Saturation-Based Uniform Interpolation for Multi-Modal Logics

Ruba Alassaf, Renate A. Schmidt, and Uli Sattler

University of Manchester
Oxford Road, Manchester
United Kingdom

Abstract

Uniform interpolation has been the subject of many recent research papers due to its link to Craig interpolation and its potential use in knowledge-based and agent-based systems. In this paper, we present a saturation-based system that computes a *local* uniform interpolant for a formula and a "keep" signature in the multi-modal logic K_n. The system works by exhaustively applying a set of rules to generate a sufficient number of local consequences, which are then filtered to remove those that contain symbols outside the keep signature. We show that the system is guaranteed to terminate and is sound and uniform interpolation complete. We further prove that we can extend the system to compute uniform interpolants for formulas in multi-modal logics of serial and reflexive frames D_n and T_n.

Keywords: Uniform Interpolation, Resolution, Bisimulations

1 Introduction

A formula ϕ' is said to be a *uniform* Σ*-interpolant* of ϕ, if for any ψ over the signature Σ, ψ is implied by ϕ' if and only if ψ is implied by ϕ. Uniform interpolation is closely related to the notion of Craig interpolation. In Craig interpolation, two formulae ϕ and ψ such that ϕ implies ψ are given, and the task is to compute an intermediate formula ϕ' such that ϕ implies ϕ' and ϕ' implies ψ. A logic is said to have the uniform interpolation (respectively, Craig interpolation) property if, for any formula and signature (respectively, two formulas), a uniform interpolant (respectively, Craig interpolant) can be computed.

The notion of uniform interpolation is stronger than Craig interpolation for two reasons. The first is that any logic that has the uniform interpolation property also has the Craig interpolation property. The second is that uniform interpolation can be used to compute Craig interpolants. This can be achieved by "keeping" consequences over the shared signature during interpolation.

Regarding applications, uniform interpolation and Craig interpolation have been investigated for a range of modal logics, and the closely related description logics, which underlie ontology languages such as OWL [2]. Ontology engineers benefit from uniform interpolation in ontology debugging, versioning and summarisation [17,13]. In agent-based applications, uniform interpolation is used to update the knowledge of an agent by making them ignorant of certain propositional formulas [1]. Knowledge-sharing applications may use uniform interpolation to facilitate knowledge exchange among agents with different domain specialisations [20].

Several variants of the uniform interpolation problem have been studied. For classical logic, the problem is reduced to the second-order quantifier elimination problem [7,8], which aims to eliminate given predicate symbols from the formula. Uniform interpolation has been investigated in description logic sometimes under the name of deductive forgetting. Uniform interpolation and forgetting are dual notions as the former aims to compute a formula by keeping a signature Σ, while the aim of latter is to compute a formula that eliminates the complement of Σ. Most studies have focused on TBox forgetting [17,13,14,16] sometimes with an ABox [15], with [19] considering on concept forgetting.

The difference between TBox forgetting and concept forgetting (or forgetting for a local modal K_n formula) is that a TBox is a set of axioms that are universally quantified, whereas a concept or a modal K_n-formula is local, i.e., a modal K_n-formula is instantiated to a particular set of individuals/worlds. Uniform interpolation of a TBox is thus a global interpolation problem whereas uniform interpolation of a concept or a K_n-formula is a local interpolation problem. There are also differences in the complexity of the two problems. It was shown that TBox forgetting is triple exponential in the size of the input [17], whereas concept forgetting is in ExpSpace [19].

We are interested in the local form of uniform interpolation for modal logics. The modal logic K was shown to have the uniform interpolation property via constructive proofs [9,21]. An implementable approach to constructing uniform interpolants was given in [3] for the modal logics K and T. Wolter [22] proved that the modal logic $S5$ has the uniform interpolation property, and that uniform interpolation for any normal mono-modal logic can be generalised to its multi-modal case. Recently, it was shown that $K45_n$ and $KD45_n$ have the uniform interpolation property [5]. It is known that $S4$ and $K4$ do not have the uniform interpolation property [10,3].

The aim of our research is to develop resolution-based systems for computing local uniform interpolants in modal logics that are suitable for implementation. In this paper, we present such a system for the multi-modal logics K_n, D_n, and T_n. We prove that the system is guaranteed to terminate, is sound and uniform interpolation complete. We established the complexity of the system, and discuss the relationship to other works. The main contributions of this paper are:

- The system is the first provably correct resolution-based system for computing uniform interpolants in the modal logics K_n, D_n and T_n. Our work extends and improves the work presented in [11] for modal logic K.
- The idea of our completeness proof is novel for resolution-based uniform interpolation systems. Completeness proofs for previous systems which

use resolution are based on proof-theoretic arguments [11,14,15,16]. Our completeness proof is based on a model-theoretic argument in which bisimulation takes a central role.

2 Getting Started

We assume the reader is familiar with the multi-modal logic K_n [4,12]. We fix $\mathcal{A}$ a set of strings, and $P = \{p, q, r, \dots\}$ a countable possibly infinite set of propositional symbols. A K_n formula ϕ over a given signature is defined inductively as follows: $\phi ::= p \mid \top \mid \bot \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \Box_a\phi \mid \Diamond_a\phi$ where $a \in \mathcal{A}$.

We use $\mathcal{F} = (\mathcal{W}, R)$ to denote a Kripke frame where $\mathcal{W}$ is a nonempty set of worlds and R a mapping from elements of $\mathcal{A}$ to binary relations over $\mathcal{W}$. To abbreviate notation, we use R_a instead of $R(a)$, and we say that u is *a-accessible* from w if $R_a(w, u)$. We use $\mathcal{M} = (\mathcal{W}, R, V)$ to denote a Kripke model where $(\mathcal{W}, R)$ is a Kripke frame, and V is a valuation function that assigns each propositional symbol p in P a subset $V(p)$ of $\mathcal{W}$.

A formula ϕ is (locally) *satisfiable in a model* $\mathcal{M}$, if there is a world w in $\mathcal{W}$ at which ϕ is true. We use $\mathcal{M}, w \models \phi$ to denote that ϕ is true at w in $\mathcal{M}$. A formula ϕ is (unconditionally) *satisfiable* if it is true at some world in some model. A formula ϕ is globally satisfied (or true) in a model $\mathcal{M}$, denoted $\mathcal{M} \models \phi$, if it is true at every w in $\mathcal{M}$. A formula ϕ is valid, denoted $\models \phi$, if it is satisfied in all models over any frame $\mathcal{F}$. A set of formulae N is globally satisfied by a model $\mathcal{M}$, denoted $\mathcal{M} \models N$, if for each formula ϕ in N, $\mathcal{M}$ globally satisfies ϕ.

In this paper, we consider two logics which extend K_n, namely D_n and T_n, the multi-modal logics of serial and reflexive frames, respectively. In a D_n model, each accessibility relation R_a is serial, i.e. for each $w \in \mathcal{W}$, and each $a \in \mathcal{A}$, there exists a $w' \in \mathcal{W}$ such that $R_a(w, w')$ is true in the model. In a T_n model, each accessibility relation R_a is reflexive, i.e. for each $w \in \mathcal{W}$, and each $a \in \mathcal{A}$, $R_a(w, w)$ is true in the model. The modal logic D_n is axiomatized by the axioms of K_n and the axiom schema $(D) = \Box_a\phi \to \Diamond_a\phi$. Similarly, T_n is axiomatized by adding the axiom schema $(T) = \Box_a\phi \to \phi$.

We are interested in the problem of computing local uniform interpolant of a formula and a signature.

Definition 2.1 Given a formula ϕ, a *uniform interpolant* of ϕ with respect to a signature Σ of propositional symbols is a formula ϕ' such that:
(i) ϕ' does not contain symbols outside of Σ, and
(ii) for any modal formula ψ over Σ, we have that for all models $\mathcal{M}$, $\mathcal{M} \models \phi \to \psi$ iff for all models $\mathcal{M}$, $\mathcal{M} \models \phi' \to \psi$.

3 Uniform Interpolation Method

We start with a high-level description of our uniform interpolation method for multi-modal logic K_n. Without loss of generality, we assume that the input formula ϕ is given in negation normal form.

Overview. The system is based on resolution with modal logic adaptations. The idea behind our approach is the following: for each propositional symbol x outside the given signature Σ, we generate a sufficient set of clauses for the given formula and subsequently eliminate any formulae that contain x. We repeat the process for all propositional symbols outside Σ.

