Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)

Forgetting in Multi-Agent Modal Logics

Liangda Fang[1,3] **Yongmei Liu**[1] **Hans van Ditmarsch**[2]
[1]Dept. of Computer Science,
Sun Yat-sen University, Guangzhou 510006, China
[2]LORIA, CNRS – Université de Lorraine, France
[3]Dept. of Computer Science, Jinan University, Guangzhou 510632, China
fangld@jnu.edu.cn, ymliu@mail.sysu.edu.cn, hans.van-ditmarsch@loria.fr

Abstract

In the past decades, forgetting has been investigated for many logics and has found many applications in knowledge representation and reasoning. However, forgetting in multi-agent modal logics has largely been unexplored. In this paper, we study forgetting in multi-agent modal logics. We adopt the semantic definition of existential bisimulation quantifiers as that of forgetting. We propose a syntactical way of performing forgetting based on the canonical formulas of modal logics introduced by Moss. We show that the result of forgetting a propositional atom from a satisfiable canonical formula can be computed by simply substituting the literals of the atom with $\top$. Thus we show that $\mathsf{K_n}$, $\mathsf{D_n}$, $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$ and $\mathsf{S5_n}$ are closed under forgetting, and hence have uniform interpolation.

1 Introduction

In the past decades, forgetting has been investigated for many logics and has found many applications in knowledge representation and reasoning (KR). Intuitively, forgetting some symbols from a knowledge base (KB) should result in a weaker KB which entails the same set of sentences that do not mention those symbols.

The seminal paper by Lin and Reiter [1994] studied forgetting in propositional and first-order logics. Over the years, forgetting in propositional logic has been used in abductive reasoning [Lin, 2001], reasoning under inconsistency [Lang and Marquis, 2010], reasoning about knowledge [Su *et al.*, 2009], epistemic planning [Herzig *et al.*, 2003], etc; and forgetting in first-order logic has been applied to progression for basic action theories in the situation calculus [Lin and Reiter, 1997; Liu and Lakemeyer, 2009]. In recent years, forgetting has been generalized to various description logics and applied to ontology reuse [Wang *et al.*, 2009; 2010; Konev *et al.*, 2009; Lutz and Wolter, 2011]. There exist some resolution-based approaches for forgetting in expressive description logics [Ludwig and Konev, 2014; Koopmann and Schmidt, 2014; 2015]. Forgetting has also been studied for logic programs and used in conflict solving [Zhang and Foo, 2006; Eiter and Wang, 2008; Wang *et al.*, 2014].

Forgetting for modal logics has also been investigated and applied to reasoning about knowledge and belief. Baral and Zhang [2005] studied knowledge update, a special form of update with the effect that the agent becomes ignorant of a propositional formula. Van Ditmarsch *et al.* [2009] presented a dynamic epistemic logic where the dynamic operator is the action of forgetting a propositional atom. Zhang and Zhou [2009] studied forgetting in propositional $\mathsf{S5}$ logic and its applications in knowledge updates and knowledge games. Liu and Wen [2011] explored forgetting in first-order $\mathsf{S5}$ logic and applied it to progression of knowledge in the situation calculus, which concerns updating the current representation of the world state and agents' epistemic state to reflect the change caused by an ontic or epistemic action. Later, Fang *et al.* [2015] extended the above results to progression of both knowledge and belief for nondeterministic actions in the single-agent propositional case.

Zhang and Zhou's semantic definition of forgetting coincides with that of existential bisimulation quantifiers. Bisimulation quantifiers were introduced by Ghilardi and Zawadowski [1995] to study uniform interpolation, which is the dual notion of forgetting, in modal logics. A logic has uniform interpolation if for any formula ϕ and any set S of symbols occurring in ϕ, there is a formula ψ using only symbols in S such that ϕ and ψ entail the same set of sentences formulated only in S. The semantics of existential bisimulation quantifiers is as follows: a model M satisfies a formula $\exists p\phi$ where p is an atom iff there is a model M' satisfying ϕ such that M and M' are bisimilar with exception on p. It turns out that any logic that is invariant under bisimulation and closed under bisimulation quantification, that is, closed under elimination of the quantification, has uniform interpolation. It is well-known that K, T and $\mathsf{S5}$ have uniform interpolation [Ghilardi, 1995; Bílková, 2007; Ghilardi and Zawadowski, 2000]. However, neither $\mathsf{K4}$ or $\mathsf{S4}$ has uniform interpolation [Bílková, 2007; Ghilardi and Zawadowski, 1995].

In the multi-agent case, D'Agostino and Lenzi [2006] gave a constructive proof that μ-calculus, an extension of $\mathsf{K_n}$ with a fixed-point operator, is closed under bisimulation quantification and hence has uniform interpolation. Studer [2009] showed that KC, which is $\mathsf{K_n}$ with common knowledge, does not have uniform interpolation, and this also holds for $\mathsf{K4C}$. Pattinson [2013] showed that all modal logics axiomatizable by formulas whose modal depth uniformly equals one, including $\mathsf{K_n}$ and $\mathsf{D_n}$, have uniform interpolation. Reasoning about knowledge and belief in the multi-agent case is certainly important from the KR point of view, since higher-order knowledge and belief (*i.e.*, knowledge and belief about other agents' knowledge and belief) play an important role in coordination and interaction. However, to the best of our knowledge, other than the above mentioned works, forgetting in multi-agent modal logics has largely been unexplored, and it remains an open problem whether $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$ or $\mathsf{S5_n}$ have uniform interpolation.

It is well-known that in propositional logic, the result of forgetting atom p from a satisfiable term (*i.e.*, a conjunction of literals) δ can be obtained from δ by substituting any occurrence of literal p or $\neg p$ with $\top$. As every propositional formula can be equivalently transformed into a disjunction of satisfiable terms, and forgetting is distributive over disjunction, propositional logic is closed under forgetting. Interestingly, the idea of literal elimination was used by D'Agostino and Lenzi to prove that μ-calculus is closed under bisimulation quantification.

In this paper, we study forgetting in multi-agent modal logics. We adopt the semantic definition of existential bisimulation quantifiers as that of forgetting. We resort to the normal forms of propositional modal logics, called canonical formulas by Moss [2007], a nice property of which is that an arbitrary modal formula is equivalent to a disjunction of a finite set of satisfiable canonical formulas. Due to the distributive law of forgetting over disjunction, forgetting from arbitrary modal formulas can be reduced to forgetting from satisfiable canonical formulas, which we show can be computed via literal elimination. Hence we show that except $\mathsf{K4_n}$ and $\mathsf{S4_n}$, which do not have uniform interpolation, the other main multi-agent modal systems are closed under forgetting.

2 Preliminaries

In this section, we introduce the background material, *i.e.*, the syntax and semantics of multi-agent modal logics, and the canonical formulas for modal logics as defined by Moss.

2.1 Multi-agent modal logics

We fix a set $\mathcal{A}$ of n agents and a countable set of atoms $\mathcal{P}$.

Definition 1 The language $\mathcal{L}_n^{\mathbf{K}}$ is generated by the BNF:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{K}_i\phi,$$

where $p \in \mathcal{P}$ and $i \in \mathcal{A}$. We use $\mathcal{L}^{pl}$ for the propositional language, *i.e.*, the language without the $\mathbf{K}_i$ modality.

