Knowledge updates: Semantics and complexity issues

Chitta Baral

Department of Computer Science and Engineering
Arizona State University
Tempe, AZ 85233, USA
chitta@asu.edu

Yan Zhang

School of Computing & Information Technology
University of Western Sydney
Penrith South DC, NSW 1797, Australia
yan@cit.uws.edu.au

Abstract

We consider the problem of how an agent's knowledge can be updated. We propose a formal method of knowledge update on the basis of the semantics of modal logic S5. In our method, an update is specified according to the minimal change on both the agent's actual world and knowledge. We discuss general minimal change properties of knowledge update and show that our knowledge update operator satisfies all Katsuno and Mendelzon's update postulates. We characterize several specific forms of knowledge update which have important applications in reasoning about change of agents' knowledge. We also examine the persistence property of knowledge and ignorance associated with knowledge update.

We then investigate the computational complexity of model checking for knowledge update. We first show that in general the model checking for knowledge update is Σ_2^P-complete, which places the problem at the same layer in the polynomial hierarchy of the traditional model based belief update (e.g. PMA). We then identify a subclass of knowledge update problems that has polynomial time complexity for model checking. We point out that some important knowledge update problems belong to this subclass. We further address another interesting subclass of knowledge update problems for which the complexity of model checking is NP-complete.

[1] This paper is based on the authors' two conference papers from IJCAI-2001 [1] and KR-2002 [2].

1 Introduction and motivation

The well-studied issues of belief updates and belief revision [13] are concerned with the update and revision aspects of an agent's belief with respect to new beliefs. The notion of belief update has been used, and often serves as a guideline [12,24], in reasoning about the effect of (*world altering*) actions on the state of the world. Thus if ϕ represents the agent's belief about the world and the agent performs an action that is supposed to make ψ true in the resulting world, then the agent's belief about the resulting world can be described by $\phi \diamond \psi$, where $\diamond$ is the update operator of choice.

Now let us consider reasoning about sensing actions [22,23], which in their pure form, when executed, do not change the world, but change the agent's knowledge about the world. Let $sense_f$ be a sensing action whose effect is that after it is executed the agent knows whether f is true or not. This can be expressed as $Kf \vee K\neg f$, where K is the modal operator $Knows$. The current theory of belief updates does not tell us how to do updates with respect to such gain in knowledge due to a sensing action. (Note that we can not just have $\psi \equiv f \vee \neg f$ and use the the notion of belief update, as $f \vee \neg f$ is a tautology). The major goal of this paper is to define a notion of *knowledge update*, analogous to belief update, where the original theory (ϕ) and the new theory (ψ) are in a language that can express knowledge. Such a notion would not only serve as a guideline to reason about pure and mixed sensing actions in presence of constraints, but also allow us to reason about actions corresponding to *forgetting* and *ignorance*.

We then investigate the computational complexity of model checking for knowledge update. We first show that in general the model checking problem for knowledge update is Σ_2^P-complete, which places the problem at the same layer of the polynomial hierarchy than the traditional model based belief update (e.g. PMA) [16]. We further identify a subclass of knowledge update problems for which model checking can be achieved in polynomial time. We show that some important knowledge update problems belong to this subclass. Finally, we address another interesting subclass of knowledge updates for which the complexity for model checking is NP-complete.

The structure of the rest of the paper is as follows. In Section 2 we start with describing the particular modal logic that we plan to use in expressing knowledge, and describe the notion of k-models analogous to 'models' in classical logic. We define closeness between k-models and use it to define a particular notion of knowledge update. In Section 3 we discuss minimal change properties of knowledge update. An interesting result shows that our knowledge update operator satisfies all Katsuno and Mendelzon's update postulates. In Section 4, we present alternative characterizations of four particular knowledge updates – *gaining knowledge, ignorance, sensing, and forgetting*, and show their equivalence to our original notion of knowledge update. Some of these alternative characterizations are based on the formulation of reasoning about sensing actions, and thus our equivalence results serve as justification of the intuitiveness of our definition of knowledge update. In Section 5 we explore sufficiency conditions that guarantee persistence of knowledge (or ignorance) during a knowledge update. From Section 6 we start to investigate model checking complexity for knowledge update. In Section 6 we first give general background on computational complexity. In Section 7, we study the model checking complexity for the general case of knowledge update. In Section 8, we define a subclass of knowledge update problems whose model checking can be achieved in polynomial time. In Section 9, we further address an interesting intractable subclass of knowledge update problems whose model checking is lower than the genral case. Finally, in Section 10, we conclude this paper with some remarks discussions.

2 Closeness between k-models and knowledge update

In this section, we describe formal definitions for knowledge update. Our formalization will be based on the semantics of the propositional modal logic S5 with a single agent. In general, under Kripke semantics, a *Kripke structure* is a triple (W, R, π), where W is a set of possible worlds, R is an equivalence relation on W, and π is a truth assignment function that assigns a propositional valuation to each world in W. Given a Kripke structure $S = (W, R, \pi)$, a *Kripke interpretation* is a pair $M = (S, w)$, where $w \in W$ is referred to the *actual world* of M. The entailment relation $\models$ between Kripke interpretations and formulas is defined to provide semantics for formulas of S5 [4].

In the case of single agent, however, we may restrict ourselves to those S5 structures in which the relation R is universal, i.e. each world is accessible from every world, and worlds are identified with the set of atoms true at the worlds [20]. To simplify a comparison between two worlds (e.g. Definition 2), we may view an atom $p \in w$ iff $w \models p$. Therefore, in our context a Kripke structure (W, R, π) is uniquely characterized by W and we may simplify a Kripke interpretation as a pair (W, w) which we call a *k-model*, where w indicates the actual world of the agent and W presents all possible worlds that the agent may access. Note that w is in W for any k-model (W, w).

In our following description, we use $a, b, c, \cdots$ to denote primitive propositional atoms; $\phi, \psi, \upsilon, \cdots$ to denote propositional formulas without including modalities (we also call them *objective* formulas); and $\alpha, \beta, \gamma, \mu, \cdots$ to denote formulas that may contain modal operator K. For convenience, we use $T \equiv \alpha_1 \wedge \cdots \wedge \alpha_k$ to represent a finite set of formulas $\{\alpha_1, \cdots, \alpha_k\}$ and call T a *(knowledge) set.*

Definition 1 *(***S5 Semantics***) Let $\mathcal{P}$ be the set of all primitive propositions in the language. The entailment relation $\models$ under normal S5 semantics is defined as follows:*

(1) $(W, w) \models p$ iff p is primitive (i.e. $p \in \mathcal{P}$) and $w \models p$;
(2) $(W, w) \models \alpha \wedge \beta$ iff $(W, w) \models \alpha$ and $(W, w) \models \beta$;
(3) $(W, w) \models \neg\alpha$ iff it is not the case that $(W, w) \models \alpha$;
(4) $(W, w) \models K\alpha$ iff $(W, w') \models \alpha$ for all $w' \in W$.

Given a formula T, $M = (W, w)$ is called a k-model of T if $M \models T$. We use $Mod(T)$ to denote the set of all k-models of T. For an objective formula ϕ, $Mod(\phi)$ simply denotes the set of worlds w where $w \models \phi$. In this case, w is also called a *model* of ϕ. For a formula α, we say that T *entails* α, denoted as $T \models \alpha$, if for every k-model M of T, $M \models \alpha$.

