



Advanced Modal Logic X

Yanjing Wang

Department of Philosophy, Peking University

Mar. 26th, 2024

Advanced Modal Logic (2024 Spring)

Every modal formula can be translated into an equivalent first-order formula with one free variable in the corresponding language.*

- Standard modal language
- Over pointed models

Some first-order formula are not expressible by the modal language, e.g., Rxx .

Can quantifiers be viewed as modalities?

Theorem (Compactness Theorem)

Every finitely satisfiable set of modal formulas is satisfiable.

Theorem (Löwenheim-Skolem Theorem)

If a set of modal formulas is satisfiable in at least one infinite model, then it is satisfiable in models of every infinite cardinality (assuming the modal language is countable).

ML as a proper fragment of FOL

We can characterize it within FOL using bisimilarity.

Theorem (van Benthem Characterization Theorem)

Let $\alpha(x)$ be a first-order formula in FOL_{τ} . $\alpha(x)$ is invariant under bisimilarity iff it is equivalent to the standard translation of a modal formula.

Note that bisimilarity is the measure of **distinguishing power**.

It follows that a FOL-definable property is *not* modal definable iff there are two bisimilar models such that one has the property and the other does not.

The first proof

The first proof

A simple characterization

A modally expressible $\alpha(x)$ needs to be invariant under modal equivalence \equiv_{ML} , namely: for any $\mathcal{M}, w, \mathcal{N}, v$:

if $\mathcal{M}, w \equiv_{\text{ML}} \mathcal{N}, v$ then $\mathcal{M} \Vdash \alpha(x)[w] \iff \mathcal{N} \Vdash \alpha(x)[v]$.

Theorem (A characterization via \equiv_{ML})

Let $\alpha(x)$ be a first-order formula with one free variable in $\text{FOL}_{\tau}(\mathbf{P})$. $\alpha(x)$ is invariant under modal equivalence iff it is equivalent to (the standard translation of) a modal formula.

Corollary

Let $\alpha(x)$ be a first-order formula with one free variable in $\text{FOL}_{\tau}(\mathbf{P})$. $\alpha(x)$ is invariant under ω -bisimilarity iff it is equivalent to (the standard translation of) a modal formula.

Do we need to restrict the language in the above corollary?

The theorem is actually more general

Given L_1 and L_2 over the same class of models, if $L_1 \cup L_2$ is compact and both L_1 and L_2 are closed under classical negation and conjunction, then $L_1 \preceq_d L_2$ iff $L_1 \preceq_e L_2$.

This is the reason various classical modal logic texts do not differentiate between distinguishing power and expressive power. However, in domains where incompact logics hold significant relevance, such as temporal logic, the distinction between the two should be elucidated more clearly.

ML as a proper fragment of FOL

The previous result characterizes ML within FOL by using \equiv_{ML} (or \leftrightarrow_{ω}), which is not as elegant as the following van Benthem characterization.

Theorem (van Benthem Characterization Theorem)

Let $\alpha(x)$ be a first-order formula in FOL_{τ} . $\alpha(x)$ is invariant under bisimilarity iff it is equivalent to the standard translation of a modal formula.

To prove this theorem based on the previous result, we only need to show that:

$\alpha(x)$ is invariant under bisimilarity iff $\alpha(x)$ is invariant for \equiv_{ML} (it is not trivial since $\equiv_{\text{ML}} \neq \leftrightarrow$ in general).

A “detour” strategy

Since $\Leftrightarrow \subseteq \equiv_{\text{ML}}$, we only need to prove that if $\alpha(x)$ is invariant under bisimilarity then it is invariant under modal equivalence. Now assume $\mathcal{M}, w \equiv_{\text{ML}} \mathcal{N}, v$, $\alpha(x)$ is invariant under bisimilarity and $\mathcal{M} \Vdash \alpha(x)[w]$ we need to show $\mathcal{N} \Vdash \alpha(x)[v]$. The strategy is as follows:

$$\begin{array}{ccccc} 1. \alpha(x) & \mathcal{M}, w & \equiv_{\text{ML}} & \mathcal{N}, v & 4. \alpha(x) \\ & \downarrow \equiv_{\text{FOL}} & & \downarrow \equiv_{\text{FOL}} & \\ 2. \alpha(x) & \mathcal{M}^*, w^* & \equiv_{\text{ML}} = \Leftrightarrow & \mathcal{N}^*, v^* & 3. \alpha(x) \end{array}$$

Based on \mathcal{M}, w and \mathcal{N}, v we construct m-saturated models \mathcal{M}^*, w^* and \mathcal{N}^*, v^* such that **FOL** formulas are preserved (thus modal formulas are preserved too). Since for m-saturated models \Leftrightarrow coincides with \equiv_{ML} , $\mathcal{M}^*, w^* \Leftrightarrow \mathcal{N}^*, v^*$.