We let $\top$ and $\bot$ represent $true$ and $false$ respectively. We use p and q to range over atoms, ϕ and ψ to range over formulas, and Φ and Ψ to range over sets of formulas. We let $\mathcal{P}(\phi)$ denote the set of atoms which appear in ϕ.

The *(modal) depth* of a formula ϕ in $\mathcal{L}_n^{\mathbf{K}}$, written $dep(\phi)$, is the depth of nesting of modal operators in ϕ. We let $\hat{\mathbf{K}}_i\phi$ stand for $\neg\mathbf{K}_i\neg\phi$. We let $\bigvee\Phi$ (resp. $\bigwedge\Phi$) denote the disjunction (resp. conjunction) of members of Φ; and we use $\bigwedge\hat{\mathbf{K}}_i\Phi$ to represent the conjunction of $\hat{\mathbf{K}}_i\phi$ where $\phi \in \Phi$.

Definition 2 A frame is a pair (S, R), where

L	K	D	T	K4	S4	K45	KD45	S5
Serial		✓	✓		✓		✓	✓
Reflexive			✓		✓			✓
Transitive				✓	✓	✓	✓	✓
Euclidean						✓	✓	✓

Table 1: The main modal systems

- S is a non-empty set of possible worlds;
- For each agent $i \in \mathcal{A}$, R_i is a binary relation on S, called the *accessibility relation* for i.

For $s \in S$, we write $R_i(s)$ for $\{t \in S \mid sR_it\}$, and call it the i-children of s.

Different *modal systems* result from different sets of conditions on the accessibility relations. We say R_i is serial if for any $s \in S$, there is $s' \in S$ s.t. sR_is'; we say R_i is Euclidean if whenever sR_is_1 and sR_is_2, we have $s_1R_is_2$.

We list the main modal systems in Table 1. $\mathsf{KD45}$ and $\mathsf{S5}$ are well-accepted as the logics for knowledge and belief, respectively. For each symbol L listed here, we use L for the single agent case, $\mathsf{L_n}$ for the case where there are n agents.

Definition 3 A Kripke model is a triple $M = \langle S, R, V\rangle$, where (S, R) is a frame, and V is a valuation map, which maps each $s \in S$ to a subset of $\mathcal{P}$. A pointed Kripke model is a pair (M, s), where M is a Kripke model and s is a world of M, called the *actual world*.

For simplicity, we often omit the word "pointed".

Definition 4 Let (M, s) be a Kripke model where $M = \langle S, R, V\rangle$. We interpret formulas in $\mathcal{L}_n^{\mathbf{K}}$ by induction:

- $M, s \models p$ iff $p \in V(s)$;
- $M, s \models \neg\phi$ iff $M, s \not\models \phi$;
- $M, s \models \phi \wedge \psi$ iff $M, s \models \phi$ and $M, s \models \psi$;
- $M, s \models \mathbf{K}_i\phi$ iff for all $t \in R_i(s)$, $M, t \models \phi$.

We use L to range over modal systems. Consider the context of a modal system L. We say ϕ is satisfiable, if there exists a model of ϕ. We say ϕ entails ψ, written $\phi \models \psi$, if for any model (M, s), $M, s \models \phi$ implies $M, s \models \psi$. We say ϕ and ψ are equivalent, written $\phi \equiv \psi$, if $\phi \models \psi$ and $\psi \models \phi$.

D'Agostino and Lenzi [2005] presented another modal language which is based on the cover modality ∇_i:

Definition 5 Let $i \in \mathcal{A}$, and Φ a finite set of formulas in $\mathcal{L}_n^{\mathbf{K}}$. We use $\nabla_i\Phi$ to denote the formula $\mathbf{K}_i(\bigvee\Phi) \wedge (\bigwedge\hat{\mathbf{K}}_i\Phi)$.

Note that $\nabla_i\emptyset \equiv \mathbf{K}_i\bot$.

Now, we can get the semantics of the cover modality.

Lemma 1 *Let (M, s) be a Kripke model where $M = \langle S, R, V\rangle$ and $s \in S$. Then $M, s \models \nabla_i\Phi$ iff the following conditions hold:*

- *for all $t \in R_i(s)$, there is $\phi \in \Phi$ s.t. $M, t \models \phi$;*
- *for all $\phi \in \Phi$, there is $t \in R_i(s)$ s.t. $M, t \models \phi$.*

Proposition 1 *[Janin and Walukiewicz, 1995] Every formula in $\mathcal{L}_n^{\mathbf{K}}$ is equivalent to a formula using only the ∇_i modalities.*

2.2 Canonical formulas

We will present our definitions and proofs via *canonical formulas* of modal logics, as defined by Moss [2007], which capture Kripke models up to a given depth.

Definition 6 (Canonical formulas) Let $P \subseteq \mathcal{P}$ be finite. We inductively define the set E_k^P as follows:

- $E_0^P = \{\bigwedge_{p \in \tilde{P}} p \wedge \bigwedge_{p \in P \setminus \tilde{P}} \neg p \mid \tilde{P} \subseteq P\}$;
- $E_{k+1}^P = \{\delta_0 \wedge \bigwedge_{i \in \mathcal{A}} \nabla_i \Phi_i \mid \delta_0 \in E_0^P \text{ and } \Phi_i \subseteq E_k^P\}$.

Let $\delta = \delta_0 \wedge \bigwedge_{i \in \mathcal{A}} \nabla_i \Phi_i \in E_{k+1}^P$. We denote δ_0 by $w(\delta)$, and call it the world of δ; we denote Φ_i by $R_i(\delta)$, and call it the i-children of δ.

Intuitively, when we restrict our attention to the atoms in P, a canonical formula of depth k provides a complete syntactic representation of a Kripke model up to depth k.

The following proposition says that every Kripke model (M, s) satisfies a unique canonical formula of a given depth k, which we call the depth k canonical formula of (M, s).

Proposition 2 *[Moss, 2007] Let (M, s) be a model and $k \in \mathbb{N}$. Let $P \subseteq \mathcal{P}$ be finite. Then, there exists a unique $\delta \in E_k^P$ s.t. $M, s \models \delta$.*

We can get the following corollary which means that every modal formula can be equivalently transformed into a disjunction of satisfiable canonical formulas whose depth is not less than that of the original formula.

Proposition 3 *Consider the context of a modal system* L. *Let $\phi \in \mathcal{L}_n^{\mathbf{K}}$ and $l \in \mathbb{N}$ s.t. $l \geq dep(\phi)$. Let $P = \mathcal{P}(\phi)$. Then there exists a unique set $\Phi \subseteq E_l^P$ s.t. $\phi \equiv \bigvee \Phi$, and for every $\delta \in \Phi$, δ is satisfiable in* L.

Now, we introduce the projection operations on canonical formulas. A canonical formula can be graphically represented as a tree. The operation $\delta^{\downarrow}$ prunes the leaves of this tree. In $\delta^{\downarrow l}$, we perform pruning l times sequentially. We call $\delta^{\downarrow}$ the 1st-cut of δ, and $\delta^{\downarrow l}$ the lth-cut of δ.

