Craig Interpolation and Uniform Interpolants in Modal Logic

Frank Wolter, University of Liverpool

Telavi
September, 2023

Plan

- Bisimulation based criterion for interpolant existence for many modal logics;
- Bisimulation based proof of CIP for many modal logics;
- Computing (uniform) interpolants in exponential time for K;
- Exponential lower bound for uniform interpolants for K;
- Note on uniform interpolants for global consequence for K.

Modal Logic

The language ML of modal logic:

$$\varphi, \psi := p_i \mid \top \mid \neg\varphi \mid \varphi \wedge \psi \mid \Diamond\varphi$$

and $\Box\varphi = \neg\Diamond\neg\varphi$ and $\bot = \neg\top$.

Modal Logic

The language ML of modal logic:

$$\varphi, \psi := p_i \mid \top \mid \neg\varphi \mid \varphi \wedge \psi \mid \Diamond\varphi$$

and $\Box\varphi = \neg\Diamond\neg\varphi$ and $\bot = \neg\top$.

ML is interpreted in models $M = (W, R, V)$, where (W, R) is a Kripke frame with worlds W and an accessibility relation $R \subseteq W \times W$ and V is a valuation with $V(p_i) \subseteq W$. Then

- $M, w \models p_i$ iff $w \in V(p_i)$;
- standard for Booleans, for instance $M, w \models \varphi \wedge \psi$ if $M, w \models \varphi$ and $M, w \models \psi$;
- $M, w \models \Diamond\varphi$ if there is $v \in W$ with wRv and $M, v \models \varphi$.

Modal Logic

We write $\varphi \models \psi$ (sometimes $\varphi \models_K \psi$) if for all pointed models M, w:

$$M, w \models \varphi \quad \text{implies} \quad M, w \models \psi$$

If we restrict the class of Kripke frames to some class corresponding to a modal logic L, then we write $\varphi \models_L \psi$ if for all models $M = (W, R, V)$ with $(W, R) \models L$ and worlds w:

$$M, w \models \varphi \quad \text{implies} \quad M, w \models \psi$$

For instance,

- $L = S4$ is the logic of all transitive and reflexive frames;
- $L = K4.3$ is the logic of all linear frames.

Interpolants in Modal Logic

A formula χ is called a Craig interpolant of φ, ψ in L if $\mathrm{sig}(\chi) \subseteq \mathrm{sig}(\varphi) \cap \mathrm{sig}(\psi)$ and

$$\varphi \models_L \chi \models_L \psi$$

L has the Craig interpolation property (CIP) if a Craig interpolant of φ, ψ exists whenever $\varphi \models_L \psi$.

A Criterion for Interpolant Existence for Compact Modal logics

Let Σ be a finite signature.

Pointed models M_1, w_1 and M_2, w_2 are Σ-indistinguishable,

$$M_1, w_1 \equiv_\Sigma M_2, w_2,$$

if $M_1, w_1 \models \varphi$ iff $M_2, w_2 \models \varphi$, for all formulas φ with $\mathrm{sig}(\varphi) \subseteq \Sigma$.

A Criterion for Interpolant Existence for Compact Modal logics

Let Σ be a finite signature.

Pointed models M_1, w_1 and M_2, w_2 are Σ-indistinguishable,

$$M_1, w_1 \equiv_\Sigma M_2, w_2,$$

if $M_1, w_1 \models \varphi$ iff $M_2, w_2 \models \varphi$, for all formulas φ with $\text{sig}(\varphi) \subseteq \Sigma$.

Theorem. The following conditions are equivalent for compact $\models_L$, any formulas φ, ψ and $\Sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$:

- there does not exist an interpolant of φ, ψ in L;
- φ and $\neg\psi$ are satisfiable in Σ-indistinguishable models.

Proof

"$\Leftarrow$" If φ and $\neg\psi$ are satisfiable in Σ-indistinguishable models, then we have

- $M_1, w_1 \models \varphi$;
- $M_2, w_2 \models \neg\psi$;
- $M_1, w_1 \equiv_\Sigma M_2, w_2$.

