Advanced Modal Logic XI

Yanjing Wang

Department of Philosophy, Peking University

Mar. 24th, 2020

Advanced Modal Logic (2020 Spring)

1. Rosen's characterization theorem

Recap again: "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.

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

Given a $\texttt{FOL}$-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 $\beta(x) : \neg\alpha(x) \wedge \exists y P y$ for some P not mentioned in $\alpha(x)$. It is easy to see that $\alpha(x)$ is valid iff $\beta(x)$ is invariant under bisimulation. 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.

What about Modal logic on finite models?

Theorem (Rosen)

*Given a **finite** similarity type τ and a **finite** set of $\mathbf{P}$. A first-order formula $\alpha(x)$ is invariant under bisimilarity on **finite** models iff $\alpha(x)$ is equivalent to a modal formula on **finite** models.*

Van Benthem's theorem does not imply the above theorem (inv under $\underline{\leftrightarrow}$ on finite models does not imply inv. under $\underline{\leftrightarrow}$)!

Theorem (Again, a weaker one, given finite τ and $\mathbf{P}$)

A first-order formula $\alpha(x)$ is invariant under $\underline{\leftrightarrow}_k$ on finite models iff $\alpha(x)$ is equivalent to an $\mathtt{ML}_k$ formula on finite models.

Proof.

Hint: Every $\equiv_k$ equivalence class is modally definable. $\square$

Finite models

We need to show that:

Lemma

A first-order formula $\alpha(x)$ is invariant under $\underline{\leftrightarrow}$ over finite models iff $\alpha(x)$ is invariant under $\underline{\leftrightarrow}_l$ for some l over finite models.

$\Leftarrow$ is trivial. To prove the $\Rightarrow$ direction, we take the following detour strategy (we need to find the right l):

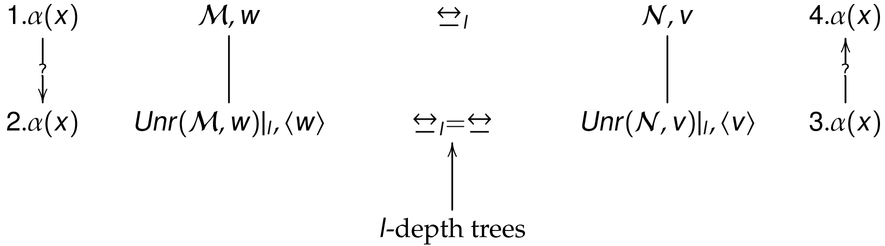

Locality

To complete the proof, we need to show:

Lemma (l-locality)

A first-order formula $\alpha(x)$ is invariant under $\underline{\leftrightarrow}$ over finite models implies that for some $l \in \mathbb{N}$, for any $\mathcal{M}, w$:
$\mathcal{M} \Vdash \alpha(x)[w] \iff Unr(\mathcal{M}, w)|_l \Vdash \alpha(x)[\langle w \rangle]$.

We already know that modal formulas are "local" in the sense that the truth value of a modal formula ϕ on a pointed model $\mathcal{M}, w$ is fully determined by some bounded "reachable" part of $\mathcal{M}, w$ (e.g., a submodel of the full unravelling of $\mathcal{M}, w$ with the first l levels only, where l is the modality depth of ϕ). We need to show the fragment of first-order logic invariant under bisimulation also has such a property of locality.

Quantifier rank

The quantifier rank $qr(\alpha)$ of a FOL-formula α is the maximum number of nested quantifiers in α:

- $qr(P(x)) = qr(x = y) = 0$
- $qr(\neg\alpha) = qr(\alpha)$
- $qr(\alpha \wedge \beta) = max(qr(\alpha), qr(\beta))$
- $qr(\forall x\alpha) = qr(\alpha) + 1$

Example

$qr(\forall y(xRy \rightarrow \exists x(yRx \wedge P(x))) = 2$

Let $mqr(\alpha(x)) = Min\{qr(\alpha'(x)) \mid \alpha'(x) \text{ is equivalent to } \alpha(x)\}$ and $mmd(\phi) = Min\{md(\phi') \mid \phi' \text{ is equivalent to } \phi\}$. Q: is $mqr(\alpha(x)) = mmd(\phi)$ if $\alpha(x)$ and ϕ are equivalent?