Definition 7 Let $P \subseteq \mathcal{P}$ be finite. Let $k \in \mathbb{N}$ and $\delta \in E_k^P$. Then, $\delta^{\downarrow}$ is inductively defined as follows:

$$\delta^{\downarrow} = \begin{cases} \delta, & \text{if } k = 0; \\ w(\delta), & \text{if } k = 1; \\ w(\delta) \wedge \bigwedge_{i \in \mathcal{A}} \nabla_i (R_i(\delta))^{\downarrow}, & \text{otherwise.} \end{cases}$$

Let Φ be a set of canonical formulas. We use $\Phi^{\downarrow}$ to denote the set $\{\phi^{\downarrow} \mid \phi \in \Phi\}$.

Definition 8 Let $P \subseteq \mathcal{P}$ be finite. Let $k, l \in \mathbb{N}$ s.t. $k \geq l$. Let $\delta \in E_k^P$. Then, $\delta^{\downarrow l}$ is inductively defined as follows:

$$\delta^{\downarrow l} = \begin{cases} \delta, & \text{if } l = 0; \\ \delta^{\downarrow}, & \text{if } l = 1; \\ (\delta^{\downarrow l-1})^{\downarrow}, & \text{otherwise.} \end{cases}$$

Similarly to $\Phi^{\downarrow}$, $\Phi^{\downarrow l}$ is the set $\{\phi^{\downarrow l} \mid \phi \in \Phi\}$. Obviously, $\delta \models \delta^{\downarrow l}$ for all $l \leq dep(\delta)$.

It is easy to prove the following, which intuitively says that the i-children of the lth-cut of a canonical formula is equal to the lth-cut of its i-children, *i.e.*, the two operations are commutative.

Proposition 4 *Let δ be a canonical formula. Let $l \in \mathbb{N}$ s.t. $l < dep(\delta)$. Then for all $i \in \mathcal{A}$, we have $R_i(\delta^{\downarrow l}) = R_i(\delta)^{\downarrow l}$.*

Example 1 Let $\delta = p \wedge \nabla_i\{\neg p \wedge \nabla_i\{p, \neg p\}\}$. Obviously, the depth of δ is 2. By Definitions 7 and 8, $\delta^{\downarrow 1} = \delta^{\downarrow} = p \wedge \nabla_i(\{\neg p \wedge \nabla_i\{p, \neg p\}\})^{\downarrow} = p \wedge \nabla_i\{\neg p\}$. Similarly, we get $R_i(\delta^{\downarrow 1}) = \{\neg p\}$, $R_i(\delta) = \{\neg p \wedge \nabla_i\{p, \neg p\}\}$, and $R_i(\delta)^{\downarrow 1} = \{\neg p\}$. Hence $R_i(\delta^{\downarrow 1}) = R_i(\delta)^{\downarrow 1} = \{\neg p\}$.

3 Forgetting

In this section, we define forgetting in multi-agent modal logics. We first review forgetting in propositional logic.

We use a subset of the atoms to denote a valuation.

Definition 9 Let $\phi \in \mathcal{L}^{pl}$ and p an atom. A formula $\psi \in \mathcal{L}^{pl}$ s.t. $\mathcal{P}(\psi) \subseteq \mathcal{P}(\phi) \setminus \{p\}$ is a result of forgetting p in ϕ, written pforget$(\phi, p) \equiv \psi$, if for any valuation P, $P \models \phi$ (meaning P satisfies ϕ) iff there is a valuation P' s.t. $P' \models \psi$ and $P \sim_p P'$, *i.e.*, $P \setminus \{p\} = P' \setminus \{p\}$.

The definition of forgetting in multi-agent modal logics is analogous to that in propositional logic. For this purpose, we use the well-known concept of bisimulation.

Definition 10 (p-bisimulation) Let (M, s) and (M', s') be two Kripke models where $M = \langle S, R, V \rangle$ and $M' = \langle S', R', V' \rangle$. A p-bisimulation between (M, s) and (M', s') is a relation $\rho \subseteq S \times S'$ s.t. $s \rho s'$, and whenever $t \rho t'$, we get:

atoms $V(t) \sim_p V'(t')$;

forth For all i, if tR_iu, then there is u' s.t. $t'R'_iu'$ and $u\rho u'$;

back For all i, if $t'R'_iu'$, then there is u s.t. tR_iu and $u \rho u'$.

We say that (M, s) and (M', s') are p-bisimilar, written $(M, s) \underline{\leftrightarrow}_p (M', s')$, if there is a p-bisimulation between them.

Clearly, $\underline{\leftrightarrow}_p$ is an equivalence relation. A nice property of p-bisimilar Kripke models is that they agree on all modal formulas wherein p does not appear.

Proposition 5 *Let $(M, s) \underline{\leftrightarrow}_p (M', s')$. Let $\phi \in \mathcal{L}_n^{\mathbf{K}}$ where p does not appear in ϕ. Then $M, s \models \phi$ iff $M', s' \models \phi$.*

Definition 11 (Forgetting) Consider the context of a modal system L. Let $\phi \in \mathcal{L}_n^{\mathbf{K}}$ and p an atom. A formula ψ s.t. $\mathcal{P}(\psi) \subseteq \mathcal{P}(\phi) \setminus \{p\}$ is a result of forgetting p in ϕ, written kforget$(\phi, p) \equiv \psi$, if for any model (M', s'), $M', s' \models \psi$ iff there is a model (M, s) of ϕ s.t. $(M, s) \underline{\leftrightarrow}_p (M', s')$.

So the semantics of forgetting is that of existential bisimulation quantifiers as presented in the introduction. Clearly, if both ψ_1 and ψ_2 are results of forgetting p in ϕ, then they are logically equivalent. In this paper, when we prove kforget$(\phi, p) \equiv \psi$, the proof of the if direction of Definition 11 is straightforward by making use of $(M, s) \underline{\leftrightarrow}_p (M', s')$ and the induction hypothesis if applicable. To save space, we only present the proof of the only-if direction.

We now analyze basic properties of forgetting.

Proposition 6 *Consider the context of a modal system* L. *If kforget$(\phi, p) \equiv \psi$, then $\phi \models \psi$; and for any η where p does not occur, $\phi \models \eta$ iff $\psi \models \eta$.*

Proposition 7 *Consider the context of a modal system* L.

1. *If $\phi \in \mathcal{L}^{pl}$ and pforget$(\phi, p) \equiv \psi$, then kforget$(\phi, p) \equiv \psi$;*
2. *kforget$(\phi_1 \vee \phi_2, p) \equiv$ kforget$(\phi_1, p) \vee$ kforget(ϕ_2, p).*

Proof: We only prove the only-if direction of Part 1. Let $M', s' \models \psi$. We let the structure of M be a copy of that of M'; let V expand V' s.t. $V(s) \models \phi$. Then we have $M, s \models \phi$ and $(M, s) \underline{\leftrightarrow}_p (M', s')$. ■

We now relate forgetting to uniform interpolation.

Definition 12 We say that a modal system L is closed under forgetting if for any formula ϕ and any atom $p \in \mathcal{P}(\phi)$, there exists an L formula ψ s.t. kforget$(\phi, p) \equiv \psi$.

Definition 13 We say a modal system L has uniform interpolation if for every formula ϕ and every $P \subseteq \mathcal{P}(\phi)$, there is a formula ψ such that $\mathcal{P}(\psi) \subseteq P$ and such that for any formula η with $\mathcal{P}(\eta) \subseteq P$, we have $\phi \models \eta$ iff $\psi \models \eta$.