The system uses special *world symbols*, or *W-symbols* for short, which are propositional symbols that help in two related ways:

(i) They are used to flatten the input formula to surface some parts of it. E.g., $\Box(\psi \vee \Diamond\phi)$ becomes $\Box W_1$, $W_1 \Rightarrow \psi \vee \Diamond W_2$ and $W_2 \Rightarrow \phi$.

(ii) They allow the inferences to be restricted to subformulae labelled with the same W-symbol. E.g., $\Box(x \wedge (\neg x \vee p))$ becomes $\Box W$, $W \Rightarrow x$ and $W \Rightarrow \neg x \vee p$. Later on, we see that one of our rules allows us to apply a resolution step on x.

Initially, we can think of W-symbols as constants representing worlds in a labelled tableau algorithm.

For a formula ϕ, a signature Σ, and an ordering $\succ$ over the symbols outside the input signature Σ, the system is provided a clause set $N_0 = \{W_0 \Rightarrow \phi\}$ as input, and applies its rules exhaustively to the formulae in the set until no rules can be applied, resulting in a clause set of the form $N_n = \{W_0 \Rightarrow \phi_1, ..., W_0 \Rightarrow \phi_m\}$. The formula $\phi' = \phi_1 \wedge \ldots \wedge \phi_m$ is then a uniform Σ-interpolant of ϕ, which is proved later.

The role of W_0 is to represent a specific world that satisfies the given formula ϕ. Any model $\mathcal{M}$ that satisfies ϕ at point w can be extended in a non-vacuous way to one that satisfies W_0 and $W_0 \Rightarrow \phi$ by setting $w \in V(W_0)$. In this extended model, $W_0 \Rightarrow \phi$ is globally witnessed as non-vacuously true.

The process of constructing a uniform interpolant is iterative with respect to the symbols outside Σ, and the ordering $\succ$ fixes the order in which these symbols are eliminated. For some uniform interpolation problems, a good ordering may allow the system to solve a problem in far fewer steps. For simplicity, and since the ordering does not improve any worst-case complexity results, we can assume this ordering is arbitrary. We use x to denote the maximal propositional symbol occurring in the current clause set N_i at the ith step.

The System. The rules of our uniform interpolant system are given in Figures 1, 2 and 3. Each rule has a premise, some conditions and a conclusion. The rules are structured with the premise above a horizontal line and the conclusion below it. The premise (respectively conclusion) can be one or more clauses depending on which rule is being applied. There are three types of rules in the system: preprocessing rules, resolution rules, and elimination rules.

The preprocessing rules and the elimination rules are replacement rules; they replace the premise in the current working clause set with the conclusion. The resolution rules are saturation rules; they keep the premise and extend the clause set with the conclusion. The rules can be applied in any order as long as the conditions for each rule are met.

The clauses obtained and handled by our system are in a normal form. They are all labelled with a W-symbol in the condition of the implication. We can have a formula or another W-symbol in the consequence of the implication. Concretely, for some W-symbols W_i and W_j, and some modal formula ψ, a clause can have the form $W_i \Rightarrow \psi$ or $W_i \Rightarrow W_j$. If ψ is a disjunction of modal formulae, we assume that it is a set, i.e., there is no repetition among the disjuncts. This is essential for the correctness of the method. We use the notation $\Rightarrow$, in contrast to $\rightarrow$, to highlight that an implication is generated by our system to maintain our normal form. Semantically, they are identical.

To describe the different types of W-symbols, we introduce some terminology and the functions *Def* and *Corr* which are used in the conditions of our system, and later on in the proofs.

Definition 3.1 Given a set N of clauses, the set S_w is the set of W-symbols introduced for subformulas appearing under a modal operator via the WI rule. We call these symbols *base W-symbols*.

We use an injective function Def that maps base W-symbols to subformulas of clauses in N, and that is extended each time we introduce a new W-symbol.

The set C_w is the set of W-symbols introduced by the Res $\Box\bigcirc$ rule. We call these symbols *combinatory W-symbols*.

We define a function $Corr$ that maps W-symbols to subsets of S_w as follows:

$$Corr(W_i) = \begin{cases} \{W_i\}, & \text{if } W_i \in S_w \\ Corr(W_n) \cup Corr(W_m), & \text{if } W_i \in C_w \text{ where } W_n \text{ and } W_m \text{ come} \\ & \text{from the premise of the Res } \Box\bigcirc \\ & \text{rule that has introduced } W_i. \end{cases}$$

It is easy to see that the definition of $Corr$ is well-founded. Intuitively, a base W-symbol is introduced to represent a subformula, and a combinatory W-symbol can be seen as a unique representative of a subset of the base W-symbols.

We now describe the three groups of rules which together make up our system. We use N to refer to the current working clause set. We assume that x is the current symbol we would like to eliminate, i.e., it is the maximal symbol in the current set with respect to the given ordering $\succ$. The W-symbol W_i is the ith W-symbol introduced during the inference process.

Preprocessing. The purpose of the preprocessing rules is to apply transformations to the members of the working clause set so that they can be handled by the other rules. Generally, the idea is to surface symbols appearing in ϕ that are not in Σ, i.e., to surface x in ϕ. The rules are applied in a lazy manner which means their application can be deferred to whenever they are necessary. The preprocessing rules are provided in Figure 1. The clausification rule distributes disjunction over conjunction. The world introduction rule performs structural transformation that flattens modal formulae.

Resolution. The purpose of the resolution rules is to deduce a sufficient number of clauses/formulas to generate a uniform interpolant. The rules are given in Figure 2.

Clausification:

$$\frac{N, W_i \Rightarrow (\phi_1 \wedge \phi_2) \vee \phi_3}{N, W_i \Rightarrow \phi_1 \vee \phi_3, W_i \Rightarrow \phi_2 \vee \phi_3}$$

provided that either ϕ_1 or ϕ_2 contain x. ϕ_3 may be empty.

World Introduction (WI):

$$\frac{N, W_i \Rightarrow \bigcirc_a \phi_1 \vee \phi_2}{N, W_i \Rightarrow \bigcirc_a W_j \vee \phi_2, W_j \Rightarrow \phi_1}$$

provided that

(i) $\bigcirc \in \{\Box, \Diamond\}$,
(ii) ϕ_1 must contain x,
(iii) if ϕ_2 contains x then x must occur under a modal operator, and
(iv) if there is a W_k such that $Def(W_k) = \phi$, then $W_j = W_k$, otherwise W_j is a fresh W-symbol, $Corr(W_j) = \{W_j\}$ and $Def(W_j) = \{\phi\}$.

ϕ_2 may be empty.

Fig. 1. The preprocessing rules of the UI_{K_n} system for the modal logic K_n.

Literal Resolution (Res):

$$\frac{W_i \Rightarrow \psi_1 \vee x \quad W_i \Rightarrow \psi_2 \vee \neg x}{W_i \Rightarrow \psi_1 \vee \psi_2}$$

ψ_1 and/or ψ_2 may be empty.

World Resolution (Res W):

$$\frac{W_i \Rightarrow \psi \quad W_j \Rightarrow W_i}{W_j \Rightarrow \psi}$$

provided that ψ contains x.

$\Box\bigcirc$ **Resolution** (Res $\Box\bigcirc$):

$$\frac{W_i \Rightarrow \psi_1 \vee \Box_a W_n \quad W_i \Rightarrow \psi_2 \vee \bigcirc_a W_m}{W_i \Rightarrow \psi_1 \vee \psi_2 \vee \bigcirc_a W_j,\ W_j \Rightarrow W_n,\ W_j \Rightarrow W_m}$$

provided that:

(i) $\bigcirc \in \{\Box, \Diamond\}$,
(ii) $Corr(W_n) \cap Corr(W_m)$ is empty,
(iii) if there is a W_k such that $Corr(W_k) = Corr(W_n) \cup Corr(W_m)$ then $W_j = W_k$, otherwise W_j is a fresh W-symbol, and $Corr(W_j) = Corr(W_n) \cup Corr(W_m)$.

ψ_1 and/or ψ_2 may be empty.

Fig. 2. The resolution rules of the UI_{K_n} system for modal logic K_n.

The literal resolution rule is the heart of our system; it computes a formula by resolving on a maximal symbol x if the premise is labelled with the same W-symbol. The world resolution rule is used to propagate formulas labelled by another W-symbol, which is essentially a resolution step between world symbols. The $\Box\bigcirc$ resolution rule is used to capture combinations of successor relations. The second and third conditions are the blocking conditions; they aim to ensure that the rule application is not redundant which is important

Positive Purification (+PUR):

$$\frac{N, W_i \Rightarrow \psi \vee x}{N, W_i \Rightarrow \psi \vee \top}$$

provided that no more non-purification rules can be applied to the named clause in the premise. ψ may be empty.

