Modal logic and the characterization theorem

Overview of this section:

- Predicate logic has an undecidable satisfiability problem and a model checking problem of high complexity (PSPACE), hence not perfectly suited for verification of systems.
- We introduce modal logic, a fragment of predicate logic whose models are Kripke structures (essentially vertex-labeled directed graphs).
- We show: modal logic has a decidable satisfiability problem and a polynomial time decidable model checking problem.
- Temporal logics used in verification (like modal logic, LTL, CTL, μ-calculus) typically do not distinguish structures that are bisimulation equivalent.
- We show bisimulation-invariant predicate logic coincides with modal logic!

Syntax of modal logic

We fix a countable set $\mathbb{P}$ of unary relational symbols.

The set ML of formulas of modal logic is the smallest set that satisfies the following:

- $p \in$ ML for each $p \in \mathbb{P}$,
- if $\varphi \in$ ML then $\neg\varphi \in$ ML,
- if $\varphi_1, \varphi_2 \in$ ML then $(\varphi_1 \vee \varphi_2) \in$ ML,
- if $\varphi_1, \varphi_2 \in$ ML then $(\varphi_1 \wedge \varphi_2) \in$ ML,
- if $\varphi \in$ ML then $\Diamond\varphi \in$ ML, and
- if $\varphi \in$ ML then $\Box\varphi \in$ ML.

Example.

$$((p_1 \wedge p_2) \vee \neg\Box(p_3 \vee \Diamond\Diamond(p_2 \wedge \neg p_4))) \in \text{ML}$$

Semantics of modal logic

A Kripke structure is a logical structure $\mathcal{A}$ over some signature $S = \mathsf{P} \cup \{E\}$, where $\mathsf{P} \subseteq \mathbb{P}$ is finite and where E is a binary relational symbol.

For each formula $\varphi \in \mathsf{ML}$ and each suitable Kripke structure $\mathcal{A}$ and each $a \in U_{\mathcal{A}}$ we define $(\mathcal{A}, a) \models \varphi$ inductively as follows:

- $(\mathcal{A}, a) \models p$ if and only if $a \in p^{\mathcal{A}}$,
- $(\mathcal{A}, a) \models \neg\varphi$ if and only if $(\mathcal{A}, a) \not\models \varphi$,
- $(\mathcal{A}, a) \models \varphi_1 \vee \varphi_2$ if and only if $(\mathcal{A}, a) \models \varphi_1$ or $(\mathcal{A}, a) \models \varphi_2$,
- $(\mathcal{A}, a) \models \varphi_1 \wedge \varphi_2$ if and only if $(\mathcal{A}, a) \models \varphi_1$ and $(\mathcal{A}, a) \models \varphi_2$,
- $(\mathcal{A}, a) \models \Diamond\varphi$ if and only if $(\mathcal{A}, b) \models \varphi$ for some $b \in U_{\mathcal{A}}$ with $(a, b) \in E^{\mathcal{A}}$, and
- $(\mathcal{A}, a) \models \Box\varphi$ if and only if $(\mathcal{A}, b) \models \varphi$ for all $b \in U_{\mathcal{A}}$ with $(a, b) \in E^{\mathcal{A}}$.

The size and subformulas of a formula

The size $|\varphi|$ of a formula φ is defined as follows:

- $|p| = 1$ if for each $p \in \mathbb{P}$,
- $|\neg\varphi| = |\varphi| + 1$,
- $|\varphi_1 \vee \varphi_2 = |\varphi_1 \wedge \varphi_2| = |\varphi_1| + |\varphi_2| + 1$, and
- $|\Diamond\varphi| = |\Box\varphi| = |\varphi| + 1$.

The set of subformulas $\mathsf{subf}(\varphi)$ of a formula φ is defined as follows:

- $\mathsf{subf}(p) = \{p\}$ for each $p \in \mathbb{P}$,
- $\mathsf{subf}(\neg\varphi) = \{\neg\varphi\} \cup \mathsf{subf}(\varphi)$,
- $\mathsf{subf}(\varphi_1 \circ \varphi_2) = \{\varphi_1 \circ \varphi_2\} \cup \mathsf{subf}(\varphi_1) \cup \mathsf{subf}(\varphi_2)$ for each $\circ \in \{\vee, \wedge\}$, and
- $\mathsf{subf}(\circ\varphi) = \{\circ\varphi\} \cup \mathsf{subf}(\varphi)$ for each $\circ \in \{\Diamond, \Box\}$.