Zhang and Zhou [2009] proposed the forgetting postulates, which correspond to the definition of uniform interpolation.

By Proposition 6, we have

Proposition 8 *If a modal system L is closed under forgetting, then L has uniform interpolation.*

As mentioned in the introduction, neither $\mathsf{K4}$ nor $\mathsf{S4}$ has uniform interpolation. Because $\mathsf{K4}$ is indeed $\mathsf{K4}_1$, which is a special case of $\mathsf{K4_n}$, $\mathsf{K4_n}$ does not have uniform interpolation. Similarly, neither has $\mathsf{S4_n}$. Hence, we get

Corollary 1 *Neither $\mathsf{K4_n}$ nor $\mathsf{S4_n}$ is closed under forgetting.*

Finally, we propose a syntactical method of forgetting.

Definition 14 (Literal elimination) Let $\phi \in \mathcal{L}_n^{\mathbf{K}}$ and p an atom. We let ϕ^p denote the formula obtained from ϕ by substituting all occurrences of $\neg p$ with $\top$ and subsequently substituting all occurrences of p with $\top$.

Merely requiring to replace $\neg p$ and p by $\top$ would be ambiguous. Because, should we then replace $\neg p$ by $\top$ or by $\bot$? The last would happen if we were to replace the p in $\neg p$ by $\top$. That would be an undesirable outcome. The 'subsequently' in Definition 14 is not ambiguous and avoids that outcome.

Example 2 Let $\delta_1 = p \wedge q \wedge \nabla_i\{\neg p \wedge \neg q\}$. Then, $\delta_1^p = \top \wedge q \wedge \nabla_i\{\top \wedge \neg q\} \equiv q \wedge \nabla_i\{\neg q\}$.

Typically, we only apply this substitution to formulas wherein negations only bind atoms. Similarly to $\Phi^{\downarrow}$, Φ^p is the set $\{\phi^p \mid \phi \in \Phi\}$.

In the following sections, we will show that when δ is a canonical formula *satisfiable* in L, kforget$_{\mathsf{L}}(\delta, p) \equiv \delta^p$.

4 Forgetting in multi-agent modal logics

In this section, we show that forgetting from satisfiable canonical formulas can be computed via literal elimination in the following multi-agent modal logics: $\mathsf{K_n}$, $\mathsf{D_n}$, $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$ and $\mathsf{S5_n}$. As an easy corollary, we have that the above logics are closed under forgetting and they have uniform interpolation.

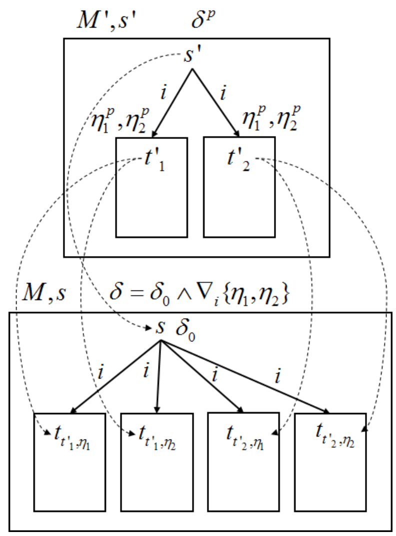

Figure 1: Illustration for the proof of Theorem 1

4.1 Forgetting in $\mathsf{K_n}$ and $\mathsf{D_n}$

In this subsection, we consider $\mathsf{K_n}$ and $\mathsf{D_n}$ cases. In fact, these results were first proved by D'Agostino and Lenzi [2006]. As mentioned in the introduction, they gave a constructive proof that μ-calculus is closed under bisimulation quantification. The core result they proved is the following:

Theorem *[D'Agostino and Lenzi, 2006] Let ϕ be a cover disjunctive formula for μ-calculus. Then $\exists p \phi \equiv \phi^p$.*

In the theorem, a cover disjunctive formula is a generalization of the notion of a disjunction of canonicial formulas that we use in this paper.

Here $\exists p$ is the bisimulation quantifier we defined in the introduction. Their proof is via an automata approach. Here we present an easy inductive proof, which serves as the basis for the proofs of all the forgetting results in this paper.

Theorem 1 (The basic theorem) *Let L be $\mathsf{K_n}$ or $\mathsf{D_n}$, and δ an L-satisfiable canonical formula. Then kforget$_{\mathsf{L}}(\delta, p) \equiv \delta^p$.*

Proof: We prove by induction on $dep(\delta)$. The base case, *i.e.*, $\delta \in E_0^P$, follows from Proposition 7. We prove the only-if direction of the induction case, *i.e.*, $\delta \in E_k^P$ where $k \geq 1$. Let $M', s' \models \delta^p$. We construct M and define a relation ρ between the worlds of M and M' as follows. Figure 1 illustrates the construction.

1. Create a new world s, let $s \rho s'$, and $V(s) \models w(\delta)$.
2. For all $i \in \mathcal{A}$, $t' \in R'_i(s')$, and $\eta \in R_i(\delta)$, if $M', t' \models \eta^p$, by the induction hypothesis, there exist $(M_{t',\eta}, t_{t',\eta})$ and $\rho_{t',\eta}$ s.t. $M_{t',\eta}, t_{t',\eta} \models \eta$ and $\rho_{t',\eta} : (M_{t',\eta}, t_{t',\eta}) \underline{\leftrightarrow}_p (M', t')$. Add a new copy of $M_{t',\eta}$ into M, let $s R_i t_{t',\eta}$, and expand ρ with $\rho_{t',\eta}$.

It is easy to show that $\rho : (M, s) \underline{\leftrightarrow}_p (M', s')$ and $M, s \models \delta$.

In the case of $\mathsf{D_n}$, it is obvious from the above construction that if M' satisfies seriality, so does M. ■

As a corollary of Propositions 3, 7 and Theorem 1, we get

Corollary 2 *$\mathsf{K_n}$ and $\mathsf{D_n}$ are closed under forgetting.*

Example 3 Let ϕ be $\hat{\mathbf{K}}_i p \wedge \hat{\mathbf{K}}_i \neg p$, meaning that agent i is ignorant about p. Intuitively, after forgetting p, the agent should still have consistent belief. We convert ϕ into a disjunction of two canonical formulas δ_1 and δ_2 where $\delta_1 = p \wedge \nabla_i\{p, \neg p\}$ and $\delta_2 = \neg p \wedge \nabla_i\{p, \neg p\}$. Then $\delta_1^p = \delta_2^p = \top \wedge \nabla_i\{\top\}$. The disjunction of them is equivalent to $\hat{\mathbf{K}}_i \top$.

4.2 Forgetting in $\mathsf{T_n}$

In this subsection, we analyze properties of $\mathsf{T_n}$ satisfiable canonical formulas, and show that forgetting via literal elimination applies to them.

We begin with a simple example which shows that Theorem 1 does not hold for unsatisfiable canonical formulas.

Example 4 Let $\delta = \neg p \wedge \nabla_i\{p\}$. Clearly, δ is a canonical formula, and it is equivalent to $\perp$ in $\mathsf{T_n}$. However, $\delta^p = \top \wedge \nabla_i\{\top\}$ which is equivalent to $\top$.

