Advanced Modal Logic II

Yanjing Wang

Department of Philosophy, Peking University

Feb. 20th, 2020

Advanced Modal Logic (2020 Spring)

Kripke frames and models

From model-theoretical point of view, a logic framework is a triple: $\langle Language, Models, Semantics \rangle$ $(\mathbf{L}, \mathbb{C}, \vDash)$.

The modal language based on a set **P** of proposition letters and a *modal similarity type* $\tau = (\mathbf{O}, \rho)$ is as follows $\langle \mathbf{ML}(\mathbf{P}, \tau) \rangle$:

$$\phi \quad ::= \quad \top \mid p \mid (\phi \wedge \phi) \mid \neg\phi \mid \nabla(\underbrace{\phi \ldots \phi}_{\rho(\nabla)})$$

where $p \in \mathbf{P}$ and $\nabla \in \mathbf{O}$. (What if $\rho(\nabla)$ is 0?)
We define $\bot$ as $\neg\top$, $\phi \to \psi$ as $\neg\phi \vee \psi$, and $\triangle(\phi_1 \ldots \phi_{\rho(\nabla)})$ as $\neg\nabla(\neg\phi_1 \ldots \neg\phi_{\rho(\nabla)})$.
When **P** is clear in the context, we use **ML** to denote the basic modal language with a single unary modality $\Box$.

Kripke models and semantics: the intuition

A Kripke *model* (relational model) can be roughly pictured as a graph with labelled directed edges and nodes:

- there are some nodes (*possible worlds*, *states*, *situations* etc.), labelled by propositional letters;
- some relations among them labelled by modal operators.

A *frame* is a model without labels on nodes.

ϕ is *necessarily* true at the current world iff ϕ is true at all the possible alternatives of the current world.

$\Box\phi$ formula holds at a world *w* iff ϕ is true on all the worlds *v* that are *reachable* from *w*.

Kripke frames and models

Talking about model, we often say “a model for *some language*”

Definition (Kripke frame and Kripke model)

A *Kripke frame* for language $\mathbf{ML}(\mathbf{P}, (\mathbf{O}, \rho))$ is a pair:

$$\mathcal{F} = \langle W, \{R_\nabla \mid \nabla \in \mathbf{O}\}\rangle$$

where:

- W is a *non-empty* set;
- $R_\nabla \subseteq \underbrace{W \times \cdots \times W}_{\rho(\nabla)+1}$ is a $(\rho(\nabla) + 1)$-ary relation on W.

Since $\mathcal{F}$ is not about $\mathbf{P}$ we also call such $\mathcal{F}$ a τ-frame given a similarity type τ. A *Kripke model* $\mathcal{M}$ for $\mathbf{ML}(\mathbf{P}, \tau)$ is a pair $\langle \mathcal{F}, V\rangle$ where $\mathcal{F}$ is a τ-frame and $V : W \to 2^{\mathbf{P}}$ is a *valuation function*. A *pointed Kripke model* $(\mathcal{M}, w)$ is a Kripke model with a *designated* point w in the set of states of $\mathcal{M}$.

$2^{\mathbf{P}}$ is the power set of $\mathbf{P}$ (essentially it is a set of functions from $\mathbf{P}$ to $\{0, 1\}$).

Kripke Semantics

Recall the language $\mathbf{ML}(\mathbf{P}, (\mathbf{O}, \rho))$:

$$\phi \quad ::= \quad \top \mid p \mid (\phi \wedge \phi) \mid \neg\phi \mid \nabla(\underbrace{\phi \ldots \phi}_{\rho(\nabla)})$$

Pay attention to the quantifiers in the last clause:

Kripke Semantics ($\rho(\nabla) > 0$)

$$\begin{array}{rcl}
\mathcal{M}, w \models \top & \Leftrightarrow & \text{always} \\
\mathcal{M}, w \models p & \Leftrightarrow & p \in V(w) \\
\mathcal{M}, w \models \neg\phi & \Leftrightarrow & \mathcal{M}, w \not\models \phi \\
\mathcal{M}, w \models (\phi \wedge \psi) & \Leftrightarrow & \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \nabla(\phi_1 \ldots \phi_{\rho(\nabla)}) & \Leftrightarrow & \forall w_1 \ldots w_{\rho(\nabla)} : (\langle w, w_1, \ldots, w_{\rho(\nabla)} \rangle \in R_\nabla \\
 & & \Longrightarrow \exists i \in [1, \rho(\nabla)] : \mathcal{M}, w_i \models \phi_i)
\end{array}$$

Why do we define it like this?