Assume χ is an interpolant in L of φ, ψ. Then from $\varphi \models_L \chi$, $M_1, w_1 \models \chi$. By $\mathrm{sig}(\chi) \subseteq \Sigma$, $M_2, w_2 \models \chi$. This contradicts $M_2, w_2 \models \neg\psi$.

Proof

"$\Rightarrow$" Assume no interpolant exists. Let

$$\varphi^{\Sigma} = \{\chi \mid \varphi \models_L \chi, \mathsf{sig}(\chi) \subseteq \Sigma\}$$

By compactness $\varphi^{\Sigma} \not\models_L \psi$. Take a model M_2, w_2 of $\varphi^{\Sigma} \cup \{\neg\psi\}$. Let

$$t^{\Sigma}_{M_2} = \{\chi \mid \mathsf{sig}(\chi) \subseteq \Sigma, M_2, w_2 \models \chi\}$$

By compactness we find a model M_1, w_1 of $t^{\Sigma}_{M_2} \cup \{\varphi\}$. By definition

$$M_1, w_1 \equiv_{\Sigma} M_2, w_2.$$

Characterise $\equiv_\Sigma$: Bisimulations

Let Σ be a finite set of propositional atoms. Let $M_1 = (W_1, R_1, V_1)$ and $M_2 = (W_2, R_2, V_2)$ be models.

Relation $\beta \subseteq W_1 \times W_2$ is a Σ-bisimulation between M_1 and M_2 if:

- $(w_1, w_2) \in \beta$ implies $w_1 \in V_1(p)$ iff $w_2 \in V_2(p)$ for all $p \in \Sigma$;
- If $(w_1, w_2) \in \beta$ and $(w_1, w_1') \in R_1$, then there exists w_2' with $(w_2, w_2') \in R_2$ and $(w_1', w_2') \in \beta$; and vice versa.

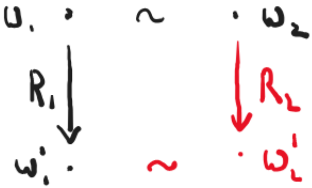

Bisimulations

M_1, x_1 and M_2, x_2 are **Σ-bisimilar**, in symbols,

$$M_1, x_1 \sim_\Sigma M_2, x_2,$$

if there exists a Σ-bisimulation β between M_1 and M_2 with $(x_1, x_2) \in \beta$.

Example

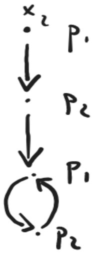

Bisimulation Characterisation

Theorem. For all finite outdegree/ω-saturated models M_1, w_1 and M_2, w_2 of the following are equivalent:

$$M_1, w_1 \sim_\Sigma M_2, w_2 \quad \text{iff} \quad M_1, w_1 \equiv_\Sigma M_2, w_2$$

The direction '$\Rightarrow$' always holds.

Criterion for Craig interpolant existence

We say that φ and ψ are satisfiable in Σ-bisimilar models if there are pointed models

- $M_1, w_1 \models \varphi$;
- $M_2, w_2 \models \psi$;

such that $M_1, w_1 \sim_\Sigma M_2, w_2$.

Criterion for Craig interpolant existence

We say that φ and ψ are satisfiable in Σ-bisimilar models if there are pointed models

- $M_1, w_1 \models \varphi$;
- $M_2, w_2 \models \psi$;

such that $M_1, w_1 \sim_\Sigma M_2, w_2$.

Theorem. The following conditions are equivalent for any L determined by an FO-definable class of frames and formulas φ, ψ and $\Sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$:

- there does not exist an interpolant of φ, ψ in L
- φ and $\neg\psi$ are satisfiable in Σ-bisimilar models.

Criterion for CIP

Theorem. Let L be determined by an FO-definable class of frames. Then L has CIP if for $\Sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$ the following are equivalent

- $\varphi \wedge \neg\psi$ is satisfiable
- φ and $\neg\psi$ are satisfiable in Σ-bisimilar models.

Criterion for CIP

Theorem. Let L be determined by an FO-definable class of frames. Then L has CIP if for $\Sigma = \mathrm{sig}(\varphi) \cap \mathrm{sig}(\psi)$ the following are equivalent

- $\varphi \wedge \neg\psi$ is satisfiable
- φ and $\neg\psi$ are satisfiable in Σ-bisimilar models.