The reason that forgetting via literal elimination does not work on δ is that δ is unsatisfiable in $\mathsf{T_n}$. This motivates us to consider only satisfiable canonical formulas.

The following proposition says that any $\mathsf{T_n}$ satisfiable canonical formula δ has the reflexive property: for any agent i and any $1 \leq l \leq dep(\delta)$, the lth-cut of δ is an i-child of its $(l-1)$th-cut.

Proposition 9 *Let δ be a $\mathsf{T_n}$ satisfiable canonical formula where $dep(\delta) \geq 1$. Let $l \in \mathbb{N}$ s.t. $1 \leq l \leq dep(\delta)$. Then, for all $i \in \mathcal{A}$, we have $\delta^{\downarrow l} \in R_i(\delta^{\downarrow l-1})$.*

Proof: Here, we only prove the base case, *i.e.*, $\delta^{\downarrow} \in R_i(\delta)$. Let $k = dep(\delta)$ and $M, s \models \delta$. Obviously, $M, s \models \delta^{\downarrow}$. Since M is reflexive, $s \in R_i(s)$. By the semantics of the ∇_i modality, there exists $\delta_i \in R_i(\delta)$ s.t. $M, s \models \delta_i$. So both $\delta^{\downarrow}$ and δ_i are the depth $k-1$ canonical formula of (M, s) (cf. Proposition 2). Hence $\delta^{\downarrow} = \delta_i \in R_i(\delta)$. ∎

Example 4 Cont'd We have $\delta^{\downarrow} = \neg p$ and $R_i(\delta) = \{p\}$. Obviously, $\delta^{\downarrow} \notin R_i(\delta)$. So δ is not satisfiable in $\mathsf{T_n}$.

Example 5 Let $\delta = p \wedge q \wedge \nabla_i\{p \wedge q, p \wedge \neg q\}$. It is a $\mathsf{T_n}$ satisfiable canonical formula. Then, $\delta^{\downarrow} = p \wedge q$, and $R_i(\delta) = \{p \wedge q, p \wedge \neg q\}$. Obviously, $\delta^{\downarrow} \in R_i(\delta)$.

Theorem 2 (The $\mathsf{T_n}$ theorem) *Let δ be a $\mathsf{T_n}$ satisfiable canonical formula. Then $kforget_{\mathsf{T_n}}(\delta, p) \equiv \delta^p$.*

Proof: The proof is the same as that of the basic theorem except the following. In the induction case, we get $(M, s) \underline{\leftrightarrow}_p (M', s')$ and $M, s \models \delta$. Let $k = dep(\delta)$. Then for any $l \leq k$, $M, s \models \delta^{\downarrow l}$. We now let sR_is for each agent i. It is obvious that (M, s) is a $\mathsf{T_n}$ model. It is easy to see that $(M, s) \underline{\leftrightarrow}_p (M', s')$ still holds. It remains to show that $M, s \models \delta$ still holds. We prove by induction on $k - l$ that for any $l \leq k$, $M, s \models \delta^{\downarrow l}$ still holds. Base case. $M, s \models w(\delta)$, which is $\delta^{\downarrow k}$. Induction case. Suppose that $M, s \models \delta^{\downarrow l}$. To show that $M, s \models \delta^{\downarrow l-1}$, it suffices to show that for each $i \in \mathcal{A}$, there exists $\eta \in R_i(\delta^{\downarrow l-1})$ s.t. $M, s \models \eta$. By Proposition 9, $\delta^{\downarrow l} \in R_i(\delta^{\downarrow l-1})$. Hence $\delta^{\downarrow l}$ is the desired η. ∎

4.3 Multi-pointed Kripke models

In the next subsection, we show that forgetting via literal elimination applies to satisfiable canonical formulas of $\mathsf{K45_n}$, $\mathsf{KD45_n}$ and $\mathsf{S5_n}$. The proof for the basic theorem does not immediately carry forward to these cases, because the model constructed in the proof may not be transitive or Euclidean. Also, we cannot simply fix the problem by adding edges as we do in the proof of the $\mathsf{T_n}$ theorem. We will overcome the problem via multi-pointed Kripke models which offers flexibility in the construction of required models. In this subsection, we introduce the basic concepts regarding multi-pointed models.

Definition 15 A multi-pointed Kripke model is a pair (M, T) where M is a Kripke model, and T is a possibly empty set of worlds of M.

Throughout this paper, we use (M, s) to denote a single-pointed model, and (M, T) a multi-pointed model.

Given a single-pointed model (M, s) and $i \in \mathcal{A}$, we can naturally obtain a multi-pointed model (M, T) where $T = R_i(s)$, *i.e.*, the i-children of s.

Similarly to the semantics of the cover modality, we have:

Definition 16 Let Φ be a set of formulas. We say a multi-pointed model (M, T) is Φ-*complete* if the following hold:

- for every $t \in T$, there exists $\phi \in \Phi$ s.t. $M, t \models \phi$;
- for every $\phi \in \Phi$, there exists $t \in T$ s.t. $M, t \models \phi$.

Obviously, $M, s \models \nabla_i \Phi$ iff $(M, R_i(s))$ is Φ-complete.

Then, we can extend Proposition 2 to multi-pointed models: we have that a multi-pointed model (M, T) corresponds to a unique set of canonical formulas of a given depth k, which we call the depth k canonical formula set of (M, T).

Proposition 10 *Let (M, T) be a multi-pointed model, and $k \in \mathbb{N}$. Let $P \subseteq \mathcal{P}$ be finite. Then, there exists a unique set $\Phi \subseteq E_k^P$ s.t. (M, T) is Φ-complete.*

We now extend the concept of bisimulation to multi-pointed models.

Definition 17 Let (M, T) and (M', T') be two Kripke models. A p-bisimulation between (M, T) and (M', T') is a relation ρ between the worlds of M and M' s.t.

- for every $t \in T$, there exists $t' \in T'$ s.t. $t \rho t'$;
- for every $t' \in T'$, there exists $t \in T$ s.t. $t \rho t'$;
- whenever $u \rho u'$, the atoms, forth and back conditions in Definition 10 hold.

We end with a constraint on multi-pointed models, which is crucial for constructing transitive and Euclidean models.

Definition 18 Let (M, T) be a Kripke model and $i \in \mathcal{A}$. We say that (M, T) is i-*equivalent* if for all $t_1, t_2 \in T$, we have $t_1 R_i t_2$, and for every $s \in S$, if there exists $t \in T$ s.t. sR_it or tR_is, then $s \in T$.

4.4 Forgetting in $\mathsf{K45_n}$, $\mathsf{KD45_n}$ and $\mathsf{S5_n}$

In this subsection, we generalize the forgetting results for satisfiable canonical formulas to $\mathsf{K45_n}$, $\mathsf{KD45_n}$ and $\mathsf{S5_n}$.

The following proposition says that any $\mathsf{K45_n}$ satisfiable canonical formula δ has the identical-children property: for any agent i and any i-child δ_i of δ, the lth-cut of δ's i-children is equal to the i-children of the $(l-1)$th-cut of δ_i.