Kripke models of unary similarity types

Given $\tau = (\mathbf{O}, \rho)$, if for all $\nabla \in \mathbf{O}$: $\rho(\nabla) = 1$ then we call τ a *unary similarity type*. In such cases we often use indexed boxes (e.g., $\Box_a$) to denote modalities and write $w \xrightarrow{a} v$ for $(w, v) \in R_a$ in the frames and models. We say a model is of a unary similarity type if its frame is of a unary similarity type. Given $\mathcal{M}$, we write $W_{\mathcal{M}}$, $R_a^{\mathcal{M}}$ (or $\xrightarrow{a}_{\mathcal{M}}$) and $V_{\mathcal{M}}$ for the corresponding components in $\mathcal{M}$.

$\mathcal{M}, w \vDash \Box_a \phi$ *iff* for all $v : w \xrightarrow{a}_{\mathcal{M}} v \implies \mathcal{M}, v \vDash \phi$

$\mathcal{M}, w \vDash \neg\Box_a\neg\phi$ *iff* it is not the case that $(\forall v : w \xrightarrow{a} v \implies \mathcal{M}, v \vDash \neg\phi)$

$\mathcal{M}, w \vDash \Diamond_a \phi$ *iff* there exists $v : w \xrightarrow{a}_{\mathcal{M}} v$ and $\mathcal{M}, v \vDash \phi$

Some properties of Kripke models

A Kripke frame $\mathcal{M}$ of a unary similarity type is said to be:

- *finite*, if $W_{\mathcal{M}}$ is finite. (In CS the relations are also required to be finitely representable.)
- *image-finite*, if for any world $w \in W_{\mathcal{M}}$, any $a \in \mathbf{O}$, $\{v \mid w \xrightarrow{a}_{\mathcal{M}} v\}$ is finite.
- *deterministic*, if for any world $w \in W_{\mathcal{M}}$, any $a \in \mathbf{O}$, $w \xrightarrow{a}_{\mathcal{M}} v$ and $w \xrightarrow{a}_{\mathcal{M}} v'$ implies $v = v'$.
- *well-founded*, if there is no infinite upward sequence: $\ldots \xrightarrow{a_2}_{\mathcal{M}} w_2 \xrightarrow{a_1}_{\mathcal{M}} w_1 \xrightarrow{a_0}_{\mathcal{M}} w_0$.
- *acyclic*, if there is no cycle.
- $\mathcal{M}, w$ is a *tree* if $\mathcal{M}$ is acyclic, w can reach (via some nodes) all the other nodes, and each non-w node has a unique predecessor.

An example

Consider the following Kripke model $\mathcal{M}$ of the basic similarity type ($\tau = (\{\Box\}, \rho)$ and $\rho(\Box) = 1$):

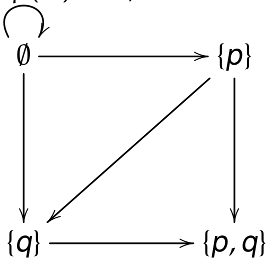

Where is $\Box\Diamond p$ true? Where is $\Box\Diamond p \wedge \Diamond\Box p$ true?