Now the basic problem of knowledge update that we would like to investigate is formally described as follows: given a *k-model* $M = (W, w)$, that is usually viewed as a *knowledge state* of an agent, and a formula μ - the agent's new knowledge that may contain modal operator K, how do we update M to another k-model $M' = (W', w')$ such that $M' \models \mu$ and M' is *minimally different* from M with respect to some criterion. To approach this problem, we first need to provide a definition of *closeness* between two k-models with respect to a given k-model.

Definition 2 *(***Closeness between k-models***) Let $M = (W, w)$, $M_1 = (W_1, w_1)$ and $M_2 = (W_2, w_2)$ be three k-models. We say M_1 is* as close to M as *M_2, denoted as $M_1 \leq_M M_2$, if:*

(1) $(w_1 \setminus w \cup w \setminus w_1) \subset (w_2 \setminus w \cup w \setminus w_2)$; or
(2) $w_1 = w_2$ and one of the following conditions holds:
  (i) $W_1 = W_2$;
  (ii) if $W \subset W_1$, then (a) there exist some ϕ and ψ such that $M \models K\phi$ and $M_2 \not\models K\phi$ and $M \not\models K\psi$ and $M_2 \models K\psi$, or (b) for any ϕ if $M \models K\phi$ and $M_1 \not\models K\phi$, then $M_2 \not\models K\phi$;
  (iii) if $W_1 \subset W$, then condition (a) above is satisfied, or (c) for any ϕ if $M \not\models K\phi$ and $M_1 \models K\phi$, then $M_2 \models K\phi$;
  (iv) if $W \not\subset W_1$ and $W_1 \not\subset W$, then conditions (b) and (c) above are satisfied.

We denote $M_1 <_M M_2$ if $M_1 \leq_M M_2$ and $M_2 \not\leq_M M_1$.

In the above definition, condition 1 simply says that the symmetric differences between w and w_1 is a proper subset of that between w and w_2, while in condition 2, (ii), (iii) and (iv) express that different preferences are applied to compare knowledge between M_1 and M_2 with respect to M (note that (i) is just for the case that M_1 and M_2 are identical). For convenience, given a k-model M, if we denote $KM = \{\phi \mid \text{for all } w \in W, w \models \phi\}$, then (a) is equivalent to $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$; (b) is equivalent to $KM \setminus KM_1 \subseteq KM \setminus KM_2$; and (c) is equivalent to $KM_1 \setminus KM \subseteq KM_2 \setminus KM$. Also (b) and (c) together present a difference on both knowledge decrease and increase between M_1 and M_2 in terms of M.

Note that during the comparison between two k-models, we give preference to the change of the actual world over the change of knowledge about the world. (The closeness criterion between actual worlds that we use is the commonly used criterion [25] based on symmetric difference). For instance, if the actual world of a k-model M_1 is closer to the actual world of M than the actual world of another k-model M_2, we will think M_1 is closer to M and the comparison of knowledge between M_1 and M_2 is ignored. Only when both M_1 and M_2 have the same actual world, we will compare knowledge of M_1 and M_2 in terms of M. This seems to be intuitive to us. In fact, the comparison between actual worlds determines the *actual distance* between two k-models to the given k-model M. If M_1 and M_2 have the same actual distance to M, the *knowledge distance* is then taken into account.

Basically, condition (ii) (or (iii) resp.) in Definition 2 defines a knowledge preference based on knowledge decrease (or increase resp). That is, if M_1 *only* looses knowledge from M (or *only* gains some knowledge to M, resp.), then M_1 is preferred over those k-models that have both knowledge decrease and increase from M, i.e. (a), and also preferred over those k-models that only loose more knowledge from M (or add more knowledge to M, resp.) than M_1 does, i.e. (b) or (c) respectively. Condition (iv), on the other hand, deals with the mixed situation that M_1 has both knowledge decrease and increase from M. In this case, a combined difference on knowledge decrease and increase is applied to determine the knowledge distance, i.e. (b) *and* (c). Conditions (ii), (iii) and (iv) can be illustrated by the following figures respectively.

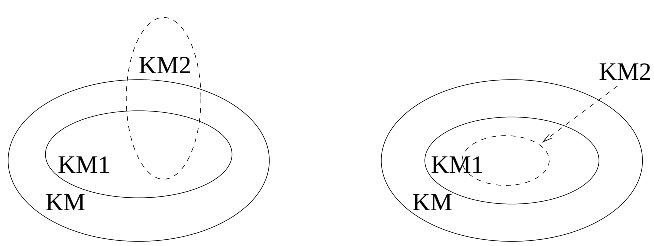

Fig. 1. $M_1 \leq_M M_2$ under the condition $w_1 = w_2$ and $W \subset W_1$: (a) or (b) holds.

Definition 3 *(*k**-model Update***) Let $M = (W, w)$ be a k-model and μ a formula. A k-model $M' = (W', w')$ is called a* possible resulting k-model *after updating M with μ if and only if the following conditions hold:*

(1) $M' \models \mu$*;*
(2) there does not exist another k-model $M'' = (W'', w'')$ such that $M'' \models \mu$ and $M'' <_M M'$.

We denote the set of all possible resulting k-models after updating M with μ as $Res(M, \mu)$.

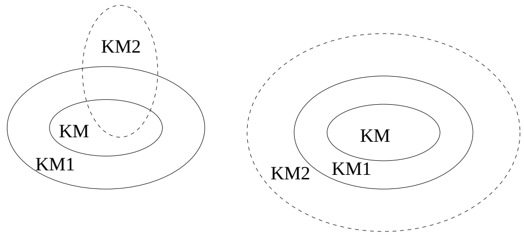

Fig. 2. $M_1 \leq_M M_2$ under the condition $w_1 = w_2$ and $W_1 \subset W$: (a) or (c) holds.

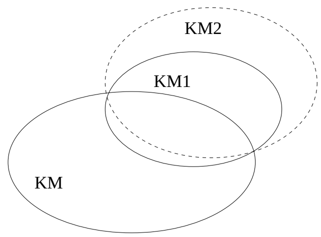

Fig. 3. $M_1 \leq_M M_2$ under the condition $w_1 = w_2$, $W \not\subset W_1$ and $W_1 \not\subset W$: (b) and (c) hold.

Example 1 *Let $T \equiv Kc \wedge \neg Ka \wedge \neg Kb \wedge K(a \vee b)$ and $\mu \equiv K\neg c$. We denote*

$$w_0 = \{a, b, c\},\ w_1 = \{a, c\},\ w_2 = \{b, c\},$$
$$w_3 = \{c\},\ w_4 = \{a, b\},\ w_5 = \{a\},$$
$$w_6 = \{b\},\ w_7 = \emptyset.$$

Clearly, $M_0 = (\{w_0, w_1, w_2\}, w_0)$ is a k-model of T. Consider the update of M_0 with μ. Let $M_1 = (\{w_4, w_5, w_6\}, w_4)$. Now we show that M_1 is a possible resulting k-model after updating M_0 with μ.