Proposition 11 *Let δ be a $\mathsf{K45_n}$ satisfiable canonical formula where $dep(\delta) \geq 2$. Let $l \in \mathbb{N}$ s.t. $1 \leq l < dep(\delta)$. Then, for all $i \in \mathcal{A}$ and $\delta_i \in R_i(\delta)$, $(R_i(\delta))^{\downarrow l} = R_i(\delta_i^{\downarrow l-1})$.*

Proof: Let $k = dep(\delta)$, $M, s \models \delta$, $i \in \mathcal{A}$ and $\delta_i \in R_i(\delta)$. Here, we only prove the base case, *i.e.*, $(R_i(\delta))^{\downarrow} = R_i(\delta_i)$. Then there exists $t \in R_i(s)$ s.t. $M, t \models \delta_i$. Since M is transitive and Euclidean, $R_i(s) = R_i(t)$. Obviously, $(M, R_i(s))$ is $(R_i(\delta))^{\downarrow}$-complete and $(M, R_i(t))$ is $R_i(\delta_i)$-complete. So $(M, R_i(s))$ is also $R_i(\delta_i)$-complete. So both $(R_i(\delta))^{\downarrow}$ and $R_i(\delta_i)$ are the depth $k-2$ canonical formula set of $(M, R_i(s))$ (cf. Proposition 10). Hence $(R_i(\delta))^{\downarrow} = R_i(\delta_i)$. ■

Now, we prove an important lemma, which is the multi-pointed extension of the only-if direction of the version of the basic theorem for $\mathsf{K45_n}$.

Lemma 2 (The $\mathsf{K45_n}$ lemma) *Let δ be a $\mathsf{K45_n}$ satisfiable canonical formula where $dep(\delta) \geq 1$. Let (M', s') be a $\mathsf{K45_n}$ model of δ^p. Then for all $i \in \mathcal{A}$, there exists a multi-pointed $\mathsf{K45_n}$ model (M, T) that is i-equivalent, $R_i(\delta)$-complete and $(M, T) \underline{\leftrightarrow}_p (M', R'_i(s'))$.*

Proof: We prove by induction on $dep(\delta)$.

Base case, *i.e.*, $R_i(\delta)$ is a set of propositional formulas. We construct (M, T) and define the p-bisimulation ρ between (M, T) and $(M', R'_i(s'))$ as follows. Figure 2 illustrates the construction. We initialize $S = \emptyset$ and $\rho = \emptyset$.

1. For all $t' \in R'_i(s')$ and $\delta_i \in R_i(\delta)$, if $M', t' \models \delta_i^p$, we create a world t s.t. $V(t) \models w(\delta_i)$. Then we add it into S and T, and let $t \rho t'$.
2. For all $t_1, t_2 \in T$, we let $t_1 R_i t_2$.
3. For $t \in T$ and $j \neq i$, we make a copy of $(M', R'_j(t'))$, denoted by $(M_{t,j}, T_{t,j})$ where t' is the original world of t. We connect t to all worlds of $T_{t,j}$ via j-edges, *i.e.*, let $t R_j u$ for $u \in T_{t,j}$. We let $u \rho u'$ when u is the copy of u'.

It is easy to verify (M, T) satisfies the requirements.

Induction step. The construction of (M, T) is similar to that of the base case except the following. In Step 3, since $M', t' \models \delta_i^p$, by the induction hypothesis, there exist $(M_{t,j}, T_{t,j})$ that is j-equivalent and $R_j(\delta_i)$-complete and $\rho_{t,j}$: $(M_{t,j}, T_{t,j}) \underline{\leftrightarrow}_p (M', R'_j(t'))$. We expand ρ with $\rho_{t,j}$. Obviously, (M, T) is transitive, Euclidean and i-equivalent, and $(M, T) \underline{\leftrightarrow}_p (M', R'_i(s'))$.

It remains to verify that (M, T) is $R_i(\delta)$-complete. Let $dep(\delta) = k$. We prove by induction on $k - l$ that for all $0 \leq l < k$, (M, T) is $R_i(\delta)^{\downarrow l}$-complete. Base case: $l = k - 1$, and $R_i(\delta)^{\downarrow l}$ is a set of propositional formulas. By Step 1, (M, T) is $R_i(\delta)^{\downarrow l}$-complete. Induction step. Suppose that (M, T) is $R_i(\delta)^{\downarrow l}$-complete. We prove that it is also $(R_i(\delta))^{\downarrow l-1}$-complete. Let $t \in T$. By the construction, there exists $\delta_i \in R_i(\delta)$ s.t. $M, t \models w(\delta_i)$. It

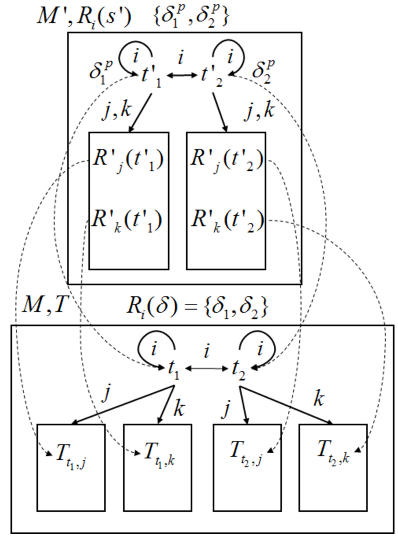

Figure 2: Illustration for the proof of Lemma 2

suffices to show that $M, t \models \delta_i^{\downarrow l-1}$. Since $w(\delta_i^{\downarrow l-1}) = w(\delta_i)$, $M, t \models w(\delta_i^{\downarrow l-1})$. By the identical-children property (Proposition 11), $(R_i(\delta))^{\downarrow l} = R_i(\delta_i^{\downarrow l-1})$. By the induction hypothesis, (M, T) is $R_i(\delta_i^{\downarrow l-1})$-complete, so $M, t \models \nabla_i R_i(\delta_i^{\downarrow l-1})$. For $j \neq i$, because $(M_{t,j}, T_{t,j})$ is $R_j(\delta_i)$-complete and $R_j(t) = T_{t,j}$, $M, t \models \nabla_j R_j(\delta_i)$, hence $M, t \models \nabla_j R_j(\delta_i)^{\downarrow l-1}$. By Proposition 4, $R_j(\delta_i)^{\downarrow l-1} = R_j(\delta_i^{\downarrow l-1})$, so $M, t \models \nabla_j R_j(\delta_i^{\downarrow l-1})$. ■

Now, we get the following theorem via Lemma 2.

Theorem 3 (The $\mathsf{K45_n}$ theorem) *Let δ be a $\mathsf{K45_n}$ satisfiable canonical formula. Then $kforget_{\mathsf{K45_n}}(\delta, p) \equiv \delta^p$.*

Proof: The proof is the same as that of the basic theorem except Step 2 of the induction case. By the $\mathsf{K45_n}$ lemma, for every $i \in \mathcal{A}$, there exist (M_i, T_i) and ρ_i : $(M_i, T_i) \underline{\leftrightarrow}_p (M', R'_i(s'))$. Add a new copy of M_i into M, let $s R_i t$ for all $t \in T_i$, and expand ρ with ρ_i. ■

From the construction of the $\mathsf{K45_n}$ lemma and the $\mathsf{K45_n}$ theorem, if the given model M' is a $\mathsf{KD45_n}$ model, then we can acquire a $\mathsf{KD45_n}$ model M. Hence we can get

Theorem 4 *Let δ be a $\mathsf{KD45_n}$ satisfiable canonical formula. Then $kforget_{\mathsf{KD45_n}}(\delta, p) \equiv \delta^p$.*

We proceed to $\mathsf{S5_n}$. The models constructed in the proofs of the $\mathsf{K45_n}$ lemma and the $\mathsf{K45_n}$ theorem may not be reflexive. We can fix it via adding edges as follows: for all $i \in \mathcal{A}$,

1. for all $t \in S$, we let $t R_i t$; (reflexive edges)
2. for all $t, u \in S$, if $t R_i u$, we let $u R_i t$. (symmetric edges)

It is easy to verify that the new models are $\mathsf{S5_n}$ models.