Task. Construct from any Σ-bisimilar $M_1, w_1 \models \varphi$ and $M_2, w_2 \models \neg\psi$ a single $M, z \models \varphi \wedge \neg\psi$.

Criterion for CIP

Theorem. Let L be determined by an FO-definable class of frames. Then L has CIP if for $\Sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$ the following are equivalent

- $\varphi \wedge \neg\psi$ is satisfiable
- φ and $\neg\psi$ are satisfiable in Σ-bisimilar models.

Task. Construct from any Σ-bisimilar $M_1, w_1 \models \varphi$ and $M_2, w_2 \models \neg\psi$ a single $M, z \models \varphi \wedge \neg\psi$.

Lots of research on algebraic reformulation (amalgamation of algebras). We here discuss the 'bisimulation product' approach introduced by Marx.

Bisimulation products

Assume $M_1 = (F_1, V_1)$ and $M_2 = (F_2, V_2)$ and β is a Σ-bisimulation between M_1 and M_2 with $(x_1, x_2) \in \beta$.

The bisimulation product $M_\beta = (F_\beta, V_\beta)$ is defined by setting

$$F_\beta = (F_1 \times F_2)_{|\beta}$$

and by setting for the projections $\pi_i : F_\beta \to F_i$:

- $V_\beta(p) = \pi_1^{-1}(V_1(p))$, for $p \in \text{var}(\varphi)$;
- $V_\beta(p) = \pi_2^{-1}(V_2(p))$, for $p \in \text{var}(\psi)$

This is well defined for $p \in \text{var}(\varphi) \cap \text{var}(\psi)$.

Bisimulation Products

The projections $\pi_i : M_\beta \to M_i$ are then actually bisimulations and so

- $M_\beta, (x_1, x_2) \models \varphi$ since $M_1, x_1 \models \varphi$;
- $M_\beta, (x_1, x_2) \models \neg\psi$ since $M_2, x_2 \models \neg\psi$.

Bisimulation Products

The projections $\pi_i : M_\beta \to M_i$ are then actually bisimulations and so

- $M_\beta, (x_1, x_2) \models \varphi$ since $M_1, x_1 \models \varphi$;
- $M_\beta, (x_1, x_2) \models \neg\psi$ since $M_2, x_2 \models \neg\psi$.

Theorem. If L is determined by an FO-definable class of frames closed under cartesian products and subframes, then L has CIP.

This is the case for all L with frames defined by by universal Horn sentences

$$\forall\vec{x}(R(\vec{x}) \wedge \cdots \wedge R(\vec{x}) \to R(\vec{x}))$$

Examples. K4, S4, S5, T.

Counterexamples for closure under bisim products

Linear frames, transitive frames satisfying

$$\forall x, y(x = y \vee R(x, y) \vee R(y, x)),$$

are not preserved under bsimiluation products:

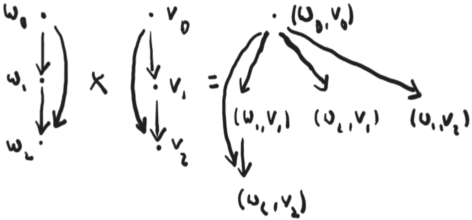

Uniform Interpolants

A formula χ is called a uniform interpolant for φ and $\Sigma \subseteq \text{sig}(\varphi)$ if it is an interpolant for φ, ψ whenever

- $\varphi \models \psi$;
- $\text{sig}(\varphi) \cap \text{sig}(\psi) \subseteq \Sigma$;
- in particular, $\text{sig}(\chi) \subseteq \Sigma$.

Uniform Interpolants

A formula χ is called a uniform interpolant for φ and $\Sigma \subseteq \text{sig}(\varphi)$ if it is an interpolant for φ, ψ whenever

- $\varphi \models \psi$;
- $\text{sig}(\varphi) \cap \text{sig}(\psi) \subseteq \Sigma$;
- in particular, $\text{sig}(\chi) \subseteq \Sigma$.

In contrast to Craig interpolants, uniform interpolants are unique up to logical equivalence as they are the logically strongest Craig interpolant.