Ultrafilter extension does not do the work

However, ultrafilter extension does not preserve the truth of $\alpha(x)$. Consider the ultrafilter extension of $(\mathbb{N}, <)$. There is a “cluster” of reflexive non-principal ultrafilters at the “end” of the chain of natural numbers. Every non-principal ultrafilter is reachable from π_0 . Thus the formula $\exists y : xRy \wedge yRy$ is satisfiable at $(\mathbb{N}, <)^{ue}, \pi_0$ but not at $(\mathbb{N}, <), 0$.

We need a model construction method which can: 1. make the models m -saturated and 2. preserve truth values of *first-order* formulas.

Ultrafilters again

Original intuition behind (ultra)filters: “small” subsets are **out** only “large” ones stay (imagine a filter in the basin).

Ultrafilters were originally used to define a collection of subsets of a set W which can be regarded as “large” subsets of W in a consistent mathematical sense (see exercises). Therefore given an index set I of a family of models, if φ holds on some of \mathcal{M}_i, w_i and $\{i \mid \mathcal{M}_i, w_i \models \varphi\}$ is **in** a (non-principal) ultrafilter over I then we can say that φ holds on “almost every” \mathcal{M}_i, w_i . We use this idea to define ultraproducts of models.

A digression: many faces of ultrafilters

0 – 1 finitely-additive measure over W

A 0 – 1 function $\mu : 2^W \rightarrow \{0, 1\}$ that satisfies:

- $\mu(W) = 1$
- If Y_1, \dots, Y_n are pairwise disjoint subsets of W , then
$$\mu\left(\bigcup_{k \leq n} Y_k\right) = \sum_{i \leq n} \mu(Y_i).$$

Such a measure, viewed as $\{X \mid \mu(X) = 1\}$, is simply an ultrafilter (prove it!)

A digression: many faces of ultrafilters

Let W be a set of voters and C be a set of candidates. Let Π_C be the permutations of candidates in C . A *vote aggregation function* (VA) is $f: \Pi_C^W \rightarrow \Pi_C$. A “fair” vote aggregation is expected to satisfy:

- **Unanimity (U)**: If all individuals have the same preference, then f produces that.
- **Irrelevant alternatives (IA)**: The relative ranking of two candidates c_i, c_j in the output of f depends only on the relative rankings of c_i, c_j in each individual input.

Arrow's impossibility theorem

If W is finite and $|C| \geq 3$, the only VAs that satisfy **U** and **IA** have a dictator. (The set of f -decisive sets is an ultrafilter over W)

Take Yifeng's course for more. (Computational) social choice theory: building computational walls for undesired possibility.

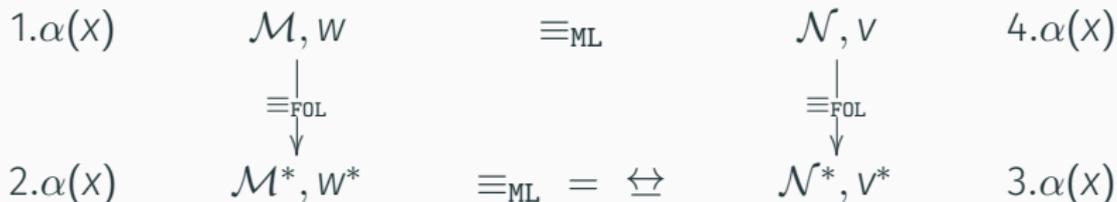
Recap: “detour” strategy

Theorem (van Benthem Characterization Theorem)

Let $\alpha(x)$ be a first-order formula in FOL_{τ} . $\alpha(x)$ is invariant under bisimilarity iff it is equivalent to the standard translation of a modal formula.

Theorem (A weaker theorem)

Let $\alpha(x)$ be a first-order formula in FOL_{τ} . $\alpha(x)$ is invariant under \equiv_{ML} iff it is equivalent to the standard translation of a modal formula.



Ultraproduct

Definition (Ultraproduct over sets)

Given a family of sets $\{W_i \mid i \in I\}$ and an ultrafilter U over the index set I , we define the equivalence relation \sim_U as

$$\sim_U = \{(f, g) \mid f, g \in \prod_{i \in I} W_i \text{ and } \{i \mid f(i) = g(i)\} \in U\}$$

The ultraproduct over these sets modulo U is the set $\{[f]_U \mid f \in \prod_{i \in I} W_i\}$. If for all i : $W_i = W$ then $\prod_U W_i = \prod_U W$ is called an ultrapower over W .