Negative Purification (−PUR):

$$\frac{N, W_i \Rightarrow \psi \vee \neg x}{N, W_i \Rightarrow \psi \vee \top}$$

provided that no more non-purification rules can be applied to the named clause in the premise. ψ may be empty.

World Elimination (ELM W):

$$\frac{N, W_i \Rightarrow \psi_1, \ldots, W_i \Rightarrow \psi_n}{N^{W_i}_{(\psi_1 \wedge \cdots \wedge \psi_n)}}$$

provided that $i \neq 0$, $\psi_1, \ldots, \psi_n$ do not contain x or any W-symbol, and N only contains W_i on the right hand side of $\Rightarrow$ clauses. The expression N^{ϕ}_{ψ} denotes the set of clauses that is obtained by replacing each occurrence of ϕ in N by ψ.

Fig. 3. The purification and elimination rules of the UI_{K_n} system for modal logic K_n.

to control the complexity, and that the system does not infinitely introduce W-symbols which is essential for termination.

Elimination. The elimination rules are responsible for eliminating symbols outside of $\Sigma \cup \{W_0\}$. They are applied once we have exhaustively applied the resolution rules to compute conclusions over Σ. The rules are given in Figure 3.

The positive and negative purification rules replace a maximal symbol x, occurring either positively or negatively, with $\top$. The world elimination rule collects modal formulas labelled with the same W-symbol, and replaces right hand side occurrences of the W-symbol with the conjunction of these formulas, effectively eliminating the W-symbol from the set of clauses.

3.1 Examples

In the following examples, we demonstrate how the UI_{K_n} system is used to compute a uniform interpolant with respect to $\Sigma = \{p, q\}$. Starting from $i = 0$, we use N_i to refer to the clause set that is obtained after applying the ith step in the derivation.

Example 3.2 Consider a formula $\phi = (\neg p \vee \Diamond x) \wedge (\neg x \vee \Box q)$.

The input to the system is the set $N_0 = \{W_0 \Rightarrow (\neg p \vee \Diamond x) \wedge (\neg x \vee \Box q)\}$. The only rule applicable to N_0 is the clausification rule which gives

$$N_1 = \{W_0 \Rightarrow \neg p \vee \Diamond x, W_0 \Rightarrow \neg x \vee \Box q\}.$$

Now we apply the world introduction rule to get

$$N_2 = \{W_0 \Rightarrow \neg p \vee \Diamond W_1, W_1 \Rightarrow x, W_0 \Rightarrow \neg x \vee \Box q\}.$$

The only applicable rules are the positive and negative purification rules. We achieve

$$N_3 = \{W_0 \Rightarrow \neg p \vee \Diamond W_1, W_1 \Rightarrow \top, W_0 \Rightarrow \top \vee \Box q\}.$$

Eliminating W_1, we obtain

$$N_4 = \{W_0 \Rightarrow \neg p \vee \Diamond \top, W_0 \Rightarrow \top \vee \Box q\}.$$

The Σ-uniform interpolant is $\phi' = (\neg p \vee \Diamond \top) \wedge (\top \vee \Box q)$.

Notice that this example illustrates the local flavour of the system. We see that the occurrences of x at two different modal levels do not interact via any resolution rule.

Example 3.3 Consider a formula $\phi = (\neg p \vee \Diamond x) \wedge \Box(\neg x \vee \Box q)$. We start with the set $N_0 = \{W_0 \Rightarrow (\neg p \vee \Diamond x) \wedge \Box(\neg x \vee \Box q)\}$. Applying the clausification rule to N_0 we get

$$N_1 = \{W_0 \Rightarrow \neg p \vee \Diamond x, W_0 \Rightarrow \Box(\neg x \vee \Box q)\}.$$

By applying the world introduction rule twice, we have

$$N_3 = \{W_0 \Rightarrow \neg p \vee \Diamond W_1, W_1 \Rightarrow x, W_0 \Rightarrow \Box W_2, W_2 \Rightarrow \neg x \vee \Box q\}.$$

The only applicable rule is the $\Box\Diamond$ rule, and it yields

$$N_4 = N_3 \cup \{W_0 \Rightarrow \neg p \vee \Diamond W_3, W_3 \Rightarrow W_1, W_3 \Rightarrow W_2\}.$$

By applying the world resolution rule twice, we obtain

$$N_6 = N_4 \cup \{W_3 \Rightarrow x, W_3 \Rightarrow \neg x \vee \Box q\}.$$

Now, we can apply the literal resolution rule which yields

$$N_7 = N_6 \cup \{W_3 \Rightarrow \Box q\}.$$

We apply the positive and negative purification rules (4 applications) and achieve

$$\begin{aligned} N_{11} = \{\, & W_0 \Rightarrow \neg p \vee \Diamond W_1, && W_1 \Rightarrow \top, && W_0 \Rightarrow \Box W_2, \\ & W_2 \Rightarrow \top \vee \Box q, && W_0 \Rightarrow \neg p \vee \Diamond W_3, && W_3 \Rightarrow W_1, \\ & W_3 \Rightarrow W_2, && W_3 \Rightarrow \top, && W_3 \Rightarrow \top \vee \Box q, \\ & W_3 \Rightarrow \Box q \,\}. \end{aligned}$$

Now, x does not appear anywhere. We eliminate the world variables W_1, W_2, W_3 via the world elimination rule.

To eliminate W_1, we look for clauses labelled with W_1, in this case we only have $W_1 \Rightarrow \top$. We remove $W_1 \Rightarrow \top$ and replace each occurrence of W_1 on the right hand side of $\Rightarrow$ with $\top$ as follows:

$$\begin{aligned} N_{12} = \{\, & W_0 \Rightarrow \neg p \vee \Diamond \top, && W_0 \Rightarrow \Box W_2, && W_2 \Rightarrow \top \vee \Box q, \\ & W_0 \Rightarrow \neg p \vee \Diamond W_3, && W_3 \Rightarrow \top, && W_3 \Rightarrow W_2, \\ & W_3 \Rightarrow \top \vee \Box q, && W_3 \Rightarrow \Box q \,\}. \end{aligned}$$

Similarly for W_2, we remove $W_2 \Rightarrow \top \vee \Box q$, and replace the other occurrences of W_2 with $\top \vee \Box q$.

$$N_{13} = \{\, W_0 \Rightarrow \neg p \vee \Diamond\top, \quad W_0 \Rightarrow \Box(\top \vee \Box q), \quad W_0 \Rightarrow \neg p \vee \Diamond W_3, \quad W_3 \Rightarrow \top, \quad W_3 \Rightarrow \top \vee \Box q, \quad W_3 \Rightarrow \Box q \,\}.$$

Finally, we eliminate W_3,

$$N_{14} = \{\, W_0 \Rightarrow \neg p \vee \Diamond\top, \quad W_0 \Rightarrow \Box(\top \vee \Box q), \quad W_0 \Rightarrow \neg p \vee \Diamond(\top \wedge (\top \vee \Box q) \wedge \Box q) \,\}.$$

The uniform interpolant is

$$\phi' = (\neg p \vee \Diamond\top) \wedge (\Box(\top \vee \Box q)) \wedge (\neg p \vee \Diamond(\top \wedge (\top \vee \Box q) \wedge \Box q)),$$

which is equivalent to $\phi' = (\neg p \vee \Diamond\Box q)$ by standard simplifications.

4 Correctness and Complexity

The first lemmas in this section are relevant to termination. We prove termination by showing that any derivation uses a finite number of symbols, and we argue that because of this, the system will stop generating new clauses.

Lemma 4.1 *The number of W-symbols introduced in a run of the UI_{K_n} system is bounded by $2^n - 1$ for n being the length of the input.*

Proof. Let S_w be the set of base W-symbols. S_w is finite because the number of modal operators in the input formula is finite, and the role of the world introduction rule is to replace each subformula (not containing a W-symbol) appearing under a modal operator with a W-symbol.

Let C_w be the set of combinatory W-symbols introduced by the RES $\Box\bigcirc$ rule. We show inductively that for each W-symbol W_i in C_w, the set $Corr(W_i)$ corresponds to a unique set of W-symbols from S_w.