Uniform Interpolants

A formula χ is called a uniform interpolant for φ and $\Sigma \subseteq \text{sig}(\varphi)$ if it is an interpolant for φ, ψ whenever

- $\varphi \models \psi$;
- $\text{sig}(\varphi) \cap \text{sig}(\psi) \subseteq \Sigma$;
- in particular, $\text{sig}(\chi) \subseteq \Sigma$.

In contrast to Craig interpolants, uniform interpolants are unique up to logical equivalence as they are the logically strongest Craig interpolant.

$\exists \mathbf{x}.\varphi$, $\mathbf{x} = \text{sig}(\varphi) \setminus \Sigma$, is a uniform interpolant in second-order modal logic, but we cannot express it in modal logic.

Uniform Interpolants and Bisimulation Quantifiers

For $\mathbf{x} \subseteq \text{sig}(\varphi)$, let $\exists^{\sim}\mathbf{x}.\varphi$ be a formula with the truth condition

- $M, w \models \exists^{\sim}\mathbf{x}.\varphi$ if exists M', w' with $M, w \sim_{\text{sig}(\varphi)\setminus\mathbf{x}} M', w'$ and $M', w' \models \varphi$.

Uniform Interpolants and Bisimulation Quantifiers

For $\mathbf{x} \subseteq \text{sig}(\varphi)$, let $\exists^{\sim}\mathbf{x}.\varphi$ be a formula with the truth condition

- $M, w \models \exists^{\sim}\mathbf{x}.\varphi$ if exists M', w' with $M, w \sim_{\text{sig}(\varphi)\setminus\mathbf{x}} M', w'$ and $M', w' \models \varphi$.

It is called bisimulation quantifier and weakens second-order quantification to quantification modulo a bisimulation. For $\mathbf{x} = \text{sig}(\varphi) \setminus \text{sig}(\psi)$:

$\neg\psi \land \exists^{\sim}\mathbf{x}.\varphi$ is sat iff there is no interpolant of φ, ψ

Uniform Interpolants and Bisimulation Quantifiers

For $\mathbf{x} \subseteq \text{sig}(\varphi)$, let $\exists^{\sim}\mathbf{x}.\varphi$ be a formula with the truth condition

- $M, w \models \exists^{\sim}\mathbf{x}.\varphi$ if exists M', w' with $M, w \sim_{\text{sig}(\varphi)\setminus\mathbf{x}} M', w'$ and $M', w' \models \varphi$.

It is called bisimulation quantifier and weakens second-order quantification to quantification modulo a bisimulation. For $\mathbf{x} = \text{sig}(\varphi) \setminus \text{sig}(\psi)$:

$$\neg\psi \land \exists^{\sim}\mathbf{x}.\varphi \text{ is sat} \quad \text{iff} \quad \text{there is no interpolant of } \varphi, \psi$$

Equivalently, $\exists^{\sim}\mathbf{x}.\varphi$ is a uniform interpolant (if expressible):

- $\exists^{\sim}\mathbf{x}.\varphi \models \psi$ iff
- there is an interpolant of φ, ψ iff
- $\varphi \models \psi$ (by CIP).

Example for bisimulation quantifiers

Let

$$\varphi = \Diamond(p \wedge x) \wedge \Diamond(p \wedge \neg x)$$

Then $M, w \models \exists x.\varphi$ if w has at least two successors satisfying p. This cannot be expressed in ML.

$M, w \models \exists^{\sim} x.\varphi$ if w has a successor satisfying p. This is expressed by $\Diamond p$.

Uniform Interpolants

Theorem $\models_K$ has uniform interpolation. Uniform interpolants can be constructed in exponential time.

The uniform interpolant for φ and Σ is equivalent to $\exists^{\sim}\mathbf{x}.\varphi$, for $\mathbf{x} = \mathrm{sig}(\varphi) \setminus \Sigma$.

Uniform Interpolants

Theorem $\models_K$ has uniform interpolation. Uniform interpolants can be constructed in exponential time.

The uniform interpolant for φ and Σ is equivalent to $\exists^{\sim}\mathbf{x}.\varphi$, for $\mathbf{x} = \text{sig}(\varphi) \setminus \Sigma$.