Ehrenfeucht-Fraïssé games (EF-games)

- The playground: two pointed models $\mathcal{M}, w$ and $\mathcal{N}, v$
- Spoiler and Defender move in turns to match the points
- Defender wins the n–round game if the play induces a partial isomorphism (isomorphism between the "played points") between the two models.

Theorem (Ehrenfeucht, 1961, simplified version)

The following two are the same:

- *Defender has a winning strategy in the m-round EF-game* $G_m(\mathcal{M}, w, \mathcal{N}, v)$
- $\mathcal{M}, w$ *and* $\mathcal{N}, v$ *satisfy the same* $\mathtt{FOL}$ *formulas (with one free variable) of quantifier rank* $\leq m$.

Ehrenfeucht-Fraïssé games (EF-games)

Example

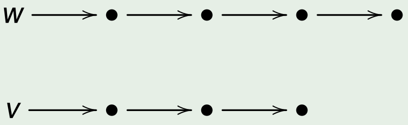

From the configuration (w, v), Spoiler needs 2 rounds to win the above game while in the corresponding bisimulation game he needs 4 rounds. (What about a 6-point one instead of the 4-point one?)
Can you express $\Diamond\Diamond\Diamond\Diamond\top$ by a $\mathtt{FOL}$ formula $\alpha(x)$ such that $qr(\alpha(x)) = 3$? In general, you can express $\Diamond^{2^n}\top$ by an equivalent $\mathtt{FOL}$ formula with quantifier rank $n+1$ (exercise).
Therefore, fixing a point, any point which is as far as 2^n away is still relevant for a $n+1$ round EF-game.

Proof strategy

We need to show that given $\alpha(x)$, there is an l for all $\mathcal{M}, w$:

$$\mathcal{M} \Vdash \alpha(x)[w] \iff Unr(\mathcal{M}, w)|_l \Vdash \alpha(x)[\langle w \rangle] \text{ for some } l$$

$$\begin{array}{ccccc}
\alpha(x) & \mathcal{M}, w & & Unr(\mathcal{M}, w)|_l, \langle w \rangle & \alpha(x) \\
\Updownarrow & \leftrightarrows & & \leftrightarrows & \Updownarrow \\
\alpha(x) & \mathcal{M}^*, w^* & ? \equiv_{\mathtt{FOL}_q} & \mathcal{N}^*, v^* & \alpha(x)
\end{array}$$

where:

$q = qr(\alpha(x))$ and $l = 2^q - 1$. Why do we pick this l?

$S = 2^{q-1} + 2^{q-2} + \cdots + 2^0, 2S = 2^q + 2^{q-1} + \cdots + 2^1 = 2^q + S - 1$

$S = ?$

Proof strategy

We need to construct *finite* $\mathcal{M}^*, w^*$ and $\mathcal{N}^*, v^*$ such that $\mathcal{M}^*, w^* \equiv_{\mathtt{FOL}_q} \mathcal{N}^*, v^*$, $\mathcal{M}^*, w^* \leftrightarroweq \mathcal{M}, w$, $\mathcal{N}^*, v^* \leftrightarroweq Unr(\mathcal{M}, w)|_l, \langle w \rangle$.
Let $PUnr_l(\mathcal{M}, w)$ be the *partial unravelling* of $\mathcal{M}, w$ w.r.t. l, i.e., $PUnr_l(\mathcal{M}, w)$ is obtained from $Unr(\mathcal{M}, w)|_l$ by replacing each leaf node $< w \ldots w' >$ with a copy of $\mathcal{M}, w'$ (unravelled one more step to avoid reflexive points). Intuitively, $PUnr(\mathcal{M}, w)|_l$ only unravels $\mathcal{M}, w$ up to level l and leave the "further" parts untouched. Since **O** is finite, we can show that if $\mathcal{M}$ is finite then there are only finitely many leaves in $Unr(\mathcal{M}, w)|_l$. Therefore if $\mathcal{M}$ is finite then $PUnr_l(\mathcal{M}, w)$ is finite too (recall that we need to construct finite models to prove Rosen's theorem, otherwise the "ultra operations" suffice).