Since the function $Corr$ has a finite range, which is the powerset of S_w, we can prove that the domain is finite by showing that $Corr$ is injective. Let W_n and W_m be two W-symbols. We show inductively that if $Corr(W_n) = Corr(W_m)$, then $W_n = W_m$. The proof is given in the appendix.

Since S_w is a finite set, and by condition (ii) of the $\Box\bigcirc$ resolution rule, the number of possible W-symbols is bounded by the number of unique combinations of symbols from S_w. The upper bound is equal to $2^{|S_w|}$-1, where $|S_w|$ is the size of S_w. □

Lemma 4.2 *The UI_{K_n} system will stop generating new clauses.*

Proof. The input formula ϕ contains a finite number of propositional symbols and modal operators. By Lemma 4.1, the system introduces a finite number of W-symbols. The resolution rules do not produce results with increased modal depth. □

Lemma 4.3 *The UI_{K_n} system will not reintroduce a W-symbol that was eliminated before.*

Proof. This is because a W-symbol is introduced to surface a formula that has symbols not in Σ, and only when a W-symbol no longer labels non-Σ symbols, the world elimination rule is allowed to be applied. □

From Lemmas 4.1, 4.2 and 4.3, we conclude the following theorem.

Theorem 4.4 (Termination) *Given a formula ϕ and a signature Σ, the UI_{K_n} system computes a formula ϕ' in a finite number of steps.*

The following lemma considers the space complexity of our system.

Lemma 4.5 *The space complexity of the UI_{K_n} calculus is double exponentially bounded in the length of the input.*

Proof. The root source of the complexity is that the system generates an exponential number of combinatory symbols. The complexity argument is built upon three claims. Using a linear number of propositional symbols, a linear number of base W-symbols, and an exponential number of combinatory W-symbols:

- C1 the size of each clause, before eliminating any W-symbols, has an exponential upper bound in the size of the input,
- C2 the number of clauses we could generate has a double exponential upper bound in the size of the input, and
- C3 the size of each clause, after eliminating W-symbols, has a double exponential upper bound in the size of the input.

We show the three claims in the appendix. □

The next lemmas show that the signature of ϕ' is Σ.

Lemma 4.6 *For a given formula ϕ and a signature Σ, the UI_{K_n} system will always be able to eliminate every W-symbol that is not W_0, using the world elimination rule.*

Proof. The normal form that is used in our system dictates that only one W-symbol can be present on the left hand side of $\Rightarrow$. Eliminating a W-symbol will replace occurrences of a W in the working clause set with a combination of clauses that do not contain W-symbols. The only situation that may prevent the elimination rule from being applied is if a clause contains the same W-symbol on both sides of the $\Rightarrow$, but it can be shown inductively that this cannot happen. □

Lemma 4.7 *For a given formula ϕ and a signature Σ, the UI_{K_n} system will always be able to eliminate symbols in the signature of ϕ that are not in Σ.*

Proof. Let x be the maximal symbol in the signature of ϕ but outside Σ. The world introduction rule is a replacement rule that aims to surface x. When no resolution rule is applicable, the calculus applies a purification rule that eliminates x. □

Theorem 4.8 (Soundness) *Given a formula ϕ and a signature Σ, the UI_{K_n} system computes a formula ϕ' such that for any formula ψ over Σ, we have that if $\models \phi' \to \psi$ then $\models \phi \to \psi$.*

Proof. Consider ϕ, ϕ' as in Theorem 4.8. We use $\mathcal{M}|_\Sigma$ for the reduction of $\mathcal{M}$ to Σ, i.e. the model obtained from $\mathcal{M}$ by ignoring all symbols outside Σ.

We say that two models $\mathcal{M}', \mathcal{M}$ are Σ-*inseparable* if $\mathcal{M}'|_\Sigma = \mathcal{M}|_\Sigma$. To show the theorem, we use the following claims:

C4 For each model $\mathcal{M}$ and world w_0 such that $\mathcal{M}, w_0 \models \phi$, there exists a model $\mathcal{M}'$ such that $\mathcal{M}$ and $\mathcal{M}'$ are Σ-inseparable, $\mathcal{M}' \models W_0 \Rightarrow \phi$, and $V(W_0) = \{w_0\}$.

C5 Each rule in our calculus is Σ-preserving, that is, if a rule is applied to a set of clauses N to produce N' and $\mathcal{M}$ is a model that globally satisfies N, then there exists a model $\mathcal{M}'$ that is Σ-inseparable from $\mathcal{M}$ that satisfies N'.

C6 Let $N = \{W_0 \Rightarrow \psi_i \mid 1 \leq i \leq n\}$ be the set obtained at the end of our derivation. For each model $\mathcal{M}$ that globally satisfies N and for each $w_0 \in V(W_0)$, we have $\mathcal{M}, w_0 \models \phi'$ with $\phi' = \bigwedge_{i \leq |N|} \psi_i$.

The claims are shown in the appendix.

Via C4, C5, and C6, we can show that for any model of the input formula, there exists a Σ-inseparable model for the output formula. And, thus, for any formula ψ over Σ, if $\phi \to \psi$ is true in all models, then $\phi' \to \psi$ is true in all models. □

For our completeness proof, we are interested in understanding models that are invariant up to the satisfaction of Σ-modal formulas. Σ-modal formulas are modal formulas described using a signature of propositional symbols Σ. For this purpose, we use the following notion.

Definition 4.9 Let $(\mathcal{M}, w)$ and $(\mathcal{M}', w')$ be two Kripke models where $\mathcal{M} = (\mathcal{W}, R, V)$ and $\mathcal{M}' = (\mathcal{W}', R', V')$. A Σ-*bisimulation* between $\mathcal{M}$ and $\mathcal{M}'$ is a relation $\rho \subseteq \mathcal{W} \times \mathcal{W}'$ such that $w \rho w'$, and whenever $u \rho u'$, the following holds:
atoms u and u' satisfy the same propositional symbols from Σ;
forth For all a, if $uR_a t$, then there is a t' such that $u'R'_a t'$ and $t \rho t'$;
back For all a, if $u'R'_a t'$, then there is a t such that $uR_a t$ and $t \rho t'$.

The following is our completeness theorem.

Theorem 4.10 (Completeness) *Given a K_n formula ϕ and a signature Σ, the UI_{K_n} system computes a K_n formula ϕ' such that, for any K_n formula ψ over Σ, we have that if $\models \phi \to \psi$ then $\models \phi' \to \psi$.*

Proof. Consider ϕ, ϕ' and ψ as in Theorem 4.10. Assume $\models \phi \to \psi$ but $\not\models \phi' \to \psi$. The assumption implies that there exists a counter model $\mathcal{M}'$ and a world w_0 such that, $\mathcal{M}', w_0 \models \phi'$ and $\mathcal{M}', w_0 \not\models \psi$.

The idea of our proof is to inductively show that a Σ-bisimilar model $\mathcal{M}$ can be defined based on $\mathcal{M}'$ such that, $\mathcal{M}, w_0 \models \phi$, and $\mathcal{M}, w_0 \not\models \psi$. This will contradict our assumption $\models \phi \to \psi$.

Let n be the number of steps used for generating ϕ' from ϕ via our calculus. Let N_k denote the set of formulae obtained after applying k steps to ϕ. For each i, starting from $\mathcal{M}_n = \mathcal{M}'$, we will construct a model $\mathcal{M}_i = (\mathcal{W}_i, R_i, V_i)$ for N_i.

We construct $\mathcal{M}_i$ by inductively extending $\mathcal{M}_{i+1}$ and ensuring that $\mathcal{M}_i$ and $\mathcal{M}_{i+1}$ are Σ-bisimilar. In this view, $N_1 = \{W_0 \Rightarrow \phi\}$, and $N_n = \{W_0 \Rightarrow \phi'_1, \ldots, W_0 \Rightarrow \phi'_m\}$ where $\phi' = \phi'_1 \wedge \ldots \wedge \phi'_m$. This will lead to a counter-model

$\mathcal{M}_0 = \mathcal{M}$ as described above.

Before we state our claims, we introduce a terminology: given a set N of clauses and W-symbols W_i and W_j, W_i is said to be *directly related* to W_j if $W_i \Rightarrow W_j \in N$ or $W_j \Rightarrow W_i \in N$. Two W-symbols W_i and W_j are *indirectly related* if there is a W-symbol W_k such that $W_k \Rightarrow \gamma_1 \vee \bigcirc_a W_i$ and $W_k \Rightarrow \gamma_2 \vee \Box_a W_j$ where γ_1, γ_2 may be empty. We say that two W-symbols are *related* if they are directly related or indirectly related.