Intuition: f, g are considered the same if they coincide “almost everywhere” ($f(i) = g(i)$ for all the i belongs to some large set).

Ultraproduct

Intuition: massage many models into one satisfying what most models satisfy; (ChatGPT) merging a collection of paintings into a new piece that retains the essence shared by most of them.

Definition (Ultraproduct over models with a binary relation)

Let $\{\mathcal{M}_i \mid i \in I\}$ be a family of models. Given an ultrafilter U over I , the *ultraproduct* over $\{\mathcal{M}_i \mid i \in I\}$ is a tuple $\langle W, \rightarrow, V \rangle$ where:

- $W = \prod_U W_{\mathcal{M}_i}$
- $|f| \rightarrow |g| \iff \{i \mid f(i) \rightarrow_{\mathcal{M}_i} g(i)\} \in U$
- $p \in V(|f|) \iff \{i \mid p \in V_{\mathcal{M}_i}(f(i))\} \in U$

If for all i : $\mathcal{M}_i = \mathcal{M}$ then $\prod_U \mathcal{M}_i = \prod_U \mathcal{M}$ is called an *ultrapower* over \mathcal{M} . (What if U is principal?)

The above is well-defined: e.g., suppose $f \sim_U g$, show that $\{i \mid p \in V_{\mathcal{M}_i}(f(i))\} \in U$ iff $\{i \mid p \in V_{\mathcal{M}_i}(g(i))\} \in U$.

Łoś's Theorem

Theorem (Łoś's Theorem (one free variable case))

Fixing a U , given any first-order formula $\alpha(x)$:

$$\prod_U \mathcal{M}_i \models \alpha(x)[|f|] \iff \{i \mid \mathcal{M}_i \models \alpha(x)[f(i)]\} \in U$$

Corollary (For ultrapower)

Fixing a U , given any first-order formula $\alpha(x)$:

$$\prod_U \mathcal{M} \models \alpha(x)[|f_w|] \iff \mathcal{M} \models \alpha(x)[w]$$

Thus the mapping $d : W \rightarrow \prod_U W$ such that $d(w) = |f_w|$ is an elementary embedding.

Ultraproduct

Theorem (Łoś 's Theorem for modal logic)

Fixing a U , given any modal formula φ :

$$\prod_U \mathcal{M}_i, |f| \models \varphi \iff \{i \mid \mathcal{M}_i, f(i) \models \varphi\} \in U$$

Corollary (For ultrapower)

Given any modal formula φ :

$$\prod_U \mathcal{M}, |f_w| \models \varphi \iff \mathcal{M}, w \models \varphi$$

For any i , $f_w(i) = w$.

An application of Łoś Theorem

Theorem (Compactness for basic modal logic)

If a set of modal logic formulas is finitely satisfiable then it is satisfiable.

Proof.

Let I be an index set such that for any finite subset Γ of Σ , there is at least an $i \in I$ such that $\mathcal{M}_i, w_i \models \Gamma$. Let

$S_\varphi = \{i \mid \mathcal{M}_i, w_i \models \varphi\}$. Let $E = \{S_\varphi \mid \varphi \in \Sigma\}$. E has the finite intersection property, thus can be extended into an ultrafilter U .

Therefore by Łoś Theorem (how?), for all φ in Σ : $\prod_U \mathcal{M}_i, |f| \models \varphi$ where $f(i) = w_i$ for all $i \in I$. □

Recap again: “detour” strategy

$$\begin{array}{ccccc} 1. \alpha(x) & \mathcal{M}, w & \equiv_{\text{ML}} & \mathcal{N}, v & 4. \alpha(x) \\ & \downarrow \equiv_{\text{FOL}} & & \downarrow \equiv_{\text{FOL}} & \\ 2. \alpha(x) & \prod_U \mathcal{M}, |f_w| & \equiv_{\text{ML}} =? \Leftrightarrow & \prod_U \mathcal{N}, |f_v| & 3. \alpha(x) \end{array}$$

The ultrapower works!

We need to show that ultrapowers over certain ultrafilters are m -saturated.

Saturated models

Intuition of saturated models: they are rich enough to realize all the complete descriptions (of a point) which is consistent with the (first-order) theory of the model.