Proof strategy

Let $\mathcal{M}^*$ be the disjoint union of $\{q$ isomorphic copies of $Unr(\mathcal{M}, w)|_l$ and another $q + 1$ isomorphic copies of $PUnr_l(\mathcal{M}, w)\}$, and w^* is the root of one of those $PUnr_l(\mathcal{M}, w)$. On the other hand, $\mathcal{N}^*$ is the disjoint union of $\{q$ isomorphic copies of $PUnr_l(\mathcal{M}, w)$ and $q + 1$ isomorphic copies of $Unr(\mathcal{M}, w)|_l\}$, and v^* is the root of one of those $Unr(\mathcal{M}, w)|_l$ Clearly, $\mathcal{M}^*, w^* \underline{\leftrightarrow} \mathcal{M}, w$, and $\mathcal{N}^*, v^* \underline{\leftrightarrow} Unr(\mathcal{M}, w)|_l, \langle w \rangle$.

Proof strategy

Intuitively, $\mathcal{M}^*, w^*$ and $\mathcal{N}^*, v^*$ are "forests" pictured as follows (q and $q+1$ denote the number of copies):

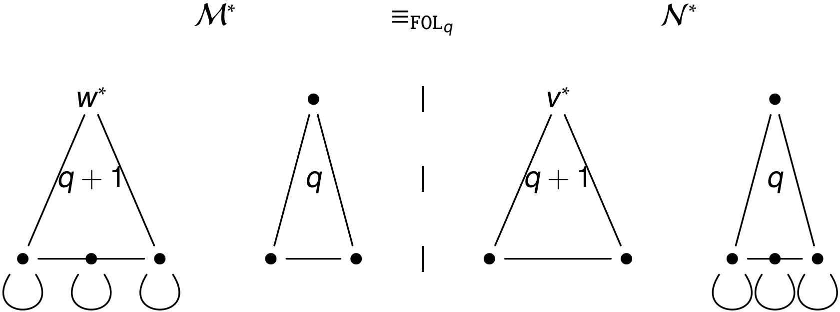

To show that $\mathcal{M}^*, w^* \equiv_{\mathtt{FOL}_q} \mathcal{N}^*, v^*$, we need a winning strategy for Defender in the q-round EF game $G_q(\mathcal{M}^*, w^*, \mathcal{N}^*, v^*)$.

Proof strategy (cf. Otto's paper)

[A winning strategy] In Round m, if Spoiler selects a world in a tree T in one model which is within the distance of 2^{q-m} to any previously selected world in T, then Defender should come up with a corresponding world in a tree of the other model with respect to the corresponding previously selected worlds in this tree. If Spoiler selects a world in a tree T in one model which is "far away" (distance $> 2^{q-m}$) from any other previously selected worlds in T, then Defender just selects a corresponding world in a "new" tree which is isomorphic to T in another model (there are enough new trees). Q: what if Spoiler selects a point in the non-tree part of a copy of $PUnr_l(\mathcal{M}, w)$?

Proof strategy

Why it works? First note that Defender can always find a matching point no matter how Spoiler plays.
We want to show that after the last play by Defender, the selection of the points form a partial isomorphism. Note that the initial selection (w^*, v^*) form a partial isomorphism. We can show that after the completion of the mth round, the selection of the points so far form a partial isomorphism **which can be extended** to the points within a bigger distance of the selected ones. Thus according to the strategy of the Defender, after the next round the selection is still a partial isomorphism.

A proof with larger l and critical distance.