We prove the following claims simultaneously by backward induction: the first claim is at the heart of our completeness, and the other three claims are invariants that will help in our induction.

Claim 1: for all $k \leq n$, there exists a model $\mathcal{M}_k$ such that $\mathcal{M}_k, w_0 \models W_0$ and $\mathcal{M}_k \models N_k$ but $\mathcal{M}_k, w_0 \not\models \psi$.

Claim 2: for all $0 \leq k < n$, $\mathcal{M}_k$ and $\mathcal{M}_{k+1}$ are Σ-bisimilar.

Claim 3: for all $k \leq n$, $\mathcal{M}_k$ has the W-symbol independence property; that is that either two W-symbols are *related* in N_k or their interpretations are disjoint.

We assume w.l.o.g. that the valuation function in $\mathcal{M}'$ maps any symbol not in Σ to the empty set.

Base case ($k = n$): To show Claim 1, we define the model $\mathcal{M}_n$ as an extension of $\mathcal{M}'$, to satisfy N_n by setting the interpretation of W_0 to be true in w_0, i.e., $V_n(W_0) = \{w_0\}$. Claim 2 holds trivially because $k = n$, and Claim 3 holds because there is only one W-symbol and that is W_0.

Step case: Assuming all claims hold for $k + 1$, we will show that there exists a Σ-bisimilar model $\mathcal{M}_k$ such that $\mathcal{M}_k \models N_k$ but $\mathcal{M}_k, w_0 \not\models \psi$, and that this model has the W-symbol independence property.

To this end, we will consider the effect of each rule and show that if it was applied as the k-th step, it is possible to construct a model $\mathcal{M}_k$ from $\mathcal{M}_{k+1}$ that satisfies our claims.

For the preprocessing rules, except for the world introduction rule, the model remains the same. For the world introduction rule, we must reset the interpretation of the W-symbol that was introduced to maintain Claim 3 and 4, i.e., we make $V_k(W_j) = \emptyset$.

Similarly, for the literal resolution rule and the world resolution rule, the model remains the same, and the argument is that N_k is a subset of N_{k+1}, so a model that satisfies N_{k+1} satisfies N_k, but for the $\Box\bigcirc$ resolution rule, we must reset the interpretation of the W-symbol that was introduced to maintain Claim 3.

We give proofs for the world elimination and purification rules.

World Elimination: Let N_{k+1} be the set obtained after eliminating the symbol W_i from N_k.

Assume without loss of generality (w.l.o.g.) that the following formulas are the only formulas that have left hand side (l.h.s.) occurrences of W_i in N_k:

$$W_i \Rightarrow \gamma_1, \ldots, W_i \Rightarrow \gamma_m \tag{1}$$

The formulas where W_i can occur on the right hand side (r.h.s.) are of the forms:

$$W_j \Rightarrow W_i \quad \text{or} \quad W_j \Rightarrow \gamma \vee \bigcirc_a W_i \quad \text{where } \bigcirc \in \{\Diamond, \Box\}. \tag{2}$$

In N_{k+1}, i.e. after the application of the world elimination rule, the formulas in Equation (1) are removed. The formulas of the forms in Equation (2) are replaced with
$W_j \Rightarrow (\gamma_1 \wedge \cdots \wedge \gamma_m)$ and $W_j \Rightarrow \gamma \vee \bigcirc_a(\gamma_1 \wedge \cdots \wedge \gamma_m)$ where $\bigcirc \in \{\Diamond, \Box\}$ respectively. Given $\mathcal{M}_{k+1}$, a model for N_{k+1}, our aim is to expand the model into $\mathcal{M}_k$ for N_k by giving W_i an appropriate assignment.

By assumption, this assignment is empty, and to achieve an appropriate interpretation, we incrementally populate the assignment in the following way.

First, to satisfy a new formula of the form $W_j \Rightarrow W_i$ in N_k, we must expand the valuation mapping of W_i in the model being constructed $\mathcal{M}_k$ to include every world in the mapping of W_j. Explicitly, the mapping of W_i must include the following worlds: $\{w \mid w \in V_{k+1}(W_j)\}$.

Second, we consider the two cases for satisfying a formula of the form $W_j \Rightarrow \gamma \vee \bigcirc_a W_i$ where $\bigcirc \in \{\Diamond, \Box\}$.

Case 1: $\bigcirc$ denotes a $\Diamond$ operator. We consider each $w \in V_{k+1}(W_j)$, if $\mathcal{M}_{k+1}, w \models \Diamond_a(\gamma_1 \wedge \cdots \wedge \gamma_m)$ then there is a world u such that $wR_a u$ is in R_a, and $\mathcal{M}_{k+1}, u \models \gamma_1 \wedge \cdots \wedge \gamma_m$. We extend the domain $\mathcal{W}_k$ with a fresh world u'. We start by making the interpretation of u' in $\mathcal{M}_k$ identical to the interpretation of u for symbols from Σ, i.e., for all p in Σ, if u is in $V_{k+1}(p)$, then we include u' in $V_k(p)$. We extend the frame with successors in the following way: for all a, and all t, if $uR_a t$ is true in $\mathcal{M}_{k+1}$, then $u'R_a t$ is made true in $\mathcal{M}_k$, and if $tR_a u$ is true in $\mathcal{M}_{k+1}$, then $tR_a u'$ is made true in $\mathcal{M}_k$. This is to ensure that the new model $\mathcal{M}_k$ maintains Σ-bisimilarity. We extend the interpretation of $V_k(W_i)$ to include u'. To maintain Claim 3, we extend the interpretation of the V_k for W-symbols related to W_i in the following way: for each W_m, if u is in $V_k(W_m)$, and W_m is related to W_i then we include u' in the valuation of u'.

Case 2: $\bigcirc$ denotes a $\Box$ operator. Similar to the first case, we consider each $w \in V_{k+1}(W_j)$. If $\mathcal{M}_{k+1}, w \models \Box_a(\gamma_1 \wedge \cdots \wedge \gamma_m)$, then for every world u connected to w via an a-successor, we check if we can make W_i true in u while maintaining Claim 3, i.e. we check if the following condition holds: for all W_m, if u is in $V_{k+1}(W_m)$, then W_m and W_i must be related. If this is not possible, we extend the domain $\mathcal{W}_k$ with a fresh world u', and we define V_k as in the case 1, but the accessibility relation is defined as follows:

$$R_{a_k} = \begin{cases} wR_{a_k}u' & \text{if for some } w,\ wR_{a_{k+1}}u \in R_{a_{k+1}} \\ u'R_{a_k}w & \text{if for some } w,\ uR_{a_{k+1}}w \in R_{a_{k+1}} \\ wR_{a_k}w' & \text{if for some } w, w' \text{ s.t. } w \neq u \text{ and } w \neq w',\ wR_{a_{k+1}}w' \in R_{a_{k+1}} \end{cases}$$

The difference between the construction in case 1 and the one here is that in case 1, $\mathcal{M}_k$ is an extension of $\mathcal{M}_{k+1}$, whereas, in this case, this would not help because of the box operator, e.g. consider $\phi = \neg x \wedge \Box(p \wedge x)$, and let $\mathcal{M}_n = (\{w_0\}, \{R(w_0, w_0)\}, V_n)$ where $V_n(p) = \{w_0\}$.

Now, we have completed the construction of the model $\mathcal{M}_k$. It is clearly a model for N_k, and is Σ-bisimilar to $\mathcal{M}_{k+1}$. The world independence property is maintained as, in our construction, we ensure that if there is a W-symbol that is unrelated to W_i in some world u, they, and their related worlds, are separated using a new world u'.

Positive Purification: Let $W_i \Rightarrow \top \vee \gamma_1$ be the clause obtained after applying purification to $W_i \Rightarrow x \vee \gamma_1$ from N_k. We define a model $\mathcal{M}_k = (\mathcal{W}_{k+1}, R_{k+1}, V_k)$ where V_k extends the map of x in the following way

$$V_k(x) = V_{k+1}(x) \cup \{w \mid w \in V_{k+1}(W_i) \textit{ and } \mathcal{M}_{k+1}, w \not\models \gamma_1\}.$$

The deleted clause $W_i \Rightarrow x \vee \gamma_1$ is now true by the construction of $\mathcal{M}_k$. It remains to check that $\mathcal{M}_k \models N_k$. Assume $\mathcal{M}_k \not\models N_k$. Since the only change was an extension to the interpretation of x, there must be a clause which contains x negatively in N_k that was true in $\mathcal{M}_{k+1}$ but is not true in $\mathcal{M}_k$. This clause must be of the form $W_j \Rightarrow \neg x \vee \gamma_2$.