Note that $|\mathsf{subf}(\varphi)| = |\varphi|$.

Model checking of modal logic

Theorem. The following problem is decidable in polynomial time:

INPUT: An ML formula φ, a suitable Kripke structure $\mathcal{A}$ and some $a \in U_{\mathcal{A}}$.

QUESTION: $(\mathcal{A}, a) \models \varphi$?

Proof (Idea only, details not difficult): For each subformula ψ of φ compute the set of $b \in U_{\mathcal{A}}$ such that $(\mathcal{A}, b) \models \psi$.

Satisfiability checking of modal logic

As expected we say that an ML formula φ is satisfiable if there exists a suitable Kripke structure $\mathcal{A}$ and some $a \in U_{\mathcal{A}}$ such that $(\mathcal{A}, a) \models \varphi$.

Theorem. (Small model property of modal logic) Assume $\varphi \in$ ML is satisfiable. Then there exists a suitable Kripke structure $\mathcal{A}$ and some $a \in U_{\mathcal{A}}$ such that

- $(\mathcal{A}, a) \models \varphi$ and
- $|U_{\mathcal{A}}| \leq 2^{|\varphi|}$.

Corollary. Satisfiability of ML is decidable.

From modal logic to predicate logic

Lemma. For each formula φ there exists a formula $\overline{\varphi}(x)$ of predicate logic such that for each suitable Kripke structure $\mathcal{A}$ and each $a \in U_{\mathcal{A}}$ we have

$$(\mathcal{A}, a) \models \varphi \qquad \Leftrightarrow \quad \mathcal{A}_{[x/a]} \models \overline{\varphi}.$$

Proof.

We define the translation $\widetilde{\varphi}$ inductively as follows:

- $\widetilde{p}(x) = p(x)$ for each $p \in \mathbb{P}$,
- $\widetilde{\neg\varphi}(x) = \neg\widetilde{\varphi}(x)$,
- $\widetilde{\varphi_1 \circ \varphi_2}(x) = \widetilde{\varphi_1}(x) \circ \widetilde{\varphi_2}(x)$ for each $\circ \in \{\vee, \wedge\}$,
- $\widetilde{\Diamond\varphi}(x) = \exists y(E(x, y) \wedge \widetilde{\varphi}(y))$, and
- $\widetilde{\Box\varphi}(x) = \forall y(E(x, y) \to \widetilde{\varphi}(y))$.

Note that $\widetilde{\varphi}$ requires at most two free variables.

Predicate logic vs. modal logic

Question. Is every property expressible in predicate logic over Kripke structures expressible in modal logic?

Answer. No! Take $\varphi(x) = \exists y E(x,y) \wedge E(y,x)$ expressing that there exists a cycle of length two (proof later).

We will concern ourselves with the following questions for the rest of this section:

- How to prove that the above property is not expressible in modal logic?
- How must we restrict the properties expressible in predicate logic to obtain the properties expressible in modal logic?

Bisimulation equivalence

Let $\mathcal{A}$ and $\mathcal{B}$ be two Kripke structures suitable for some finite signature S. A bisimulation between $\mathcal{A}$ and $\mathcal{B}$ is a relation $R \subseteq U_{\mathcal{A}} \times U_{\mathcal{B}}$ such that for each $(a, b) \in R$ the following holds:

- $a \in p^{\mathcal{A}}$ if and only if $b \in p^{\mathcal{B}}$ for each $p \in \mathbb{P}$,
- for each $(a, a') \in E^{\mathcal{A}}$ there exists some $(b, b') \in E^{\mathcal{B}}$ such that $(a', b') \in R$, and
- for each $(b, b') \in E^{\mathcal{B}}$ there exists some $(a, a') \in E^{\mathcal{A}}$ such that $(a', b') \in R$.

Given $a \in U_{\mathcal{A}}$ and $b \in U_{\mathcal{B}}$ we say $(\mathcal{A}, a)$ and $(\mathcal{B}, b)$ are bisimilar (we write $(\mathcal{A}, a) \sim (\mathcal{B}, b)$ for short) if $(a, b) \in R$ for some bisimulation R between $\mathcal{A}$ and $\mathcal{B}$.