We can change the distances to make it easier to prove. Let $l = 3^{q+1} - 1$ and let the critical distance in the strategy of Defender for round m be $2 \cdot 3^{q-m}$. If $q = 0$ then it is clear that Defender can win. Below we assume that $q > 0$. To prove that Defender can guarantee her win, we will show the following (claim (2) is there to make the IH strong enough for the proof):

After m rounds ($0 \leq m \leq q$): (1) the selected points form a partial isomorphism (2) if $m < q$ then this partial isomorphism can be extended to *all the* points within the distance $3^{q-m} - 1$ (not $2 \cdot 3^{q-(m+1)}$!) from the selected ones.

We show this by induction on m. When $m = 0$, (w^*, v^*) clearly forms a partial isomorphism, and the corresponding points of $\mathcal{M}^*, w^*$ and $\mathcal{N}^*, v^*$ (up to distance $3^{q-0} - 1 < 3^{q+1} - 1$ from w^* and v^*) respectively are isomorphic thus can be included in the extended partial isomorphism. Suppose (1) and (2) hold for $m = k < q$ then when $m = k + 1 \leq q$, there are two cases to be considered according to the strategy of Defender at the $k + 1$ round: (i) Spoiler selects some point outside the "$2 \cdot 3^{q-(k+1)}$-zone" and Defender responds with a corresponding point in a new isomorphic copy in the other model (there are enough new copies to choose); (ii) Spoiler selects one point in the $2 \cdot 3^{q-(k+1)}$-zone of some selected point, and Defender certainly has some nice point to respond to spoiler's challenge due to claim (2) in the **IH** for $m = k$ (and the fact that $2 \cdot 3^{q-(k+1)} \leq 3^{q-k} - 1$ when $k + 1 \leq q$).

Now we show the resulting selection is again a partial isomorphism. It is trivial in the case (ii) due to **IH**. For the case (i), since $k < q$, $2 \cdot 3^{q-(k+1)} > 1$, thus the point that Spoiler selects is not directly linked to any selected points. Therefore the point in a new copy that Defender selects is safe w.r.t. the point that Spoiler selects to extend the previous partial isomorphism. Now we look at claim (2) for which we can only consider the case when $m = k + 1 < q$. For case (i), since the point that Spoiler selects (call it s) is outside $2 \cdot 3^{q-(k+1)}$ of any selected point we can show that the points which are within $3^{q-(k+1)} - 1$-zones of s and *any* selected point are not directly linked. To see this, note that $3^{q-(k+1)} - 1 + 3^{q-(k+1)} - 1 + 1 < 2 \cdot 3^{q-(k+1)}$ which is the critical distance. Therefore, the $3^{q-(k+1)} - 1$-neighbourhood of the corresponding point of s in the new copy fits perfectly with the corresponding neighbourhood of s to extend the partial isomorphism. For case (ii), the $(3^{q-(k+1)} - 1)$-neighbourhood of the selected point by Defender is clearly at most $(2 \cdot 3^{q-(k+1)}) + 3^{q-(k+1)} - 1 = 3^{q-k} - 1$ far from some selected points after round k. We can see such points are covered by the claim (2) of **IH** for $m = k$, thus we can extend the partial isomorphism safely. Finally, let $m = q$, the claim (1) is what we want.

Rosen's theorem

Lemma

A first-order formula $\alpha(x)$ is invariant under $\underline{\leftrightarrow}$ over finite models iff $\alpha(x)$ is invariant under $\underline{\leftrightarrow}_k$ for some k over finite models.

Theorem

A first-order formula $\alpha(x)$ is invariant under $\underline{\leftrightarrow}_k$ on finite models iff $\alpha(x)$ is equivalent to a $\mathtt{ML}_k$ formula on finite models.

Theorem (Rosen, we can leave out the finite **P** and **O**)

A first-order formula $\alpha(x)$ is invariant under bisimulation (on finite models) iff $\alpha(x)$ is equivalent to a modal formula (on finite models).

If we relax the constraint on finite models the proof still works thus it can be viewed as a proof to the original vB theorem.