Advanced Modal Logic X

Yanjing Wang

Department of Philosophy, Peking University
Mar. 19th, 2020
Advanced Modal Logic (2020 Spring)

1. Van Benthem Characterization Theorem

A simple characterization: recap

Theorem (A characterization via $\equiv_{\mathtt{ML}}$)

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

A simple characterization

Proof.

Let $MOC(\alpha(x)) = \{ST_x(\phi) \mid \alpha(x) \Vdash ST_x(\phi) \text{ and } \phi \in \mathtt{ML}\}$.
Claim 1: If $MOC(\alpha(x)) \Vdash \alpha(x)$ then there is a modal formula ϕ such that $ST_x(\phi)$ is equivalent to $\alpha(x)$.
Claim 2: $MOC(\alpha(x)) \Vdash \alpha(x)$ is indeed true if $\alpha(x)$ is invariant for modal equivalence.
The first claim can be proved by an argument based on compactness.
For the second claim: suppose $\mathcal{M}, w \Vdash MOC(\alpha(x))$ then we collect all the modal formulas that are true on $\mathcal{M}, w$ and then show the set of their first-order correspondences together with $\alpha(x)$ is satisfiable on some model $\mathcal{N}, v$ (again by a compactness argument). Since $\alpha(x)$ can not distinguish modally equivalent models then $\alpha(x)$ holds on $\mathcal{M}, w$. □

ML as a proper fragment of FOL

The previous result characterizes ML within FOL by using $\equiv_{\mathtt{ML}}$ (or $\underline{\leftrightarrow}_{\omega}$), 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 $\mathtt{FOL}_{\tau}$. $\alpha(x)$ 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 $\equiv_{\mathtt{ML}}$ (it is not trivial since $\equiv_{\mathtt{ML}} \neq \underline{\leftrightarrow}$ in general).

A “detour” strategy

Since $\leftrightarrows \subseteq \equiv_{\mathtt{ML}}$, we only need to prove that if $\alpha(x)$ is invariant under bisimilarity then it is invariant under modal equivalence. Now assume $\mathcal{M}, w \equiv_{\mathtt{ML}} \mathcal{N}, v$, $\alpha(x)$ 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:

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

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

Ultrafilter extension

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 (why?). 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

Another intuition behind (ultra)filters: “small” subsets are out only “large” subsets stay (imagine a filter in the basin that we use everyday).

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 $\{i \mid \mathcal{M}_i, w_i \vDash \phi\}$ 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 $\mu : 2^W \to \{0,1\}$ that satisfies:

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

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 $f : \Pi_C^W \to \Pi_C$. 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 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 X)

See Lai Wei's slides for more. (Computational) social choice theory: setting computational walls.

Recap: “detour” strategy

Theorem (van Benthem Characterization Theorem)

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

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

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 such that if most models satisfy something then this merged one also satisfies something.

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 $\langle W, \to, V\rangle$ where:

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

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 $\{i \mid p \in V_{\mathcal{M}_i}(g(i))\} \in U$.
How to generalize this to an arbitrary similarity type?

Łós's Theorem

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

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

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

Corollary (For ultrapower)

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

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

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

Ultraproduct

Theorem (Łós's Theorem for modal logic)

Fixing a U, given any modal formula ϕ:

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

Corollary (For ultrapower)

Given any modal formula ϕ:

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

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 $S_\phi = \{i \mid \mathcal{M}_i, w_i \models \phi\}$. Let $E = \{S_\phi \mid \phi \in \Sigma\}$. E has the finite intersection property, thus can be extended into an ultrafilter U. Therefore by Łoś Theorem (how?), for all ϕ in Σ: $\prod_U \mathcal{M}_i, |f| \models \phi$ where $f(i) = w_i$ for all $i \in I$. □

Recap again: "detout" strategy

1.$\alpha(x)$	$\mathcal{M}, w$	$\equiv_{\mathtt{ML}}$	$\mathcal{N}, v$	4.$\alpha(x)$
	$\downarrow \equiv_{\mathtt{FOL}}$		$\downarrow \equiv_{\mathtt{FOL}}$	
2.$\alpha(x)$	$\prod_U \mathcal{M}, \lvert f_w \rvert$	$\equiv_{\mathtt{ML}} \quad =? \quad \underline{\leftrightarrow}$	$\prod_U \mathcal{N}, \lvert f_v \rvert$	3.$\alpha(x)$

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. To make our description powerful enough, we may use new constants (parameters) to denote particular points in the given model directly.

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 $w_1, \ldots, w_k$ in $\mathcal{M}$ s.t. $k < \gamma \leq |\mathcal{M}|$ and any set Σ of formulas $\alpha(x, \mathbf{w}_1, \ldots, \mathbf{w}_k)$:

$$\Sigma \text{ is finitely satisfiable in } \mathcal{M} \text{ then it is satisfiable in } \mathcal{M}.$$

($\alpha(x, \mathbf{w}_1, \ldots, \mathbf{w}_k)$ is sat. in $\mathcal{M}$ iff $\mathcal{M} \Vdash \exists x \alpha(x, x_1 \ldots, x_m)[w_1, \ldots, w_k]$)

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 $\Sigma' = \{R\mathbf{w}y\} \cup ST_y(\Sigma)$. Due to 2-saturation Σ' is satisfiable. □

Construct saturated models

Definition (Countably incomplete ultrafilter)

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

A principal ultrafilter is not countably incomplete. So a countably incomplete ultrafilter must be non-principal. Such ultrafilter exists (non-principal ultrafilter over $\mathbb{N}$)!
However, the existence of *non-principal* countably *complete* ultrafilter is not provable in ZFC.

Theorem

If U is a countably incomplete ultrafilter over I then the ultrapower $\prod_U \mathcal{M}$ is ω-saturated thus it is m-saturated (also work for n-types for arbitrary similarity types).

Recap again: "detour" strategy

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

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

Now we 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.