Since making the clause $W_i \Rightarrow x \vee \gamma_1$ true in $\mathcal{M}_k$ made $W_j \Rightarrow \neg x \vee \gamma_2$ false, we can infer that there is a common world w in $V_k(W_i) \cap V_k(W_j)$ such that $\mathcal{M}_k, w \models x \vee \gamma_1$ but $\mathcal{M}_k, w \not\models \neg x \vee \gamma_2$.

There are four cases to consider.

Case 1: $i = j$. By literal resolution, this implies that $W_i \Rightarrow \gamma_1 \vee \gamma_2$ is in N_k and N_{k+1}, and is satisfied by $\mathcal{M}_{k+1}$ but not by $\mathcal{M}_k$. This is impossible because the two models agree up to x. γ_1 and γ_2 may contain x but only under a modal operator which, by how we defined $\mathcal{M}_k$ for the World Elimination rule, means that they cannot be realised by w.

Case 2: W_i and W_j are directly related in N_k. Without loss of generality, let us assume $W_i \Rightarrow W_j$ is in N_k. This implies that (i) $V(W_i) \subseteq V(W_j)$, and (ii) by world resolution, $W_i \Rightarrow \neg x \vee \gamma_2$ is in N_k. The problem is now reduced to what has been discussed in the first case. Therefore, we can similarly conclude that this case cannot happen.

Case 3: W_i and W_j are indirectly related in N_k. Without loss of generality, assume there exists a W-symbol W_k such that $W_k \Rightarrow \gamma_1 \vee \bigcirc_a W_i$ and $W_k \Rightarrow \gamma_2 \vee \Box_a W_j$ are in N_k. By the $\Box\bigcirc$ resolution rule, this implies that $W_k \Rightarrow \gamma_1 \vee \gamma_2 \vee \bigcirc W_{ij}$, $W_{ij} \Rightarrow W_i$ and $W_{ij} \Rightarrow W_j$ is in N_k. This problem is now reduced to what has been discussed in case 2, and hence it cannot occur.

Case 4: $i \neq j$ and W_i and W_j are not related in N_k. By Claim 3, their interpretations are disjoint, i.e., $V_k(W_i) \cap V_k(W_j)$ is empty. Hence, this case cannot occur.

Finally, since we have not changed the structure of the model nor the interpretation of the elements in Σ, $\mathcal{M}_{k+1}$ and $\mathcal{M}_k$ are Σ-bisimilar. The third claim is maintained since the interpretation of the W-symbols has not changed.

As for the case of *negative purification*, we define a model $\mathcal{M}_k$ to be an identical copy of $\mathcal{M}_{k+1}$. What remains is to show that $\mathcal{M}_k \models W_i \Rightarrow \neg x \vee \gamma_1$. The proof is analogous to the proof given for positive purification. $\square$

5 Extensions

We consider extensions of K_n, namely D_n and T_n, and show that our UI_{K_n} system could be used or extended to cover them.

Multi-Modal Logic D_n. We claim that the UI_{K_n} system computes uniform interpolants for the multi-modal logic D_n. To show our claim, we prove the completeness theorem.

Theorem 5.1 (Completeness) *Given a D_n formula ϕ and a signature Σ, the UI_{K_n} system computes a D_n formula ϕ' such that, for any D_n formula ψ over Σ, we have that if $\models \phi \rightarrow \psi$ then $\models \phi' \rightarrow \psi$.*

Proof. The proof uses a similar argument to the one for K_n except, here we assume that the accessibility relations in the initial model are serial.

In addition to the three claims from the proof of Theorem 4.10, we prove the following claim:

Claim 4: For all $k \leq n$, each accessibility relation in $\mathcal{M}_k$ is serial.

Observe that $\mathcal{M}_k$ uses the frame underlying the model $\mathcal{M}_{k+1}$ in the proofs given for each of the rules, except for the world elimination rule. Therefore, for these rules, Claim 4 is established. Consider the case of the world elimination rule. We construct $\mathcal{M}_k$ to be Σ-bisimilar to $\mathcal{M}_{k+1}$. By the induction hypothesis, $\mathcal{M}_{k+1}$ is a serial model, and by the "forth" condition of definition of Σ-bisimilation, $\mathcal{M}_k$ must be serial, too. □

Multi-Modal Logic T_n. In this part, we introduce the UI_{T_n} system, and show that it is terminating, sound and uniform interpolation complete for T_n. The system extends UI_{K_n} by generalising the world resolution rule and adding the reflexivity rule as a new saturation rule. These rules are shown in Figure 4.

World Resolution (RES W):

$$\frac{W_i \Rightarrow \psi_1 \quad W_j \Rightarrow \psi_2 \vee W_i}{W_j \Rightarrow \psi_1 \vee \psi_2}$$

provided that ψ_1 contains x. ψ_2 may be a W-symbol.

Reflexivity (T):

$$\frac{W_i \Rightarrow \psi \vee \Box_a W_m}{W_i \Rightarrow \psi \vee W_m}$$

provided that W_m is a base symbol.

Fig. 4. The reflexivity rule and the world resolution rule of the UI_{T_n} system for modal logic T_n.

The soundness theorem and proof are analogous to Theorem 4.8 and its proof. The termination and complexity arguments are identical to the ones in Theorem 4.4 and Lemma 4.5. It remains to show the completeness theorem.

Theorem 5.2 (Completeness) *Given a formula ϕ and a signature Σ, the UI_{T_n} system computes a formula ϕ' such that, for any formula ψ over Σ, we have that if $\models \phi \rightarrow \psi$ then $\models \phi' \rightarrow \psi$.*

Proof. This proof is an extension of the proof of Theorem 4.10. We assume that each accessibility relation is reflexive in $\mathcal{M}'$. We generalise our definition of *directly related*. Given N a set of clauses and W-symbols W_i and W_j, W_i is said to be *directly related* to W_j if, for some γ, $W_i \Rightarrow \gamma \vee W_j \in N$ or $W_j \Rightarrow \gamma \vee W_i \in N$.

Our aim is to prove the same three claims and the following claim.

Claim 4: For all $k \leq n$, each accessibility relation R_i is reflexive in $\mathcal{M}_k$.

Base case ($k = n$): To show Claim 4, we observe that a change happened to the valuation of W_0, and that this does not change the frame underlying the model. Therefore, the new model $\mathcal{M}_n$ remains reflexive.

Because the set of reflexive models is a subset of the models considered in the proof for Theorem 4.10, the three claims remain true for all the shared rules. We focus on showing the claims for the reflexivity rule and the generalised world resolution rule, and show that Claim 4 is true for all the remaining rules.

Step case: Assuming all claims hold for $k + 1$, we show that there exists a model $\mathcal{M}_k$ that satisfies all four claims.

For all saturation rules, including the reflexivity rule and the world resolution rule, the model $\mathcal{M}_k$ is defined to be identical to $\mathcal{M}_{k+1}$. For the world elimination rule, we repeat the model construction of $\mathcal{M}_k$ as shown in the proof of Theorem 4.10, with the consideration that now we have the general form $W_j \Rightarrow \gamma \vee W_i$, but we extend it to make $R_a(u', u')$ true for each $R_a \in R_k$. Observe that this change still maintains the Σ-bisimulation, but will make $\mathcal{M}_k$ reflexive. Clearly, $\mathcal{M}_k$ is a model for N_k, and the world independence property is maintained.

For the positive and negative purification rules, observe that the frame underlying the model $\mathcal{M}_k$ is identical to that of $\mathcal{M}_{k+1}$, so Claim 4 holds. $\square$

6 Related Work

Fang et al. [6] use what Moss [18] called canonical formulas. They exploit the fact that an arbitrary modal formula is equivalent to a disjunction of a finite set of satisfiable canonical formulas, and prove that a uniform interpolant can be constructed via literal elimination. Although the proof is constructive, as the authors explicitly mention, the method is unpractical because the size of a canonical formula is non-elementary.

Other uniform interpolation methods can be divided into two groups based on whether they use a conjunctive normal form (CNF) (e.g., [11,15]) or a disjunctive normal form (DNF) (e.g., [3,19]). Uniform interpolants can be easily computed for formulae in disjunctive normal form, in fact the method in [19] was shown to have an exponential worst case space complexity, which they prove to be a tight upper bound.

Comparing the two types of approaches, we notice that methods that use the conjunctive normal form, which are often based on resolution, struggle with the following type of problem.