Example 6 In Figure 3, given a model (M, T) shown in Figure 2, we add blue edges so that it is an $\mathsf{S5_n}$ model. Note that it is not necessary to add edges for four submodels $(M_{t_1,j}, T_{t_1,j})$, $(M_{t_1,k}, T_{t_1,k})$, $(M_{t_2,j}, T_{t_2,j})$ and

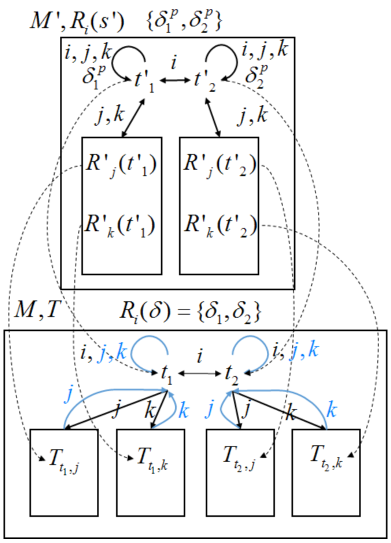

Figure 3: Adding reflexive and symmetric edges

$(M_{t_2,k}, T_{t_2,k})$. All of them are $\mathsf{S5_n}$ models since they are acquired by the induction hypothesis. We just add the edges as follows:

1. Let $t_1 R_j t_1$, $t_1 R_k t_1$, $t_2 R_j t_2$ and $t_2 R_k t_2$;
2. For $u \in T_{t_1,j}$ (resp. $T_{t_1,k}$), let $uR_j t_1$ (resp. $uR_k t_1$).
3. For $u \in T_{t_2,j}$ (resp. $T_{t_2,k}$), let $uR_j t_2$ (resp. $uR_k t_2$).

All that remains now is to prove that the constructed models satisfy the original formulas. Since the logic $\mathsf{S5_n}$ is both $\mathsf{K45_n}$ and $\mathsf{T_n}$ logics, any $\mathsf{S5_n}$ satisfiable canonical formula has the reflexive and identical-children properties (Propositions 9 and 11). Hence, we get

Theorem 5 (The $\mathsf{S5_n}$ theorem) *Let δ be an $\mathsf{S5_n}$ satisfiable canonical formula. Then $kforget_{\mathsf{S5_n}}(\delta, p) \equiv \delta^p$.*

Proof: Let $k = dep(\delta)$. Here, we only verify that $M, s \models \delta$ when $k \geq 1$. We prove by induction on $k - l$ that for any $l \leq k$, $M, s \models \delta^{\downarrow l}$ still holds. Base case. $M, s \models w(\delta)$, which is $\delta^{\downarrow k}$. Induction case. Suppose that $M, s \models \delta^{\downarrow l}$. To show that $M, s \models \delta^{\downarrow l-1}$, it suffices to show that for each $i \in \mathcal{A}$, $M, s \models \nabla_i R_i(\delta^{\downarrow l-1})$, *i.e.*, $(M, R_i(s))$ is $R_i(\delta^{\downarrow l-1})$-complete. Here, we only prove that for every $t \in R_i(s)$, there exists $\eta \in R_i(\delta^{\downarrow l-1})$ s.t. $M, t \models \eta$. The other direction can be similarly proved.

Suppose that $t = s$. By the reflexive property, $\delta^{\downarrow l} \in R_i(\delta^{\downarrow l-1})$. Let $\eta = \delta^{\downarrow l}$. This, together with the hypothesis, imply that $M, s \models \eta$ where $\eta \in R_i(\delta^{\downarrow l-1})$.

Suppose that $t \in S \setminus \{s\}$. Before adding the edges, by construction, there exists $\delta_i \in R_i(\delta)$ s.t. $M, t \models \delta_i$. Let $\eta = \delta_i^{\downarrow l-1}$. After adding the edges, that $M, t \models w(\eta) \wedge \bigwedge_{j \in \mathcal{A} \setminus \{i\}} \nabla_j R_j(\eta)$ holds. Now, we prove that $M, t \models \nabla_i R_i(\eta)$, *i.e.*, $(M, R_i(t))$ is $R_i(\eta)$-complete. By transitive and Euclidean properties, we get $R_i(t) = R_i(s)$. By the identical children property, $R_i(\eta) = R_i(\delta^{\downarrow l})$. These, together with the hypothesis, imply that $M, t \models \nabla_i R_i(\eta)$. ■

As a corollary of Theorems 2-5, we get

Corollary 3 $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$ *and* $\mathsf{S5_n}$ *are closed under forgetting.*

By Proposition 8 and Corollary 3, we can get the next corollary which settles the open problems regarding uniform interpolation in multi-agent modal logics.

Corollary 4 $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$ *and* $\mathsf{S5_n}$ *do have uniform interpolation.*

5 Conclusions

In this paper, we have studied forgetting in multi-agent modal logics. We adopted the semantic definition of existential bisimulation quantifiers as that of forgetting. We showed that except the two modal systems $\mathsf{K4_n}$ and $\mathsf{S4_n}$, which have been shown lack of uniform interpolation, the other main multi-agent modal systems, namely $\mathsf{K_n}$, $\mathsf{D_n}$, $\mathsf{T_n}$, $\mathsf{K45_n}$, $\mathsf{KD45_n}$, and $\mathsf{S5_n}$, are closed under forgetting.

To achieve the above results, we presented a syntactical method of forgetting: to forget an atom from an arbitrary modal formula, first transform the formula to a disjunction of canonical formulas, then remove all the unsatisfiable ones, and lastly substitute any literal of the atom with $\top$. To prove that this syntactical method of forgetting produces the desired result, given a model satisfying the formula obtained from a satisfiable canonical formula via literal elimination, we construct a model satisfying the original formula and p-similar to the given model. To this end, we analyze properties of satisfiable canonical formulas in different logics, and use multi-pointed models to gain flexibility in model construction.

In multi-agent systems, background information is formalized as common knowledge of propositional formulas. We have also shown that the extensions of the above logics with propositional common knowledge are closed under forgetting. These results were omitted for lack of space.

It is well-known that the size of a canonical formula is non-elementary. So the proposed syntactic method of forgetting is obviously an unpractical one. Thus one topic for future research is to investigate more efficient approaches for computing forgetting. Yet another one is to explore forgetting for distributed knowledge and more general cases of common knowledge. It would also be interesting to apply forgetting to progression and abductive reasoning of knowledge and belief in multi-agent systems.