Since $(w_0 \setminus w_4 \cup w_4 \setminus w_0) = \{c\}$, we first consider any possible k-model $M' = (W', w')$ such that $(w_0 \setminus w' \cup w' \setminus w_0) \subset \{c\}$. Clearly, the only possible w' would be w_0 itself. Let $M' = (W', w_0)$, where W' is a subset of $\{w_0, \cdots, w_7\}$. However, since $c \in w_0$, there does not exist any W' such that $M' \models K\neg c$. Therefore, from Definition 2, only condition 2 can be used to find a possible M' such that $M' <_M M_1$. So we assume $M' = (W', w_4)$. On the other hand, from M_0 and M_1, it is easy to see that $KM_0 = \{c, a \vee b\}$ and $KM_1 = \{\neg c, a \vee b\}$[2]. Then we have $KM_0 \setminus KM_1 = \{c\}$ and $KM_1 \setminus KM_0 = \{\neg c\}$. Ignoring the detailed verifications, we can show that there does not exist such $M' = (W', w_4)$ satisfying $KM_0 \setminus KM' = KM' \setminus KM_0 = \emptyset$.

Based on the k-model update, updating a formula (knowledge set) T in terms of another formula μ is then achieved by updating every k-model of T with μ.

[2] For simplicity, here we only consider the *prime* formulas ϕ in KM in the sense that if $\phi \in KM$, then there is not another ψ such that $\models \psi \supset \phi$ and $\psi \in KM$.

Definition 4 *(***Knowledge update***) Let T and μ be two formulas. The* update *of T with μ, denoted as $T \diamond \mu$, is defined by $Mod(T \diamond \mu) = \bigcup_{M \in Mod(T)} Res(M, \mu)$.*

3 Minimal change of knowledge update

In this section, we investigate minimal change properties of knowledge update. Specifically, we examine the relationship between knowledge update and the classical Katsuno and Mendelzon's update postulates. TiFirstly, the following proposition presents some useful results about knowledge update.

Proposition 1 *Let $M_1 = (W_1, w_1)$ and $M_2 = (W_2, w_2)$ be two k-models. Then the following properties hold:*

(1) $\phi \in KM_1$ iff $W_1 \subseteq Mod(\phi)$;
(2) $W_1 \subseteq W_2$ iff $KM_2 \subseteq KM_1$;
(3) $KM_1 = KM_2$ iff $W_1 = W_2$;
(4) Let $M' = (W_1 \cup W_2, w')$, then $KM' = KM_1 \cap KM_2$;
(5) Let $w' \in W_1 \cap W_2$ and $M' = (W_1 \cap W_2, w')$, then $KM_1 \cup KM_2 \subseteq KM'$.

Proof: (1). ($\Rightarrow$) From $\phi \in KM$, we have for all $w' \in W$, $w' \models \phi$ (note that ϕ is a formula without containing modal operator K). That is, $w' \in Mod(\phi)$. So $W \subseteq Mod(\phi)$.

($\Leftarrow$) Suppose $W \subseteq Mod(\phi)$. Then we have for any $w' \in W$, $w' \models \phi$. That is, $\phi \in KM$.

(2). Let $\bigwedge KM_1$ and $\bigwedge KM_2$ be the conjunctions of all formulas in KM_1 and KM_2 respectively. Then it is clear that $Mod(\bigwedge KM_1) = W_1$ and $Mod(\bigwedge KM_2) = W_2$. So we have $KM_2 \subseteq KM_1$ iff $Mod(\bigwedge KM_1) \subseteq Mod(\bigwedge KM_2)$ iff $W_1 \subseteq W_2$.

(3). In the proof of 2, we stated that $Mod(\bigwedge KM_1) = W_1$ and $Mod(\bigwedge KM_2) = W_2$. So $KM_1 = KM_2$ iff $Mod(\bigwedge KM_1) = Mod(\bigwedge KM_2)$ iff $W_1 = W_2$ [3].

(4). According to the definition of M', we have $\phi \in KM'$ iff for all $w' \in W_1 \cup W_2$, $w' \models \phi$ iff for all $w' \in W_1$, $w' \models \phi$ and for all $w'' \in W_2$, $w'' \models \phi$ iff $\phi \in KM_1$ and $\phi \in MK_2$ iff $\phi \in KM_1 \cap KM_2$.

(5). According to the definition of M', from $\phi \in KM_1 \cup KM_2$, we have $\phi \in KM_1$ or $\phi \in KM_2$. So for any $w' \in W_1 \cap W_2$, we have $w' \models \phi$. That

[3] This property also indicates that knowledge represented by a k-model $M = (W, w)$ is independent of the actual world w.

is, $\phi \in KM'$. ■

Given a set of k-models $\mathcal{S}$ and a k-model M, let $\leq_M$ be an ordering on $\mathcal{S}$ as we defined in Definition 2. By $Min(\mathcal{S}, \leq_M)$ we mean the set of all elements in $\mathcal{S}$ that are minimal with respect to ordering $\leq_M$. The following proposition simply shows that $\leq_M$ is a partial ordering.

Proposition 2 *Let M be a k-model. Then $\leq_M$ defined in Definition 2 is a partial ordering.*

Proof: From Definition 2, it is clear that $\leq_M$ is reflexive and antisymmetric. Now we prove $\leq_M$ is also transitive. Let $M = (W, w)$, $M_1 = (W_1, w_1)$, $M_2 = (W_2, w_2)$ and $M_3 = (W_3, w_3)$ be k-models, and $M_1 \leq_M M_2$ and $M_2 \leq_M M_3$. Now we prove $M_1 \leq_M M_3$.
Case 1. Suppose $M_1 \leq_M M_2$ is due to condition (1) in Definition 2, i.e. $(w_1 \setminus w \cup w \setminus w_1) \subset (w_2 \setminus w \cup w \setminus w_2)$. Consider $M_2 \leq_M M_3$. According to Definition 2, either condition (1) or (2) is satisfied. If condition (1) is satisfied, then $(w_2 \setminus w \cup w \setminus w_2) \subset (w_3 \setminus w \cup w \setminus w_3)$. This follows $(w_1 \setminus w \cup w \setminus w_1) \subset (w_3 \setminus w \cup w \setminus w_3)$. So $M_1 \leq_M M_3$. If condition (2) is satisfied, it means $w_2 = w_3$, it also follows $(w_1 \setminus w \cup w \setminus w_1) \subset (w_3 \setminus w \cup w \setminus w_3)$, and therefore $M_1 \leq_M M_3$.
Case 2. Now suppose $M_1 \leq_M M_2$ is due to condition (2) in Definition 2, i.e. $w_1 = w_2$ and one of conditions (i), (ii), (iii), or (iv) is satisfied. If $M_2 \leq_M M_3$ is due to condition (1) in Definition 2, i.e. $(w_2 \setminus w \cup w \setminus w_2) \subset (w_3 \setminus w \cup w \setminus w_3)$, it follows that $M_1 \leq_M M_3$ because $w_1 = w_2$. Suppose $M_2 \leq_M M_3$ is due to condition (2) in Definition 2, that is, $w_2 = w_3$ and one of conditions (i), (ii), (iii) or (iv) is satisfied. Here we only consider the following three cases, while all other cases can be proved in a similar way.
Case 2.1. Both $M_1 \leq_M M_2$ and $M_2 \leq_M M_3$ are due to condition (2) and (ii) in Definition 2. Under this case, we can only have (a) $KM_3 \subset KM_2 \subset KM_1 \subset KM$; or (b) $KM_2 \subset KM_1 \subset KM$ but $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$. Clearly, in either case, we have $M_1 \leq_M M_3$.
Case 2.2. $M_1 \leq_M M_2$ is due to condition (2) and (ii) and $M_2 \leq_M M_3$ are due to condition (2) and (iii) in Definition 2. By analyzing Definition 2, it concludes that this situation will never occur. This is because from $M_1 \leq_M M_2$, we can only have either $KM_2 \subset KM$ or $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$, and from $M_2 \leq_M M_3$, we can only have $KM \subset KM_2$. Obviously, these two cases conflict with each other.
Case 2.3. $M_1 \leq_M M_2$ is due to condition (2) and (ii) and $M_2 \leq_M M_3$ are due to condition (2) and (iv) in Definition 2. In this case, we will have $KM_1 \subset KM$ and $KM_2 \setminus KM \neq \emptyset$ and $KM \setminus KM_2 \neq \emptyset$. This implies $M_1 \leq_M M_3$. ■