(i) $\Box(x \vee q1) \vee \Box(\neg x \vee q2)$ (ii) $\Box(\neg x \vee r1) \vee \Box(x \vee r2)$

This is due to the $\Box\Box$ resolution rule. We claim that using our method, this problem generates a number of clauses which is double exponentially bounded by the input. Other resolution methods have a rule for combining $\Box$ operators and work in a similar way [11,15]. The work that the other rules in resolution systems do is comparable to those that use transformation to disjunctive normal form. However, even though the complexity of our system is less optimal than [19], we perform better in general cases of the following example:

(i) $\neg x \vee \Box(\neg x \vee q1) \vee \Box(\neg x \vee q2)$ (ii) $\Box(x \vee r1) \vee \Box(x \vee r2)$

Because x, the symbol that we want to eliminate, appears at level zero in the first clause but only under a modal operator in the second clause, no resolution rule is applicable. Hence, $\neg x$ in the first clause can be purified, and by extending the system with standard simplification rules, the clause is replaced with $\top$.

Different to [11], we use additional propositional symbols to flatten our input, and different to [14], we do not use unification for first-order variables, which was used there because the problem is slightly different: they look at global (TBox) uniform interpolation with local (ABox) formulae. More broadly, completeness proofs for resolution-based uniform interpolation systems are traditionally shown via proof-theoretic arguments. Our proofs show that a model-theoretic argument, using Σ-bisimulation, can be made for proving uniform interpolation completeness of resolution-based systems. We notice that the proofs for D_n and T_n required only very modest extensions.

7 Conclusion

We presented a resolution-based method to compute uniform interpolants for the multi-modal logic K_n. We proved that our method terminates, and is sound and uniform interpolation complete. The space complexity was proven to be at most double exponential in the length of the input. We showed that the method can be used for computing uniform interpolants in D_n, and can be extended to compute uniform interpolants in T_n.

For future work, we would like to study logics which are known to have the uniform interpolation property, and show that the presented system can be extended to solve the uniform interpolation problem for more modal logics. Furthermore, it would be interesting to implement the UI_{K_n} system and perform an empirical comparison between this system and a system that transforms the input into disjunctive normal form (e.g., [19,3]).

Appendix

A Proofs

Lemma A.1 *The number of W-symbols introduced by the UI_{K_n} is bounded by 2^n-1 for n being the length of the input.*

Proof. Let S_w be the set of base W-symbols. S_w is finite because the number of modal operators in the input formula is finite, and the role of the world introduction rule is to replace each subformula (not containing a W-symbol) appearing under a modal operator with a W-symbol.

Let C_w be the set of combinatory W-symbols introduced by the RES$\Box\bigcirc$ rule. We show inductively that for each W-symbol W_i in C_w, the set $Corr(W_i)$ corresponds to a unique set of W-symbols from S_w.

Since the function $Corr$ has a finite range, which is the powerset of S_w, we can prove that the domain is finite by showing that $Corr$ is injective. Let W_n and W_m be two W-symbols. We will show inductively that if $Corr(W_n) = Corr(W_m)$, then $W_n = W_m$.

Base case: W_n is in S_w. The assumption is that $Corr(W_n) = Corr(W_m)$. Since W_n was produced as a result of an application of the world introduction rule, we have that $Corr(W_n) = \{W_n\}$. By assumption, this means that $Corr(W_m) = \{W_n\}$. Since $Corr(W_m)$ corresponds to a singleton set, it must be introduced via the world introduction rule as well, and hence, $W_n = W_m$.

Step case: W_n is in C_w. The cardinality of $Corr(W_n)$ must be greater than 1; this is because when W_n is introduced after a $\Box\bigcirc$ rule application, $Corr(W_n)$ is defined as the combination of two sets that share no elements. Assuming that $Corr(W_n) = Corr(W_m)$ entails that W_m must be in C_w as two equal sets have the same cardinality. Since W_n and W_m are in C_w, they must have been introduced after applying the $\Box\bigcirc$ rule. Assume w.l.o.g. that W_m was introduced first. Let $Corr(W_n) = Corr(W_i) \cup Corr(W_j)$ for a W_i and W_j that uniquely correspond to subsets of S_w symbols. By condition (iii) of the $\Box\bigcirc$ rule, there cannot be a W_k such that $Corr(W_k) = Corr(W_i) \cup Corr(W_j)$, therefore W_n must be equal to W_m. Since S_w is a finite set, and by condition (ii) of the $\Box\bigcirc$ resolution rule, the number of possible W-symbols is bounded by the number of unique combinations of symbols from S_w. The upper bound is equal to $2^{|S_w|}$-1, where $|S_w|$ is the size of S_w. $\Box$

Lemma A.2 *The space complexity of the UI_{K_n} calculus is double exponentially bounded in the length of the input.*

Proof. The root source of the complexity is that the system generates an exponential number of combinatory symbols. The complexity argument is built upon three claims. Using a linear number of propositional symbols, a linear number of base W-symbols, and an exponential number of combinatory W-symbols:

- C1 the size of each clause, before eliminating any W-symbols, has an exponential upper bound in the size of the input,
- C2 the number of clauses we could generate has a double exponential upper bound in the size of the input, and
- C3 the size of each clause, after eliminating W-symbols, has a double exponential upper bound in the size of the input.

For C1, before eliminating any W-symbols, a clause has propositional variables, a disjunction ψ of subformulas of ϕ, and a maximally exponential number of W-symbols appearing under a box or a diamond operator, or both. The number of propositional variables is linear in the size of the input (since they must appear in the input), the size of ψ is linear too for the same reason, and since the number of W-symbols is in $O(2^n)$ (n is the size of the input), then the size of a clause before eliminating any W-symbol is in $O(2^n)$.

For C2, since for each clause we have an exponential number of (propositional or W-symbol) variables that can either appear positively or negatively, the number of clauses that could be generated before eliminating any W-symbols is at most double exponential in the size of the input. (Note that we assume that the system incorporates simplification rules such as subsumption elimination.)

For C3, since the number of clauses is at most double exponential, and each clause has a size that is of at most an exponential size, eliminating W-symbols will result in a formula that is maximally bounded by $O(2^n * 2^{2^n})$. □

Theorem A.3 (Soundness) *Given a formula ϕ and a signature Σ, the UI_{K_n} system computes a formula ϕ' such that for any formula ψ over Σ, we have that if $\models \phi' \to \psi$ then $\models \phi \to \psi$.*

Proof. Consider ϕ, ϕ' as in Theorem 4.8. We use $\mathcal{M}|_\Sigma$ for the reduction of $\mathcal{M}$ to Σ, i.e. the model obtained from $\mathcal{M}$ by ignoring all symbols outside Σ. We say that two models $\mathcal{M}', \mathcal{M}$ are Σ-*inseparable* if $\mathcal{M}'|_\Sigma = \mathcal{M}|_\Sigma$. To show the theorem, we show the following claims:

C4 For each model $\mathcal{M}$ and world w_0 such that $\mathcal{M}, w_0 \models \phi$, there exists a model $\mathcal{M}'$ such that $\mathcal{M}$ and $\mathcal{M}'$ are Σ-inseparable, $\mathcal{M}' \models W_0 \Rightarrow \phi$, and $V(W_0) = \{w_0\}$.

C5 Each rule in our calculus is Σ-preserving, that is, if a rule is applied to a set of clauses N to produce N' and $\mathcal{M}$ is a model that globally satisfies N, then there exists a model $\mathcal{M}'$ that is Σ-inseparable from $\mathcal{M}$ that satisfies N'.

C6 Let $N = \{W_0 \Rightarrow \psi_i \mid 1 \leq i \leq n\}$ be the set obtained at the end of our derivation. For each model $\mathcal{M}$ that globally satisfies N and for each $w_0 \in V(W_0)$, we have $\mathcal{M}, w_0 \models \phi'$ with $\phi' = \bigwedge_{i \leq |N|} \psi_i$.

For C4, take a model $\mathcal{M}$ and a world w_0 such that $\mathcal{M}, w_0 \models \phi$. Extending $\mathcal{M}$ by setting $V(W_0) = \{w_0\}$ results in a Σ-inseparable model in which $W_0 \Rightarrow \phi$ is globally satisfied.

For C5, we consider two of our rules: the world introduction rule and the $\Box\bigcirc$ resolution rule. Proofs for the remaining rules are standard.