Bisimulation as a game

Consider the following bisimulation game from $\langle(\mathcal{A}_1, a_1), (\mathcal{A}_2, a_2)\rangle$ (on signature S) played between Attacker and Defender:

- Attacker chooses some $i \in \{1, 2\}$ and some $(a_i, a_i') \in E^{\mathcal{A}_i}$.
- Defender answers with some $(a_{3-i}, a_{3-i}') \in E^{\mathcal{A}_{3-i}}$.
- The game continues in $\langle(\mathcal{A}_1, a_1'), (\mathcal{A}_2, a_2')\rangle$.

Who wins a play?

- If along the play there is some pair $\langle(\mathcal{A}_1, x_1), (\mathcal{A}_2, x_2)\rangle$ and a $p \in S$ such that $x_1 \in p^{\mathcal{A}_1} \not\Leftrightarrow x_2 \in p^{\mathcal{A}_2}$, then Attacker wins!
- If the play ends such that Defender cannot answer Attacker's move (no successor), then Attacker wins.
- If the play ends $\langle(\mathcal{A}_1, x_1), (\mathcal{A}_2, x_2)\rangle$, where x_1, x_2 are both dead ends, then Defender wins.
- Defender wins each infinite play.

Finite approximants

For each $\ell \geq 0$ we define the finite approximant $\sim_\ell$ between $\mathcal{A}$ and $\mathcal{B}$ (over signature S) as follows:

$$\sim_0 = \{(a, b) \in U_\mathcal{A} \times U_\mathcal{B} \mid \forall p \in S \cap \mathbb{P} : a \in p^\mathcal{A} \Leftrightarrow b \in p^\mathcal{B}\},$$

$$\begin{aligned}\sim_{\ell+1} = \{a \sim_\ell b \mid\ & \forall (a, a') \in E^\mathcal{A} \exists (b, b') \in E^\mathcal{B} : a' \sim_\ell b' \wedge \\ & \forall (b, b') \in E^\mathcal{B} \exists (a, a') \in E^\mathcal{A} : a' \sim_\ell b'\}\end{aligned}$$

One easily sees that $\sim_\ell$ is an equivalence relation for each $\ell \in \mathbb{N}$.

Moreover $\sim \subseteq \sim_\ell$ for each $\ell \in \mathbb{N}$.

Bisimulation as a game

Theorem. Defender has a winning strategy from $\langle(\mathcal{A}, a), (\mathcal{B}, b)\rangle$ if and only if $(\mathcal{A}, a) \sim (\mathcal{B}, b)$.

Theorem. Defender has a winning strategy from $\langle(\mathcal{A}, a), (\mathcal{B}, b)\rangle$ in the ℓ round game if and only if $(\mathcal{A}, a) \sim_\ell (\mathcal{B}, b)$.

Fact. Bisimulation is insensitive to disjoint sums: We have $(\mathcal{A}, a) \sim (\mathcal{B}, b)$ if and only if $(\mathcal{A} + \mathcal{C}, a) \sim (\mathcal{B}, b)$, where $\mathcal{A} + \mathcal{C}$ denotes the disjoint sum of $\mathcal{A}$ and $\mathcal{C}$.

$\sim_\ell$ and ML$_\ell$

For each $\varphi \in$ ML, let us define the modal depth md(φ) as follows:

- md$(p) = 0$ for each $p \in \mathbb{P}$,
- md$(\neg\varphi) =$ md(φ),
- md$(\varphi_1 \vee \varphi_2) =$ md$(\varphi_1 \wedge \varphi_2) = \max\{md(\varphi_1),$ md$(\varphi_2)\}$, and
- md$(\Diamond\varphi) =$ md$(\Box\varphi) =$ md$(\varphi) + 1$.

For each $\ell \geq 0$ define ML$_\ell$ $=\{\varphi \in$ ML $\mid$ md$(\varphi) = k\}$.

Lemma. Let $\ell \in \mathbb{N}$. Let $\mathcal{A}$ and $\mathcal{B}$ be Kripke structures over a finite signature S and let $a \in U_\mathcal{A}$ and $b \in U_\mathcal{B}$. Then we have:

(1) $\sim_\ell$ has finitely many equivalence classes.