Theorem 1 *Let T and μ be two formulas. Then $Mod(T \diamond \mu) = \bigcup_{M \in Mod(T)} Min(Mod(\mu), \leq_M)$.*

Proof: To prove the result, we only need to show that for each k-model M of T, $Res(M, \mu) = Min(Mod(\mu), \leq_M)$. Let $M' \in Res(M, \mu)$. Since $M' \models \mu$, $M' \in Mod(\mu)$. On the other hand, according to Definition 3, for any $M'' \in Mod(\mu)$, we have $M'' \not<_M M'$. That is, $M' \in Min(Mod(\mu), \leq_M)$. So $Res(M, \mu) \subseteq Min(Mod(\mu), \leq_M)$. Similarly, we can show $Min(Mod(\mu), \leq_M) \subseteq Res(M, \mu)$. ∎

The above theorem provides an important characterization on knowledge update in terms of a particular minimal change criterion. Now the question we are interested in is whether our knowledge update operator satisfies some classical properties of belief (knowledge base) update. In recent years, belief update has been extensively studied by many researchers and its difference from belief revision is well understood [9,18,27]. From the observation of semantic difference between belief update and revision, Katsuno and Mendelzon argued that the original revision postulates proposed by Gardenfors *el al.* [5] are not quite suitable for update, and ignoring such difference may lead to unreasonable solutions [13]. Instead, Katsuno and Mendelzon proposed alternative postulates for any update operator $\diamond$ as follows.

- (U1) $T \diamond \mu$ implies μ.
- (U2) If T implies μ then $T \diamond \mu \equiv T$.
- (U3) If both T and μ are satisfiable then $T \diamond \mu$ is also satisfiable.
- (U4) If $T_1 \equiv T_2$ and $\mu_1 \equiv \mu_2$ then $T \diamond \mu_1 \equiv T_2 \diamond \mu_2$.
- (U5) $(T \diamond \mu) \wedge \alpha$ implies $T \diamond (\mu \wedge \alpha)$.
- (U6) If $T \diamond \mu_1$ implies μ_2 and $T \diamond \mu_2$ implies μ_1 then $T \diamond \mu_1 \equiv T \diamond \mu_2$.
- (U7) If T is complete then $(T \diamond \mu_1) \wedge (T \diamond \mu_2)$ implies $T \diamond (\mu_1 \vee \mu_2)$.
- (U8) $(T_1 \vee T_2) \diamond \mu \equiv (T_1 \diamond \mu) \vee (T_2 \diamond \mu)$.

Under the context of S5 modal logic, we may think all formulas occurring in the above postulates are S5 formulas. The following theorem shows that our knowledge update operator satisfies all these postulates.

Theorem 2 *Knowledge update operator $\diamond$ defined in Definition 4 satisfies Katsuno and Mendelzon's update postulates (U1)-(U8).*

Proof: From Definitions 3 and 4, it is easy to verify $\diamond$ satisfies postulates (U1)-(U4). Now we prove $\diamond$ satisfies (U5). To prove that $(T \diamond \mu) \wedge \alpha$ implies $T \diamond (\mu \wedge \alpha)$, it is sufficient to prove that for each k-model of T, say M, $Res(M, \mu) \cap Mod(\alpha) \subseteq Res(M, \mu \wedge \alpha)$. In particular, we need to show for any $M' \in Res(M, \mu) \cap Mod(\alpha)$, $M' \in Res(M, \mu \wedge \alpha)$. Suppose $M' \notin Res(M, \mu \wedge \alpha)$. According Definition 3, we have (1) $M' \not\models \mu \wedge \alpha$; or (2) there exists another k-model M'' such that $M'' \models \mu \wedge \alpha$ and $M'' <_M M'$. If it is case (1), it follows that $M' \notin Res(M, \mu) \cap Mod(\alpha)$. Then the result holds. If it is case (2), it also implies that $M'' \models \mu$ and $M'' <_M M'$. That means, $M' \notin Res(M, \mu)$ from

Defintion 3. The result still holds.

Now we prove $\diamond$ satisfies (U6). Similarly, to prove $\diamond$ satisfies (U6), we only need to prove for any k-model of T, say M, if $Res(M, \mu_1) \subseteq Mod(\mu_2)$ and $Res(M, \mu_2) \subseteq Mod(\mu_1)$, then $Res(M, \mu_1) = Res(M, \mu_2)$. We first prove $Res(M, \mu_1) \subseteq Res(M, \mu_2)$. Let $M' \in Res(M, \mu_1)$. Then $M' \models \mu_2$. Suppose $M' \notin Res(M, \mu_2)$. This follows that there exists another $M'' \in Res(M, \mu_2)$ such that $M'' <_M M'$. Also note that $M'' \models \mu_1$. This contradicts the fact that $M' \notin Res(M, \mu_2)$. This proves $Res(M, \mu_1) \subseteq Res(M, \mu_2)$. Similarly, we can prove $Res(M, \mu_2) \subseteq Res(M, \mu_1)$.

Now we prove $\diamond$ satisfies (U7). Since T is complete, it follows that T has a unique k-model M. So we only need to prove $Res(M, \mu_1) \cap Res(M, \mu_2) \subseteq Res(M, \mu_1 \vee \mu_2)$. Let $M' \in Res(M, \mu_1) \cap Res(M, \mu_2)$. Suppose $M' \notin Res(M, \mu_1 \vee \mu_2)$. Then there exists a k-model $M'' \in Res(M, \mu_1 \vee \mu_2)$ such that $M'' <_M M'$. Note that $M'' \models \mu_1 \vee \mu_2$. If $M'' \models \mu_1$, it will follow that $M' \notin Res(Mu, \mu_1)$, otherwise, $M' \notin Res(Mu, \mu_2)$. In either case, we have $M' \notin Res(M, \mu_1) \cap Res(M, \mu_2)$. This proves the result.

Finally, the fact that $\diamond$ satisfies (U8) is obtained straightforward from Definitions 3 and 4. ■

4 Characterizing specific knowledge updates

While the previous section studies general minimal change properties of our knowledge update, alternative characterizations of knowledge update can be described for several specific forms. These specific forms present important features of knowledge update, and their alternative characterizations are convenient when the use of the notion of knowledge update becomes an overkill. For example, the alternative characterization of *sensing update* below is a much simpler characterization that is used in reasoning about sensing actions [22,23].

4.1 Gaining knowledge update

We first introduce a notation that will be useful in our following discussions. Let W be a set of worlds and $w \in W$. By $W^{(w,\phi)}$, we denote the set $\{w' \mid w' \in W \text{ and } w' \models \phi \text{ iff } w \models \phi\}$.

Proposition 3 *Given T and $K\phi$ where ϕ is objective and $T \models \phi$. Then $M' = (W', w')$ is a k-model of $T \diamond K\phi$ if and only if there exists a k-model*

$M = (W, w)$ of T such that $w = w'$ and $W' = W^{(w,\phi)}$.