World Introduction: Let N' be the set obtained by replacing $W_i \Rightarrow \phi_1 \vee \bigcirc_a \phi_2$ in N with $W_i \Rightarrow \phi_1 \vee \bigcirc_a W_j$ and $W_j \Rightarrow \phi_2$.

Given $\mathcal{M}$, a model for N, our aim is to expand the model into $\mathcal{M}'$ for N' by giving W_j an appropriate assignment. Consider the following case distinction.

Case 1: $\bigcirc$ denotes a $\Diamond$ operator. We consider each $w \in V(W_i)$, if $\mathcal{M}, w \models \Diamond_a \phi_2$ then, there exists a world u, connected to w via an a-successor, such that $\mathcal{M}, u \models \phi_2$. We include u in $V'(W_j)$, i.e. ,

$$V'(W_j) = \{u \mid \mathcal{M}, u \models \phi_2 \text{ and } \exists\, w \text{ s.t. } w \in V(W_i) \text{ and } wR_a u \text{ is true in } \mathcal{M}\}.$$

By giving W_j the above interpretation, $W_i \Rightarrow \phi_1 \vee \Diamond_a W_j$ and $W_j \Rightarrow \phi_2$ become globally satisfiable in $\mathcal{M}'$.

Case 2: $\bigcirc$ denotes a $\Box$ operator. Similar to the first case, we consider each $w \in V(W_i)$. If $\mathcal{M}, w \models \Box_a \phi_2$, then for every world u connected to w via an a-successor, W_j must be made true in u. We extend the definition of $V'(W_j)$ to include u.

By giving W_j the above interpretation, $W_i \Rightarrow \phi_1 \vee \Box_a W_j$ and $W_j \Rightarrow \phi_2$ become globally satisfiable in $\mathcal{M}'$.

Now, we have completed the construction of the model $\mathcal{M}'$. We show that $\mathcal{M}'$ is a model for N'. The model $\mathcal{M}'$ satisfies all formulas in N' that do not include W_j. Indeed, this is because the interpretation of all symbols apart from W_j has not been changed. From Cases 1 and 2, it globally satisfies all clauses that contain W_j. Therefore, $\mathcal{M}'$ is a model for N'. The two models are Σ-inseparable because they are defined over the same frame, and the valuation function was not changed for propositional symbols in Σ.

$\Box\bigcirc$ **resolution:** Let N' be the set obtained by applying the $\Box\bigcirc$ resolution rule to $W_i \Rightarrow \gamma_1 \vee \Box_a W_n$ and $W_i \Rightarrow \gamma_2 \vee \bigcirc_a W_m$ attaining $W_i \Rightarrow \gamma_1 \vee \gamma_2 \vee \bigcirc_a W_j$ and $W_j \Rightarrow W_n$ and $W_j \Rightarrow W_m$.

Given $\mathcal{M}$, a model for N, our aim is to expand the model into $\mathcal{M}'$ for N' by giving W_j an appropriate assignment. We consider the two cases:

Case 1: $\bigcirc$ denotes a $\Diamond$ operator. We consider each $w \in V(W_i)$, if $\mathcal{M}, w \models \Box_a W_n$ and $\mathcal{M}, w \models \Diamond_a W_m$ then, there exists a world u, connected to w via an a-successor, such that $\mathcal{M}, u \models W_m$, and $\mathcal{M}, u \models W_n$. We include u in $V'(W_j)$.

Case 2: $\bigcirc$ denotes a $\Box$ operator. We consider each $w \in V(W_i)$. If $\mathcal{M}, w \models \Box_a W_n$ and $\mathcal{M}, w \models \Box_a W_m$, then for every world u connected to w via an a-successor, W_j must be made true in u. We include u in $V'(W_j)$.

We show that $\mathcal{M}'$ is a model for N'. Since the only change was to the interpretation of W_j, $\mathcal{M}'$ satisfies all clauses in N' without W_j. From Cases 1 and 2, $\mathcal{M}'$ satisfies clauses that contain W_j. The two models are Σ-inseparable because they are defined over the same frame, and the valuation function was not changed for propositional symbols in Σ.

For C6, the argument is trivial. Consider N as in the claim. Let $\mathcal{M}$ be a model that globally satisfies N, and w_0 be a world in $V(W_0)$. By the definition of global satisfiability, all clauses in N are satisfiable at w_0. By the semantics of implication, we have that $\mathcal{M}, w_0 \models \psi_i$ for $1 \leq i \leq |N|$. By the semantics of conjunction, $\mathcal{M}, w_0 \models \bigwedge_{i \leq |N|} \psi_i$.

We showed via C4, C5, and C6 that for any model of the input formula, there exists a Σ-inseparable model for the output formula. And, thus, for any formula ψ over Σ, if $\phi \rightarrow \psi$ is true in all models, then $\phi' \rightarrow \psi$ is true in all models. $\Box$

References

[1] Baral, C. and Y. Zhang, *Knowledge updates: Semantics and complexity issues*, Artificial Intelligence **164** (2005), pp. 209–243.

[2] Bechhofer, S., F. van Harmelen, J. Hendler, I. Horrocks, D. McGuinness, P. Patel-Schneijder and L. A. Stein, *OWL Web Ontology Language Reference*, Recommendation,

World Wide Web Consortium (W3C) (2004), see `http://www.w3.org/TR/owl-ref/`.
[3] Bílková, M., *Uniform interpolation and propositional quantifiers in modal logics*, Studia Logica (2007), pp. 1–31.
[4] Blackburn, P., M. d. Rijke and Y. Venema, "Modal Logic," Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2001.
[5] Fang, L., Y. Liu and H. Van Ditmarsch, *Forgetting in Multi-agent Modal Logics*, in: *Proc. IJCAI 2016* (2016), pp. 1066–1073.
[6] Fang, L., Y. Liu and H. Van Ditmarsch, *Forgetting in multi-agent modal logics*, Artificial Intelligence **266** (2019), pp. 51–80.
[7] Gabbay, D. M. and H. J. Ohlbach, *Quantifier elimination in second-order predicate logic*, in: *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, KR'92 (1992), pp. 425–435.
[8] Gabbay, D. M., R. A. Schmidt and A. Szałas, "Second-Order Quantifier Elimination: Foundations, Computational Aspects and Applications," College Publications, 2008.
[9] Ghilardi, S., *An algebraic theory of normal forms*, Annals of Pure and Applied Logic **71** (1995), pp. 189–245.
[10] Ghilardi, S. and M. Zawadowski, *Undefinability of propositional quantifiers in the modal system S4*, Studia Logica **55** (1995), pp. 259–271.
[11] Herzig, A. and J. Mengin, *Uniform interpolation by resolution in modal logic*, in: *European Workshop on Logics in Artificial Intelligence*, Springer, 2008, pp. 219–231.
[12] Horrocks, I., U. Hustadt, U. Sattler and R. Schmidt, *Computational modal logic*, in: *Handbook of Modal Logic*, Elsevier, 2007 pp. 181–245.
[13] Konev, B., D. Walther and F. Wolter, *Forgetting and uniform interpolation in large-scale description logic terminologies.*, in: *IJCAI* (2009), pp. 830–835.
[14] Koopmann, P., "Practical uniform interpolation for expressive description logics," Ph.D. thesis, The University of Manchester (2015).
[15] Koopmann, P. and R. A. Schmidt, *Uniform interpolation and forgetting for $\mathcal{ALC}$ ontologies with ABoxes*, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015, pp. 175–181.
[16] Ludwig, M. and B. Konev, *Practical uniform interpolation and forgetting for $\mathcal{ALC}$ TBoxes with applications to logical difference*, in: *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014*, 2014, pp. 318–327.
[17] Lutz, C. and F. Wolter, *Foundations for uniform interpolation and forgetting in expressive description logics*, in: *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011, pp. 989–995.
[18] Moss, L. S., *Finite models constructed from canonical formulas*, Journal of Philosophical Logic **36** (2007), pp. 605–640.
[19] Ten Cate, B., W. Conradie, M. Marx, Y. Venema et al., *Definitorially complete description logics.*, KR (2006), pp. 79–89.
[20] Toluhi, D., R. Schmidt and B. Parsia, *Concept description and definition extraction for the anemone system*, in: *International Workshop on Engineering Multi-Agent Systems*, Springer, 2021, pp. 352–372.
[21] Visser, A., *Bisimulations, model descriptions and propositional quantifiers*, Logic Group Preprint Series **161** (1996).
[22] Wolter, F., *Fusions of modal logics revisited*, in: *Advances in Modal Logic* (1998), pp. 361–379.