Example. $\Diamond p$ is the uniform interpolant for $\Diamond(p \wedge x) \wedge \Diamond(p \wedge \neg x)$ and $\Sigma = \{p\}$

Motivation for proof

For every propositional formula there exists an equivalent formula in DNF. We can assume it takes the form

$$\varphi = \bigvee_{i \in I} at_i$$

with each at_i a satisfiable conjunction of literals.

Motivation for proof

For every propositional formula there exists an equivalent formula in DNF. We can assume it takes the form

$$\varphi = \bigvee_{i \in I} at_i$$

with each at_i a satisfiable conjunction of literals.

Then $\exists \mathbf{x}.\varphi \equiv \bigvee_{i \in I} at_i^{-\mathbf{x}}$, where $at_i^{-\mathbf{x}}$ is obtained from at_i by dropping $\mathbf{x}$.

Motivation for proof

For every propositional formula there exists an equivalent formula in DNF. We can assume it takes the form

$$\varphi = \bigvee_{i \in I} at_i$$

with each at_i a satisfiable conjunction of literals.

Then $\exists \mathbf{x}.\varphi \equiv \bigvee_{i \in I} at_i^{-\mathbf{x}}$, where $at_i^{-\mathbf{x}}$ is obtained from at_i by dropping $\mathbf{x}$.

Proof. Clearly $\exists \mathbf{x}.\varphi \models \bigvee_{i \in I} at_i^{-\mathbf{x}}$.
Conversely, assume $v \models \bigvee_{i \in I} at_i^{-\mathbf{x}}$. Take $i \in I$ with $v \models at_i^{-\mathbf{x}}$.
As at_i is sat, we can expand v to v' so that $v' \models at_i$. Hence $v \models \exists \mathbf{x}.\varphi$.

Generalisation to ML

Let Φ be a finite set of formulas. Set

$$\nabla\Phi = \bigwedge_{\chi\in\Phi} \Diamond\chi \wedge \Box \bigvee_{\chi\in\Phi} \chi$$

Formulas in disjunctive form are defined recursively by

$$\varphi, \psi := \top \mid \bot \mid \text{at} \wedge \nabla\Phi \mid \varphi \vee \psi$$

with at a satisfiable conjunction of literals and Φ formulas in disjunctive form.

Disjunctive Form

Theorem. [ten Cate et al. 2006] For every ML-formula one can construct an equivalent ML-formula in disjunctive form in exponential time.

Starting with negation normal form the crucial step is dealing with conjunctions. Here use distributive law and for

$$\Diamond\chi_1 \wedge \cdots \wedge \Diamond\chi_n \wedge \Box\chi'_1 \wedge \cdots \wedge \Box\chi'_m \quad \Rightarrow \quad \nabla\{\chi_i \wedge \bigwedge_{j \leq m} \chi'_j \mid i \leq n\}$$

Disjunctive Form

Theorem. [ten Cate et al. 2006] For every ML-formula one can construct an equivalent ML-formula in disjunctive form in exponential time.

Starting with negation normal form the crucial step is dealing with conjunctions. Here use distributive law and for

$$\Diamond\chi_1 \wedge \cdots \wedge \Diamond\chi_n \wedge \Box\chi'_1 \wedge \cdots \wedge \Box\chi'_m \quad \Rightarrow \quad \nabla\{\chi_i \wedge \bigwedge_{j \leq m} \chi'_j \mid i \leq n\}$$

Now for φ in disjunctive form $\exists^{\sim}\mathbf{x}.\varphi \equiv \varphi^{-\mathbf{x}}$ with $\varphi^{-\mathbf{x}}$ obtained from φ by dropping $\mathbf{x}$.

Exponential lower bound for uniform interpolants in K

Let $\mathbf{x} = x_1, \ldots, x_n$ and $\mathbf{p} = p_1, \ldots, p_n$. We define φ such that $\exists^{\sim}\mathbf{x}.\varphi$ says that there is a successor world and

not all satisfiable types *at* of literals over $\mathbf{p}$ are realized in a successor world.

Exponential lower bound for uniform interpolants in K

Let $\mathbf{x} = x_1, \ldots, x_n$ and $\mathbf{p} = p_1, \ldots, p_n$. We define φ such that $\exists^{\sim}\mathbf{x}.\varphi$ says that there is a successor world and

not all satisfiable types *at* of literals over $\mathbf{p}$ are realized in a successor world.