Proof: Since $T \models \phi$, we have $W^{(w,\phi)} = \{w' \mid w' \in W \text{ and } w' \models \phi\}$. Let $M' = (W^{(w,\phi)}, w)$. Firstly, from the condition 1 of Definition 2, it is easy to see that for any $M'' = (W'', w'')$ where $w'' \neq w$, $M' <_M M''$. Therefore, from Theorem 1, to prove the result, it is sufficient to prove that for any k-model $M'' = (W'', w)$ where $W'' \neq W^{(w,\phi)}$, $M' \leq_M M''$. We consider the following possible cases.
<u>*Case 1*</u>. $W'' \subset W^{(w,\phi)}$. Since $W^{(w,\phi)} \subseteq W$, from Proposition 1, we have $KM \subseteq KM' \subseteq KM''$, and hence $KM' \setminus KM \subseteq KM'' \setminus KM$. From Definition 2, it follows $M' \leq_M M''$.
<u>*Case 2*</u>. $W^{(w,\phi)} \subset W''$ (proper set inclusion). Without loss of generality, we assume $W'' = W^{(w,\phi)} \cup \{w_i\}$, where $w_i \models \phi$. Clearly, $w_i \notin W$ otherwise we will have $w_i \in W^{(w,\phi)}$ and then $W'' = W^{(w,\phi)}$. Since $W'' \not\subseteq W$ and $W \not\subseteq W''$, from Proposition 1, we have $KM \not\subseteq KM''$ and $KM' \not\subseteq KM$. Then it must be the case that $KM \setminus KM'' \neq \emptyset$ and $KM'' \setminus KM \neq \emptyset$. From Definition 2 (i.e. (iii) in condition 2), we know that $M' \leq_M M''$.
<u>*Case 3*</u>. $W^{(w,\phi)} \not\subset W''$ and $W'' \not\subset W^{(w,\phi)}$. Without loss of generality, we can assume that $W'' = W^{(w,\phi)} \cup \{w_i\} \setminus \{w_j\}$, where $w_j \in W^{(w,\phi)}$. Since we require that $M'' \models K\phi$, it follows that $w_i \models \phi$. Also, from the construction of $W^{(w,\phi)}$, we know that $w_i \notin W$ otherwise it reduces to the case that $W'' \subseteq W^{(w,\phi)}$. Therefore, $W'' \not\subseteq W$ and $W \not\subseteq W''$. From the above discussion, it follows that $KM \setminus KM'' \neq \emptyset$ and $KM'' \setminus KM \neq \emptyset$. So from Definition 2, we know that $M' \leq_M M''$. ■

The above proposition reveals an important property about knowledge update: to know some fact, the agent only needs to restrict the current possible worlds in each of her k-models, if this fact itself is already entailed by her current knowledge set. We call this kind of knowledge update *gaining knowledge update*.

Example 2 *Let $T \equiv a \wedge \neg Ka \wedge Kb$. Suppose $w_0 = \{a, b\}$, $w_1 = \{a\}$, $w_2 = \{b\}$ and $w_3 = \emptyset$. Then T has one k-model $M = (\{w_0, w_2\}, w_0)$. Updating M with Ka, according to our k-model update definition, we have a unique resulting k-model $M' = (\{w_0\}, w_0)$. Indeed, this result is also obtained from Proposition 3.*

4.2 Ignorance update

As a contrary case to the gaining knowledge update, we now characterize an agent ignoring a fact from her knowledge set which we call *ignorance update*, i.e. updating T with $\neg K\phi$. From Definition 1, it is easy to see that $T \diamond \neg\phi \models \neg K\phi$. However, it should be noted that updating $T \diamond \neg\phi$ can *not* be used to achieve $T \diamond \neg K\phi$. Consider a k-model $M = (\{\{a,b\},\{a\}\},\{a,b\})$. Updating M with $\neg Ka$ we have a possible resulting k-model $M' = (\{\{a,b\},\{a\},\{b\}\},\{a,b\})$, while updating M with $\neg a$ will lead to a possible result $M'' = (\{\{a,b\},\{a\},\{b\}\},\{b\})$. Note that both M' and M'' entail $\neg Ka$, but $M' <_M M''$ according to Definition 2.

Proposition 4 *Given T and $\mu \equiv \neg K\phi$ where ϕ is objective. $M' = (W', w')$ is a k-model of $T \diamond \neg K\phi$ if and only if there exists a k-model $M = (W, w)$ of T such that*

(i) if $M \models K\phi$, then $w' = w$ and $W' = W \cup \{w^*\}$, where $w^* \models \neg\phi$;
(ii) otherwise, $w' = w$ and $W' = W$.

Proof: Let $\mu \equiv \neg K\phi$ and $M = (W, w)$ be a k-model of T. Then it is easy to see that for two k-models $M' = (W', w')$ and $M'' = (W'', w'')$ such that $M' \models \mu$ and $M'' \models \mu$, and $w' = w$ and $w'' \neq w$, $M' <_M M''$. So M'' can not be a k-model of $T \diamond \mu$. In other words, a k-model of $T \diamond \mu$ must have a form $M' = (W', w)$.

From Theorem 1, to prove the result, we only need to show that for any k-model $M'' = (W'', w)$ such that $M'' \models \mu$ and $W'' \neq W \cup \{w^*\}$ where $w^* \models \neg phi$, $M' \leq_M M''$.

Suppose $M \models K\phi$. Let $M' = (W \cup \{w^*\}, w)$, where $w^* \models \neg\phi$. We first show that for any k-model $M'' = (W'', w)$ such that $M'' \models \neg K\phi$ and W'' does not have a form of $W \cup \{w_i\}$, $M' \leq_M M''$.

<u>*Case 1*</u>. Suppose $W' \subset W''$. This implies that $KM'' \subseteq KM' \subseteq KM$ from Proposition 1. So $M' \leq_M M''$ according to Definition 2 (condition (b)).
<u>*Case 2*</u>. Suppose $W'' \subset W'$. Without loss of generality, we assume that $W'' = W \cup \{w^*\} \setminus \{w_j\}$ where $w_j \in W$. This follows that $W \not\subset W''$ and $W'' \not\subset W$. So it is the case that $KM \setminus KM'' \neq \emptyset$ and $KM'' \setminus KM \neq \emptyset$. On the other hand, we have $W \subseteq W'$, from Definition 2 (i.e. condition (a)), we have $M' \leq_M M''$.
<u>*Case 3*</u>. Suppose $W'' \not\subset W'$ and $W' \not\subset W''$. Without loss of generality, we can assume that $W'' = W \cup \{w^*, w_i\} \setminus \{w_j\}$, where $w_j \in W$ and $w^*, w_i \notin W$. Again, this results to the situation that $W \not\subset W''$ and $W'' \not\subset W$. From the above discussion, it implies that $M' \leq_M M''$.