In a nutshell: saturation means big enough to include all the logically consistent potential points (应有尽有). To make our description powerful enough, we may use new constants (parameters) to denote particular points in the given model.

Saturation is usually defined by complete 1-types (potential FO-descriptions of a point; “1” means there is only one free variable in the formulas) here we have an alternative definition closer to m -saturation. The equivalence between various definitions is based on: Φ is a type w.r.t \mathcal{M} iff Φ is consistent with $Th(\mathcal{M})$ iff Φ is finitely realized in \mathcal{M} .

γ -saturated models: alternative definition

Definition (saturated models (γ is a natural number or ω))

A model \mathcal{M} is (1-type) γ -saturated iff for any finite tuple w_1, \dots, w_k in \mathcal{M} s.t. $k < \gamma \leq |\mathcal{M}|$ and any set Σ of formulas $\alpha(x, \mathbf{w}_k)$:

Σ is finitely satisfiable in \mathcal{M} then it is satisfiable in \mathcal{M} .

($\alpha(x, \mathbf{w}_k)$ is sat. in \mathcal{M} iff $\mathcal{M} \models \exists x \alpha(x, x_1, \dots, x_m)[w_1, \dots, w_k]$)

Proposition

ω -saturated models are m -saturated.

Proof.

We only show that 2-saturated models are m (odally)-saturated models with binary relations. Given a set of formulas Σ .

Suppose it is finitely satisfiable in the set of successors of a world w . Let $\Sigma' = \{Rwy\} \cup ST_y(\Sigma)$. Due to 2-saturation Σ' is satisfiable. Also work for n -types in general. □

Construct saturated models

Definition (Countably incomplete ultrafilter)

An ultrafilter over I is *countably incomplete* if it is *not* closed under countable intersections.

Countably incomplete ultrafilters must be non-principal, e.g., non-principal ultrafilter over \mathbb{N} . Existence of *non-principal* countably *complete* ultrafilter is not provable in ZFC.

Theorem

If U is a countably incomplete ultrafilter over I then the ultraproduct $\prod_U \mathcal{M}_i$ is ω -saturated, thus it is m -saturated.

Given countable $\Sigma = \{\alpha_1(x), \alpha_2(x) \dots\}$, suppose it is f -satisfiable in $\prod_U \mathcal{M}_i$. Since U is countably incomplete, find $I \supset I_1 \supset I_2 \dots$ s.t. $\bigcap I_k = \emptyset$. Define $X_0 = I$ and $X_k = I_k \cap \{i \mid \mathcal{M}_i \models \exists x \bigwedge_1^k \alpha_i(x)\}$. Thus $X_k \in U$ and $X_0 \supset X_1 \supset X_2 \dots$ s.t. $\bigcap X_k = \emptyset$. Each i “disappear” in X_k after some point $n_i \in \mathbb{N}$. Define $f(i)$ be a w in \mathcal{M}_i such that $\mathcal{M}_i \models \bigwedge_1^{n_i} \alpha_k[w]$. Since $X_n \subseteq \{i \mid \mathcal{M}_i \models \alpha_n[f(i)]\}$, $\prod_U \mathcal{M}_i[f]_U \models \Sigma$.

Recap again: “detour” strategy

$$\begin{array}{ccccc} 1.\alpha(x) & \mathcal{M}, w & \equiv_{\text{ML}} & \mathcal{N}, v & 4.\alpha(x) \\ & \downarrow \equiv_{\text{FOL}} & & \downarrow \equiv_{\text{FOL}} & \\ 2.\alpha(x) & \prod_U \mathcal{M}, |f_w| & \equiv_{\text{ML}} =? \Leftrightarrow & \prod_U \mathcal{N}, |f_v| & 3.\alpha(x) \end{array}$$

where U is a countably incomplete ultrafilter (to make sure the ultrapower is ω -saturated.)

Now we have a criterion

Given a FOL-formula with one free variable, is it decidable to test whether it is invariant under bisimulation?

The answer is: No (in general)! Consider the formula $\beta(x) : \neg\alpha(x) \wedge \exists yPy$ for some P not mentioned in $\alpha(x)$. It is easy to see that $\alpha(x)$ is valid iff $\beta(x)$ is invariant under bisimilarity. But first-order logic is undecidable in general.

Now we complete the proof

However, the previous proof is:

- highly non-trivial,
- using Axiom of Choice thus non-constructive,
- using heavy constructions w.r.t. FOL, not “modal” enough,
- using compactness of FOL.

We will present an alternative more *modal* proof, which also works for the theorem when restricted to **finite models** (thus the “**ultra tools**” are not available at all).