi-agent KB must satisfy some $\mathbf{O}_B \beta$
- Allows for **incomplete knowledge** about other agent's knowledge

Reduction to classical reasoning:

- Oracle for **FOL validity**
- Turing reduction: calls oracle many times
- Would need FO-K45 oracle if it weren't for

Summary

Multi-agent KB:

- Based on Levesque's **logic of only-knowing**
- Every model of a multi-agent KB must satisfy some $\mathbf{O}_B \beta$
- Allows for **incomplete knowledge** about other agent's knowledge

Reduction to classical reasoning:

- Oracle for **FOL validity**
- Turing reduction: calls oracle many times
- Would need FO-K45 oracle if it weren't for

Implementation options:

- FOL theorem prover (e.g., Vampire)
- Limited belief logic