Now we show that for any k-model M'' that is of the form $M'' = (W \cup \{w_i\}, w)$ and w_i is any world such that $w_i \models \neg\phi$ (note $M \models K\phi$), $M' \not\leq_M M''$ and $M'' \not\leq_M M'$. Suppose $M' \leq_M M''$. Since $W \subset W \cup \{w^*\}$, then according to Definition 2, condition (a) or (b) should be satisfied. As $W \subset W \cup \{w_i\}$, condition (a) can not be satisfied. So condition (b) must be satisfied. That is, for any ψ such that $M \models K\psi$ and $M' \not\models K\psi$, $M'' \models K\psi$. However, this implies that $KM \setminus KM' \subseteq KM \setminus KM''$, and also $KM \cap KM'' \subseteq KM \cap KM'$. From

Proposition 1 (Results 2 and 4), it follows that $W \cup W' \subseteq W \cup W''$, that is, $W \cup \{w^*\} \subseteq W \cup \{w_i\}$. Obviously, this is not true. Similarly, we can show that $M'' \not\leq_M M'$. That means, both M' and M'' are in $Res(M, \mu)$. This completes our proof. ■

Example 3 *Suppose* $T \equiv \neg Ka \wedge \neg Kb \wedge K(a \vee b) \wedge Kc$ *and the agent wants to ignore* c*. Let* $w_0 = \{a,b,c\}$*,* $w_1 = \{a,c\}$*,* $w_2 = \{b,c\}$*,* $w_3 = \{c\}$*,* $w_4 = \{a,b\}$*,* $w_5 = \{a\}$*,* $w_6 = \{b\}$*,* $w_7 = \emptyset$*. Clearly,* T *has three* k*-models:* $M_0 = (\{w_0, w_1, w_2\}, w_0)$*,* $M_1 = (\{w_0, w_1, w_2\}, w_1)$*, and* $M_2 = (\{w_0, w_1, w_2\}, w_2)$*. From Proposition 3,* $T \diamond \neg Kc$ *has the following twelve* k*-models:* $(\{w_0, w_1, w_2, w_i\}, w_j)$*, where* $i = 4, 5, 6, 7$ *and* $j = 0, 1, 2$*.*

4.3 Sensing update

Now we consider the case when μ is of the form $K\phi \vee K\neg\phi$ where ϕ is objective. Updating T with this type of μ is particularly useful in reasoning about sensing actions [22,23] where $K\phi \vee K\neg\phi$ represents the effect of a sensing action after its execution, the agent will know either ϕ or its negation. We refer to such an update as a *sensing update*. The following proposition characterizes the update of T with a formula of the form $K\phi \vee K\neg\phi$. It is interesting to note that the sufficient and necessary condition for a k-model of $T \diamond (K\phi \vee K\neg\phi)$ is similar to the one presented in Proposition 3.

Proposition 5 *Given* T *and* $\mu \equiv K\phi \vee K\neg\phi$ *where* ϕ *is objective.* $M' = (W', w')$ *is a* k*-model of* $T \diamond (K\phi \vee K\neg\phi)$ *if and only if there exists a* k*-model* $M = (W, w)$ *of* T *such that* $w = w'$ *and* $W' = W^{(w,\phi)}$ *or* $w = w'$ *and* $W' = W^{(w,\neg\phi)}$

Proof: Let $\mu \equiv K\phi \vee K\neg\phi$ and $M' = (W^{(w,\phi)}, w)$. From the definition of $W^{(w,\phi)}$, it is easy to see that $M' \models \mu$. Consider any $M'' = (W'', w'')$ where $M'' \models \mu$ and $w'' \neq w$. According to the condition 1 of Definition 2, $M' <_M M''$. So M'' can not be a k-model of $T \diamond \mu$. In other words, each k-model of $T \diamond \mu$ must have a form of $M'' = (W'', w)$. Then from Theorem 1, to prove the result, it is sufficient to prove for k-model $M'' = (W'', w)$ where $W'' \neq W^{(w,\phi)}$ or $W'' \neq W^{(w,\neg\phi)}$, $M' \leq_M M''$. This can be shown in the same way as in the proof of Proposition 3. ■

Example 4 *Suppose* $T \equiv Kb \wedge \neg Ka \wedge \neg K\neg a$ *represents the current knowledge of an agent. Note that* T *implies that the agent does not have any knowledge about* a*. Consider the update of* T *with* $\mu \equiv Ka \vee K\neg a$ *which can be thought of as the agent trying to reason – in the planning or plan verifica-*

tion stage – about a sensing action[4] *that will give her the knowledge about* a*. Let* $w_0 = \{a, b\}, w_1 = \{b\}, w_2 = \{a\}$ *and* $w_3 = \emptyset$*. It is easy to see that* $M_0 = (\{w_0, w_1\}, w_0)$ *and* $M_1 = (\{w_0, w_1\}, w_1)$ *are two* k*-models of* T*. Then according to the above proposition, it is obtained that* $M_0' = (\{w_0\}, w_0)$ *and* $M_1' = (\{w_1\}, w_1)$ *are the two* k*-models of* $T \diamond \mu$*.*

4.4 Forgetting update

As another important type of knowledge update, we consider the update of T with $\mu \equiv \neg K\phi \wedge \neg K\neg\phi$. This update can be thought of as the result of an agent *forgetting* her knowledge about the fact ϕ. We will refer to such an update as a *forgetting update*. The following proposition shows that in order to forget ϕ from T, for each k-model of the current knowledge set, the agent only needs to expand the set of possible worlds of this model with exactly *one specific* world.

Proposition 6 *Given* T *and* $\mu \equiv \neg K\phi \wedge \neg K\neg\phi$ *where* ϕ *is objective.* $M' = (W', w')$ *is a* k*-model of* $T \diamond \mu$ *if and only if there exists a* k*-model* $M = (W, w)$ *of* T *such that*

(i) if $M \models K\phi$*, then* $w' = w$ *and* $W' = W \cup \{w^*\}$*, where* $w^* \models \neg\phi$*;*
(ii) if $M \models K\neg\phi$*, then* $w' = w$ *and* $W' = W \cup \{w^*\}$*, where* $w^* \models \phi$*;*
(iii) otherwise, $w' = w$ *and* $W' = W$*.*

Proof: The proof is similar to the proof of Proposition 4. Let $\mu \equiv \neg K\phi \wedge \neg K\neg\phi$ and $M = (W, w)$ be a k-model of T. Then it is easy to see that for two k-models $M' = (W', w')$ and $M'' = (W'', w'')$ such that $M' \models \mu$ and $M'' \models \mu$, and $w' = w$ and $w'' \neq w$, $M' <_M M''$. So M'' can not be a k-model of $T \diamond \mu$. In other words, a k-model of $T \diamond \mu$ must have a form $M' = (W', w)$.

From Theorem 1, to prove the result, we only need to show that for any k-model $M'' = (W'', w)$ such that $M'' \models \mu$ and $W'' \neq W \cup \{w^*\}$, $M' \leq_M M''$.

Let $M' = (W \cup \{w^*\}, w)$, where $w^* \models \neg\phi$ if $M \models K\phi$ and $w^* \models \phi$ if $M \models K\neg\phi$. We first prove that for any k-model $M'' = (W'', w)$ such that $M'' \models \mu$ and W'' does not have a form of $W \cup \{w_i\}$, $M' \leq_M M''$.

Suppose $M \models K\phi$. Clearly $M' \models \mu$.
<u>*Case 1*</u>. Consider a k-model $M'' = (W'', w)$ where $W' = W \cup \{w^*\} \subset W''$. Note that $M'' \models \mu$ as well. However, from Proposition 1, we have $KM'' \subseteq$

[4] Such reasoning is necessary in creating plans with sensing actions or verifying such plans. On the other hand after the execution of a sensing action the agent exactly knows either a or $\neg a$, and can simply use the notion of belief update.