Model checking problems

Two model checking problems:

1. Local model checking: testing whether $\mathcal{M}, w \models \phi$;
2. Global model checking: compute the set $\{w \in W_{\mathcal{M}} \mid \mathcal{M}, w \models \phi\}$

Let $l_R(X) = \{w \in W_{\mathcal{M}} \mid \forall v : w \to_{\mathcal{M}} v \implies v \in X\}$, we can define the *extension* of formulas in $\mathcal{M}$:

$$
\begin{array}{rclrcl}
[\![\top]\!]^{\mathcal{M}} & = & W_{\mathcal{M}} & [\![p]\!]^{\mathcal{M}} & = & \{w \mid p \in V(w)\} \\
[\![\neg\phi]\!]^{\mathcal{M}} & = & W \backslash [\![\phi]\!]^{\mathcal{M}} & [\![(\phi \wedge \psi)]\!]^{\mathcal{M}} & = & [\![\phi]\!]^{\mathcal{M}} \cap [\![\psi]\!]^{\mathcal{M}} \\
[\![\Box\phi]\!]^{\mathcal{M}} & = & l_R[\![\phi]\!]^{\mathcal{M}} & & &
\end{array}
$$

An algorithm for global model checking ϕ on $\mathcal{M}$: labelling the states of $\mathcal{M}$ by the sub-formulas of ϕ (according to their complexity) that are true at each state. What about matrix representations?

An example

Consider the following Kripke model $\mathcal{M}$ of the basic similarity type ($\tau = \langle\{\Box\}, \rho\rangle$ and $\rho(\Box) = 1$):

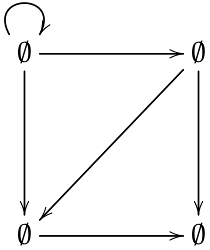

Where is $\Diamond\Box p$ true? Where is $\Box\Diamond p \wedge \Diamond\Box p$ true? Can we characterize each state by a modal formula?

Satisfiability and Validity

Modal formulas can talk about models or frames at local or global levels, this induces 4 cases:

- ϕ is *satisfiable* at $\mathcal{M}, w$ if $\mathcal{M}, w \models \phi$.
- ϕ is *valid* in $\mathcal{M}$ ($\mathcal{M} \models \phi$) if for all w in $\mathcal{M}$: $\mathcal{M}, w \models \phi$.
- ϕ is *valid* in a *pointed* frame $\mathcal{F}, w$ ($\mathcal{F}, w \models \phi$) if for any model based on $\mathcal{F}, w$: $\mathcal{M}, w \models \phi$.
- ϕ is *valid* in a frame $\mathcal{F}$ ($\mathcal{F} \models \phi$) if for any model based on $\mathcal{F}$: $\mathcal{M} \models \phi$.

	model	frame
local	$\mathcal{M}, w \models \phi$	$\mathcal{F}, w \models \phi$
global	$\mathcal{M} \models \phi$	$\mathcal{F} \models \phi$

Similarly, we can define the validity of a formula w.r.t. classes of models or frames.

When can two pointed models be seen as the same thing?

What do you think? It depends on the measurement...

- Structural equivalences ($\approx_S$), e.g., isomorphism.
- Logical equivalences ($\equiv_{\mathbf{L}}$) (two models cannot be told apart by any formula of a certain logical language **L**). More precisely, for all $\phi \in \mathbf{L}$, $\mathcal{M}, w \models \phi \iff \mathcal{N}, v \models \phi$.
- Matching the two: logical characterization of the structural equivalence: $\mathcal{M}, w \equiv_{\mathbf{L}} \mathcal{N}, v \iff \mathcal{M}, w \approx_S \mathcal{N}, v$.

If $\mathcal{M}, w \approx_S \mathcal{N}, v$ implies $\mathcal{M}, w \equiv_{\mathbf{L}} \mathcal{N}, v$ then we say **L** is *invariant* under $\approx_S$ (or $\approx_S$ is an *invariance* for **L**).

Further reading

Max Cresswell: Modal Logic before Kripke. Organon F 26(3),(2019)
B. Jack Copeland: The Genesis of Possible Worlds Semantics. J. Philosophical Logic 31(2): 99-137 (2002)
B. Jack Copeland: Meredith, Prior, and the History of Possible Worlds Semantics. Synthese 150(3): 373-397 (2006)