Define

$$\varphi = \bigwedge_{i=1}^{n} (x_i \leftrightarrow \Diamond x_i \leftrightarrow \Box x_i)) \wedge \Box \bigvee_{i \leq n} (\neg(x_i \leftrightarrow p_i)$$

Exponential lower bound for uniform interpolants in K

Let $\mathbf{x} = x_1, \ldots, x_n$ and $\mathbf{p} = p_1, \ldots, p_n$. We define φ such that $\exists^{\sim}\mathbf{x}.\varphi$ says that there is a successor world and

not all satisfiable types *at* of literals over $\mathbf{p}$ are realized in a successor world.

Define

$$\varphi = \bigwedge_{i=1}^{n}(x_i \leftrightarrow \Diamond x_i \leftrightarrow \Box x_i)) \wedge \Box \bigvee_{i \leq n}(\neg(x_i \leftrightarrow p_i)$$

So

$$\Diamond \top \wedge \neg(\bigwedge_{at} \Diamond at)$$

is the uniform interpolant for φ and $\mathbf{p}$.

Exponential lower bound for uniform interpolant in K

Assume there is a uniform interpolant χ with number of subformulas $< 2^n$. Then

$$\chi \equiv \Diamond\top \wedge \neg(\bigwedge_{at} \Diamond at)$$

We can refute χ in some M, w in which w has a successor. By the finite model property proof for K there is M', w with

- $M', w \not\models \chi$.
- at least one but $< 2^n$ successor nodes of w,

Exponential lower bound for uniform interpolant in K

Assume there is a uniform interpolant χ with number of subformulas $< 2^n$. Then

$$\chi \equiv \Diamond\top \wedge \neg(\bigwedge_{at} \Diamond at)$$

We can refute χ in some M, w in which w has a successor. By the finite model property proof for K there is M', w with

- $M', w \not\models \chi$.
- at least one but $< 2^n$ successor nodes of w,

Then M' does not realize some at in any successor of w. So $M', w \models \Diamond\top \wedge \neg(\bigwedge_{at} \Diamond at)$. Contradiction.

Size of Craig interpolants

It remains open whether one can prove an exponential lower bound on the size of Craig interpolants, if the size of a formula is the defined as the number of its subformulas.

Size of Craig interpolants

It remains open whether one can prove an exponential lower bound on the size of Craig interpolants, if the size of a formula is the defined as the number of its subformulas.

If $|\varphi|$ is the number of symbols in φ, we obtain an exponential lower bound for Craig interpolants using, for instance,

Theorem [van Ditmarsch, Iliev] In ML, ∇ is exponentially more succinct than $\Diamond$.

(Represent the witness formulas using abbreviations for $\nabla\Phi$.)

Uniform interpolants for global consequence

Let $\varphi \models_{glo} \psi$ if

$$M \models \varphi \quad \Rightarrow \quad M \models \psi$$

We have seen that no uniform interpolant exists for

$$\varphi = (A \to B) \wedge (B \to \Diamond B), \quad \Sigma = \{A\}$$

Uniform interpolants for global consequence

Let $\varphi \models_{glo} \psi$ if

$$M \models \varphi \quad \Rightarrow \quad M \models \psi$$

We have seen that no uniform interpolant exists for

$$\varphi = (A \to B) \land (B \to \Diamond B), \quad \Sigma = \{A\}$$

Theorem [Lutz and W, 2011] Uniform interpolant existence is 2ExpTime complete for the global consequence. If a uniform interpolant exists, then there exists one of triple exponential size.

Uniform interpolants for global consequence

Let $\varphi \models_{glo} \psi$ if

$$M \models \varphi \quad \Rightarrow \quad M \models \psi$$

We have seen that no uniform interpolant exists for

$$\varphi = (A \rightarrow B) \wedge (B \rightarrow \Diamond B), \quad \Sigma = \{A\}$$

Theorem [Lutz and W, 2011] Uniform interpolant existence is 2ExpTime complete for the global consequence. If a uniform interpolant exists, then there exists one of triple exponential size.

Lots of work on computing uniform interpolants in description logic using resolution-based methods (Schmidt, Koopmann and others).