(2) $(\mathcal{A}, a) \sim_\ell (\mathcal{B}, b)$ iff $(\mathcal{A}, a) \models \varphi \Leftrightarrow (\mathcal{B}, b) \models \varphi$ for all $\varphi \in$ ML$_\ell$.

(3) Each equivalence class of $\sim_\ell$ is definable by some ML$_\ell$ formula.

Trees

A Kripke structure $\mathcal{A}$ is a tree (structure) if $(U_\mathcal{A}, E^\mathcal{A})$ is a directed tree, i.e. $\mathcal{A}$ is acyclic, the symmetric closure of $E^\mathcal{A}$ is connected and each node has at most one incoming edge.

A tree $\mathcal{A}$ has depth ℓ if each path in $\mathcal{A}$ has length at most ℓ.

For $\ell \geq 0$ we say $(\mathcal{A}, a)$ is ℓ-locally a tree structure if $\mathcal{A} \upharpoonright N_\ell(a)$ is a tree structure.

Lemma.

(1) $(\mathcal{A}, a) \sim_\ell (\mathcal{B}, b)$ iff $(\mathcal{A} \upharpoonright N_\ell(a), a) \sim_\ell (\mathcal{B} \upharpoonright N_\ell(b), b)$.

(2) If $\mathcal{A}$ and $\mathcal{B}$ are trees of depth ℓ, then

$$(\mathcal{A}, a) \sim_\ell (\mathcal{B}, b) \text{ iff } (\mathcal{A}, a) \sim (\mathcal{B}, b).$$

Unravellings

The unravelling of $\mathcal{A}$ at some $a \in U_{\mathcal{A}}$ is the tree $\mathcal{A}_a^*$, where

- $U_{\mathcal{A}_a^*} = \{\pi \mid \pi$ is a finite path in $\mathcal{A}$ starting at $a\}$.
- $E^{\mathcal{A}_a^*} = \{(\pi, \pi') \in (U_{\mathcal{A}_a^*})^2 \mid \exists (u, v) \in E^{\mathcal{A}} : \pi' = \pi(x, y)\}$.

Lemma. Let $\mathcal{A}$ be a Kripke structure and let $a \in U_{\mathcal{A}}$. Then we have

- $(\mathcal{A}_a^*, a) \sim (\mathcal{A}, a)$.
- $(\mathcal{A}_a^* \upharpoonright N_\ell(a), a) \sim_\ell (\mathcal{A}, a)$.

Bisimulation invariance and locality

A predicate logic formula $F(x)$ over a Kripke signature is bisimulation invariant if the following holds for all suitable $(\mathcal{A}, a), (\mathcal{B}, b)$:

$$(\mathcal{A}, a) \sim (\mathcal{B}, b) \Longrightarrow \big(\mathcal{A}_{[x/a]} \models F \Leftrightarrow (\mathcal{B}_{[x/b]} \models F\big)$$

A predicate logic formula $F(x)$ over a Kripke signature is ℓ-local if for all suitable $(\mathcal{A}, a)$ we have

$$\mathcal{A}_{[x/a]} \models F \iff \mathcal{A} \restriction N_\ell(a)_{[x/a]} \models F$$

The Characterization Theorem

Theorem (van Benthem/Rosen, proof by Otto). The following are equivalent for any predicate logic formula $F(x)$ over a Kripke signature with $\mathrm{qr}(F) = q$:

- $F(x)$ is bisimulation-invariant.
- $F(x)$ is logically equivalent to some ML_ℓ formula, where $\ell = 2^q - 1$.

The same holds when restricted to the class of finite Kripke structures.

Proof Outline of Characterization Theorem

We prove the Characterization Theorem in three steps:

(1) Any bisimulation invariant $F(x)$ of predicate logic is ℓ-local for $\ell = 2^q - 1$, where $q = \mathsf{qr}(F)$.

(2) Any bisimulation invariant $F(x)$ that is ℓ-local is even invariant under ℓ-bisimulation equivalence $\sim_\ell$.

(3) Any property invariant under ℓ-bisimulation equivalence is definable in ML_ℓ.