$KM' \subseteq KM$. So $M' \leq_M M''$ according to Definition 2 (i.e. condition (b)).
Case 2. Suppose $W'' \subset W'$. Without loss of generality, we assume that $W'' = W \cup \{w^*\} \setminus \{w_j\}$ where $w_j \in W$. This follows that $W \not\subset W''$ and $W'' \not\subset W$. So it is the case that $KM \setminus KM'' \neq \emptyset$ and $KM'' \setminus KM \neq \emptyset$. On the other hand, we have $W \subseteq W'$, from Definition 2 (i.e. condition (a)), we have $M' \leq_M M''$.
Case 3. Now suppose $W'' \not\subset W'$ and $W' \not\subset W'$. Without loss of generality, we can assume that $W'' = W \cup \{w^*, w_i\} \setminus \{w_j\}$, where $w_j \in W$ and $w^*, w_i \notin W$. Again, this results to the situation that $W \not\subset W''$ and $W'' \not\subset W$. From the above discussion, it implies that $M' \leq_M M''$.

Following the same way as above, we can prove that under the condition that $M \models K\neg\phi$ and $M' = (W \cup \{w^*\}, w)$ where $w^* \models \phi$, for any k-model $M'' = (W'', w)$ such that $M'' \models \mu$ and W'' does not have a form of $W \cup \{w_j\}$ $M' \leq_M M''$.

Now we show that for any k-model M'' that is of the form $M'' = (W \cup \{w_i\}, w)$ and w_i is any world such that $w_i \models \neg\phi$ if $M \models K\phi$ or $w_i \models \phi$ if $M \models K\neg\phi$, $M' \not\leq_M M''$ and $M'' \not\leq_M M'$. Suppose $M' \leq_M M''$. Since $W \subset W \cup \{w^*\}$, then according to Definition 2, condition (a) or (b) should be satisfied. As $W \subset W \cup \{w_i\}$, condition (a) can not be satisfied. So condition (b) must be satisfied. That is, for any ψ such that $M \models K\psi$ and $M' \not\models K\psi$, $M'' \models K\psi$. However, this implies that $KM \setminus KM' \subseteq KM \setminus KM''$, and also $KM \cap KM'' \subseteq KM \cap KM'$. From Proposition 1 (Results 2 and 4), it follows that $W \cup W' \subseteq W \cup W''$, that is, $W \cup \{w^*\} \subseteq W \cup \{w_i\}$. Obviously, this is not true. Similarly, we can show that $M'' \not\leq_M M'$. That means, both M' and M'' are in $Res(M, \mu)$. This completes our proof. ∎

Example 5 *Suppose $T \equiv Kb \wedge (Ka \vee K\neg a)$ represents the current knowledge of an agent. After executing a forgetting action the agent now would like to update her knowledge with $\mu \equiv \neg Ka \wedge \neg K\neg a$. Let $w_0 = \{a, b\}, w_1 = \{b\}, w_2 = \{a\}, w_3 = \emptyset$. It is easy to see that $M_0 = (\{w_0\}, w_0)$ and $M_1 = (\{w_1\}, w_1)$ are the two k-models of T. Then using Proposition 6, we conclude that $M'_0 = (\{w_0, w_1\}, w_0)$, $M'_1 = (\{w_0, w_3\}, w_0)$, $M'_2 = (\{w_1, w_0\}, w_1)$, and $M'_3 = (\{w_1, w_2\}, w_1)$ are the four k-models of $T \diamond \mu$. Note that $(\{w_0, w_2\}, w_0)$ cannot be a k-model of $T \diamond \mu$ according to Proposition 6.*

5 Persistence of knowledge and ignorance

Like most systems that do dynamic modeling, the knowledge update discussed previously is non-monotonic in the sense that while adding new knowledge into a knowledge set, some previous knowledge in the set might be lost. However, it is important to investigate classes of formulas that are persistent with respect to an update, as this may partially simplify the underlying inference problem [26]. Furthermore, characterizing persistence is also an important issue in non-monotonic epistemic logic reasoning because it plays an essential role in the way of how different states of agent's knowledge can be compared [3,10,11].

Given T and μ, a formula α is said to be *persistent* with respect to the update of T with μ, if $T \models \alpha$ implies $T \diamond \mu \models \alpha$. If α is of the form $K\phi$, we call this persistence as *knowledge persistence*, while if α is of the form $\neg K\phi$, we call it *ignorance persistence*. The question that we address now is that under what conditions, a formula α is persistent with respect to the update of T with μ.

As the update of T with μ is achieved based on the update of every k-model of T with μ, our task reduces to the study of persistence with respect to a k-model update. This is defined in the following definition.

Definition 5 *(***Persistence with respect to k-model update***) Let μ and α be two formulas and M be a k-model. α is* persistent *with respect to the update of M with μ if for any $M' \in Res(M, \mu)$, $M \models \alpha$ implies $M' \models \alpha$.*

Clearly a formula α is persistent with respect to the update of T with μ if and only if for each k-model M of T, α is persistent with respect to the update of M with μ. To characterize the persistence property with respect to k-model updates, we first define a preference ordering on k-models in terms of a formula.

Definition 6 *(***Formula closeness***) Let μ be a formula and M_1 and M_2 be two k-models. We say M_1 is* as close *to μ as M_2, denoted as $M_1 \leq_\mu M_2$, if one of the following conditions holds:*

(1) $M_1 \in Mod(\mu)$;
(2) if $M_1, M_2 \notin Mod(\mu)$, then for any $M \in Mod(\mu)$, $M_1 \leq_M M_2$.

We denote $M_1 <_\mu M_2$ if $M_1 \leq_\mu M_2$ and $M_2 \not\leq_\mu M_1$.

Intuitively, the above definition specifies a partial ordering to measure the closeness between two k-models to a formula. In particular, if M_1 is a k-model of μ, then M_1 is closer to μ than all other k-models (i.e. condition 1). If neither M_1 nor M_2 is a k-model of μ, then the comparison between M_1 and M_2 with respect to μ is defined based on the k-model preference ordering $\leq_M$ for each k-model M of μ (i.e. condition 2). Note that if both M_1 and M_2 are k-models of μ, we have $M_1 \leq_\mu M_2$ and $M_2 \leq_\mu M_1$, and both of them are equally close to μ.

Example 6 *Let $\mu \equiv Ka \wedge Kb$, $w_0 = \{a, b\}$, $w_1 = \{b\}$, $w_2 = \{a\}$ and $w_3 = \emptyset$. Clearly, μ has one k-model $M = (\{w_0\}, w_0)$. Consider two k-models $M_1 = (\{w_0, w_1\}, w_0)$ and $M_2 = (\{w_1, w_2\}, w_1)$. Now let us compare which one of* *them is closer to μ. Since neither M_1 nor M_2 is a k-model of μ, we can use condition 2 in Definition 6 to compare M_1 and M_2. According to Definition 2, it is easy to see that $M_1 \leq_M M_2$ as $w_0 \setminus w_1 \cup w_1 \setminus w_0 = \{a\} \neq \emptyset$. Therefore, we conclude $M_1 \leq_\mu M_2$. Furthermore, we also have $M_1 <_\mu M_2$.*

Proposition 7 *Let μ be a formula. For any two k-models M_1 and M_2, if $M_1 \leq_\mu M_2$, then $M_2 \models \mu$ implies $M_1 \models \mu$.*

Proof: Suppose $M_2 \models \mu$. Then $M_2 \in Mod(\mu)$. From Definition 6, we know that for any other k-model M', $M_2 \leq_\mu M'$. So $M_2 \leq_\mu M_1$. But we have $M_1 \leq_\mu M_2$. This implies that both M_1 and M_2 are equally close to μ. Hence, $M_1 \models \mu$. ■

Given a formula μ and a sequence of k-models $M_1, \cdots, M_k$, if the relation $M_1 \leq_\mu M_2 \leq_\mu \cdots \leq_\mu M_k$ holds, then it means that M_i is closer to μ than M_j, where $i < j$. Now under this condition, if there is another formula α which satisfies the property that $M_j \models \alpha$ implies $M_i \models \alpha$ whenever $i < j$, we say that formula α is *persistent* with respect to formula μ. In other words, when k-models move closer to μ, α's truth value is preserved in these k-models. The following definition formalizes this idea.