Acknowledgments

Liangda Fang and Yongmei Liu acknowledge support from the Natural Science Foundation of China under Grant Nos. 61572535 and 61573386. Hans van Ditmarsch acknowledges support from ERC project EPS313360. He is also affiliated to IMSc (Institute of Mathematical Sciences), Chennai, India, and to Zhejiang University, China.

References

[Baral and Zhang, 2005] Chitta Baral and Yan Zhang. Knowledge updates: Semantics and complexity issues. *Artificial Intelligence*, 164(1-2):209–243, 2005.

[Bílková, 2007] Marta Bílková. Uniform Interpolation and Propositional Quantifiers in Modal Logics. *Studia Logica*, 85(1):1–31, 2007.

[D'Agostino and Lenzi, 2005] Giovanna D'Agostino and Giacomo Lenzi. An axiomatization of bisimulation quantifiers via the μ-calculus. *Theoretical Computer Science*, 338(1):64–95, 2005.

[D'Agostino and Lenzi, 2006] Giovanna D'Agostino and Giacomo Lenzi. On modal μ-calculus with explicit interpolants. *Journal of Applied Logic*, 4(3):256–278, 2006.

[Eiter and Wang, 2008] Thomas Eiter and Kewen Wang. Semantic forgetting in answer set programming. *Artificial Intelligence*, 172(14):1644–1672, 2008.

[Fang *et al.*, 2015] Liangda Fang, Yongmei Liu, and Ximing Wen. On the Progression of Knowledge and Belief for Nondeterministic Actions in the Situation Calculus. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2015)*, pages 2955–2963, 2015.

[Ghilardi and Zawadowski, 1995] Silvio Ghilardi and Marek W. Zawadowski. A sheaf representation and duality for finitely presenting heyting algebras. *Journal of Symbolic Logic*, 60(3):911–939, 1995.

[Ghilardi and Zawadowski, 2000] Silvio Ghilardi and Marek W. Zawadowski. From Bisimulation Quantifiers to Classifying Toposes. In *Proceedings of the Third Workshop on Advances in Modal Logic (AiML-2000)*, volume 3, pages 193–220, 2000.

[Ghilardi, 1995] Silvio Ghilardi. An algebraic theory of normal forms. *Annals of Pure and Applied Logic*, 71(3):189–245, 1995.

[Herzig *et al.*, 2003] Andreas Herzig, Jérôme Lang, and Pierre Marquis. Action representation and partially observable planning using epistemic logic. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-2003)*, pages 1067–1073, 2003.

[Janin and Walukiewicz, 1995] David Janin and Igor Walukiewicz. Automata for the modal μ-calculus and related results. In *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science (MFCS-1995)*, pages 552–562, 1995.

[Konev *et al.*, 2009] Boris Konev, Dirk Walther, and Frank Wolter. Forgetting and Uniform Interpolation in Large-Scale Description Logic Terminologies. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-2009)*, pages 830–835, 2009.

[Koopmann and Schmidt, 2014] Patrick Koopmann and Renate A. Schmidt. Count and Forget: Uniform Interpolation of $\mathcal{SHQ}$-Ontologiess. In *Proceedings of Seventh International Joint Conference on Automated Reasoning (IJCAR-2014)*, pages 434–448, 2014.

[Koopmann and Schmidt, 2015] Patrick Koopmann and Renate A. Schmidt. Uniform Interpolation and Forgetting for $\mathcal{ALC}$ Ontologies with ABoxes. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-2015)*, pages 175–181, 2015.

[Lang and Marquis, 2010] Jérôme Lang and Pierre Marquis. Reasoning under inconsistency: A forgetting-based approach. *Artificial Intelligence*, 174(12-13):799–823, 2010.

[Lin and Reiter, 1994] Fangzhen Lin and Raymond Reiter. Forget It! In *Proceedings of AAAI Fall Symposium on Relevance*, pages 154–159, 1994.

[Lin and Reiter, 1997] Fangzhen Lin and Raymond Reiter. How to progress a database. *Artificial Intelligence*, 92(1-2):131–167, 1997.

[Lin, 2001] Fangzhen Lin. On strongest necessary and weakest sufficient conditions. *Artificial Intelligence*, 128(1-2):143–159, 2001.

[Liu and Lakemeyer, 2009] Yongmei Liu and Gerhard Lakemeyer. On First-Order Definability and Computability of Progression for Local-Effect Actions and Beyond. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-2009)*, pages 860–866, 2009.

[Liu and Wen, 2011] Yongmei Liu and Ximing Wen. On the Progression of Knowledge in the Situation Calculus. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-2011)*, pages 976–982, 2011.

[Ludwig and Konev, 2014] Michel Ludwig and Boris Konev. Practical Uniform Interpolation and Forgetting for $\mathcal{ALC}$ TBoxes with Applications to Logical Difference. In *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning (KR-2014)*, pages 318–327, 2014.

[Lutz and Wolter, 2011] Carsten Lutz and Frank Wolter. Foundations for Uniform Interpolation and Forgetting in Expressive Description Logics. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-2011)*, pages 989–995, 2011.

[Moss, 2007] Lawrence S. Moss. Finite models constructed from canonical formulas. *Journal of Philosophical Logic*, 36(6):605–640, 2007.

[Pattinson, 2013] Dirk Pattinson. The Logic of Exact Covers: Completeness and Uniform Interpolation. In *Proceedings of the Twenty-Eighth ACM/IEEE Symposium on Logic in Computer Science (LICS-2013)*, pages 418–427, 2013.

[Studer, 2009] Thomas Studer. Common knowledge does not have the beth property. *Information Processing Letters*, 109(12):611–614, 2009.

[Su *et al.*, 2009] Kaile Su, Abdul Sattar, Guanfeng Lv, and Yan Zhang. Variable Forgetting in Reasoning about Knowledge. *Journal of Artificial Intelligence Research*, 35(2):677–716, 2009.

[van Ditmarsch *et al.*, 2009] Hans van Ditmarsch, Andreas Herzig, Jérôme Lang, and Pierre Marquis. Introspective forgetting. *Synthese*, 169(2):405–423, 2009.

[Wang *et al.*, 2009] Kewen Wang, Zhe Wang, Rodney Topor, Jeff Z. Pan, and Grigoris Antoniou. Concept and role forgetting in $\mathcal{ALC}$ ontologies. In *Proceedings of the Eighth International Semantic Web Conference (ISWC-2009)*, pages 666–681, 2009.

[Wang *et al.*, 2010] Zhe Wang, Kewen Wang, Rodney Topor, and Jeff Z. Pan. Forgetting for knowledge bases in DL-Lite. *Annals of Mathematics and Artificial Intelligence*, 58(1-2):117–151, 2010.

[Wang *et al.*, 2014] Yisong Wang, Yan Zhang, Yi Zhou, and Mingyi Zhang. Knowledge Forgetting in Answer Set Programming. *Journal of Artificial Intelligence Research*, 50(1):31–70, 2014.

[Zhang and Foo, 2006] Yan Zhang and Norman Y. Foo. Solving logic program conflict through strong and weak forgettings. *Artificial Intelligence*, 170(8-9):739–778, 2006.

[Zhang and Zhou, 2009] Yan Zhang and Yi Zhou. Knowledge forgetting: Properties and applications. *Artificial Intelligence*, 173(16-17):1525–1537, 2009.