Definition 7 *(**$\leq_\mu$-persistence**) Let α, μ be two formulas. We say that α is* $\leq_\mu$-persistent *if for any two k-models M_1 and M_2, $M_2 \models \alpha$ and $M_1 \leq_\mu M_2$ implies $M_1 \models \alpha$.*

Now we have the following important relationship between $\leq_\mu$-persistence and k-model update persistence.

Theorem 3 *Let α and μ be two formulas and M be a k-model. α is persistent with respect to the update of M with μ if α is $\leq_\mu$-persistent.*

Proof: Let M' be a k-model in $Res(M, \mu)$. Then we have $M' \in Mod(\mu)$. So for any k-model M'', we have $M' \leq_\mu M''$. So $M' \leq_\mu M$. Now suppose α is μ-persistent. It follows that $M \models \alpha$ implies $M' \models \alpha$. As M' is an arbitrary k-model in $Res(M, \mu)$, we can conclude that α is persistent with respect to the update of M with μ. ■

From Theorem 2, we have that $\leq_\mu$-persistence is a sufficient condition to guarantee a formula persistence with respect to a k-model update. As will be shown next, we can provide a unique characterization for μ-persistence. We first define the notion of ordering preservation as follows.

Definition 8 *(**Ordering Preservation**) Given two formulas α and β. We say that ordering $\leq_\alpha$* preserves *ordering $\leq_\beta$ if for any two k-models M_1 and*

M_2, $M_1 \leq_\alpha M_2$ *implies* $M_1 \leq_\beta M_2$.

The intuition behind ordering preservation is clear. That is, if $\leq_\alpha$ preserves $\leq_\beta$, then for any two k-models M_1 and M_2, whenever M_1 is closer to α than M_2, M_1 will be closer to β than M_2 as well. Finally, we have the following important result to characterize μ-persistence.

Theorem 4 *Given two formulas α and μ, α is $\leq_\mu$-persistent if and only if $\leq_\mu$ preserves $\leq_\alpha$.*

Proof: ($\Rightarrow$) Suppose α is $\leq_\mu$-persistent. That is, for any two k-models M_1 and M_2, $M_1 \leq_\mu M_2$ and $M_2 \models \alpha$ implies $M_1 \models \alpha$. So under the constraint that α is μ-persistent, whenever $M_1 \leq_\mu M_2$, we have $M_1 \leq_\alpha M_2$. That means, $\leq_\mu$ preserves $\leq_\alpha$.

($\Leftarrow$) Suppose $\leq_\mu$ preserves $\leq_\alpha$. From Definition 8, we have that for any two k-models M_1 and M_2, $M_1 \leq_\mu M_2$ implies $M_1 \leq_\alpha M_2$. Now suppose $M_1 \leq_\mu M_2$. So we have $M_1 \leq_\alpha M_2$. From Proposition 6, we have that $M_2 \models \alpha$ implies $M_1 \models \alpha$. From this it follows that α is $\leq_\mu$-persistent. ■

6 Background on computational complexity

In the rest of this paper, we consider complexity issues of knowledge update. In particular, we investigate the computational complexity of model checking for knowledge update. For this purpose, we will restrict the underlying language to be finite.

We first introduce basic notions from complexity theory and refer to [7] for further details. Two important complexity classes are P and NP. The class of P includes those decision problems solvable by a polynomial-time deterministic Turing machine. The class of NP, on the other hand, consists of those decision problems solvable by a polynomial-time nondeterministic Turing machine.

Let $\mathcal{C}$ be a class of decision problems. The class $P^{\mathcal{C}}$ consists of the problems solvable by a polynomial-time deterministic Truing machine with an oracle for a problem from $\mathcal{C}$, while the class $NP^{\mathcal{C}}$ includes the problems solvable by a nondeterministic Turing machine with an oracle for a problem in $\mathcal{C}$. By co-$\mathcal{C}$ we mean the class consisting of the complements of the problems in $\mathcal{C}$.

The classes Σ_k^P and Π_k^P of the *polynomial hierarchy* are defined as follows:

$\Sigma_0^P = \Pi_0^P = P$, and
$\Sigma_k^P = NP^{\Sigma_{k-1}^P}$, $\Pi_k^P =$co-Σ_k^P for all $k > 1$.

It is easy to see that $NP = \Sigma_1^P$ and co-$NP = \Pi_1^P$. A problem A is *complete* for a class $\mathcal{C}$ if $A \in \mathcal{C}$ and for every problem B in $\mathcal{C}$ there is a polynomial transformation of B to A.

The prototypical Σ_k^P-complete and Π_k^P-complete problems are deciding the validity of quantified Boolean formulas (QBFs) of the form:

$$Q_1X_1Q_2X_2 \cdots Q_kX_kE, k \geq 1, \tag{1}$$

where E is a Boolean expression using propositional atoms over alphabets $X_1, X_2, \cdots$, and X_k, and the Q_i's are alternating qualifiers from $\{\forall, \exists\}$ ($1 \leq i \leq k$). If $Q_1 = \exists$, then deciding the validity of (1) is Σ_k^P-complete, while deciding the validity of (1) is Π_k^P-complete if $Q_1 = \forall$.

Let X and Y be two finite set of propositional atoms where X and Y have the same cardinality, i.e. $|X| = |Y|$. For convenience, we use notion $X \equiv Y$ to stand for formula $(x_1 \equiv y_1) \wedge (x_2 \equiv y_2) \wedge \cdots \wedge (x_m \equiv y_m)$. Consequently, $X \equiv \neg Y$ stands for formula $(x_1 \equiv \neg y_1) \wedge (x_2 \equiv \neg y_2) \wedge \cdots \wedge (x_m \equiv \neg y_m)$. We also use $\neg X$ to denote the set $\{\neg x_i \mid x_i \in X\}$ (or formula $\bigwedge_{x_i \in X} \neg x_i$), and use notion $\bigvee \neg X$ to stand for formula $\bigvee_{x_i \in X} \neg x_i$. For a given formula α, we use $|\alpha|$ to denote the length of α.

The problem of *model checking* for knowledge update is described as follows: Given a knowledge set T, a formula μ, and a k-model M, deciding whether $M \in Mod(T \diamond \mu)$. It is well known that the model checking problem for traditional belief revision and update is located at the lower end of the polynomial hierarchy from P to Σ_2^P depending on specific revision/update operators and additional restrictions (if any) [16].

7 Complexity of model checking: General case

In this section, we investigate the complexity of model checking for the general case of knowledge update.

Lemma 1 *Let $M = (w, W), M_1 = (w_1, W_1)$ and $M_2 = (w_2, W_2)$ be three k-models.*

(1) Deciding whether $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$ has time complexity $\mathcal{O}(|W| \times |W_2|)$.
(2) Deciding whether $KM \setminus KM_1 \subseteq KM \setminus KM_2$ has time complexity $\mathcal{O}(|W_1| \times (|W| + |W_2|))$.

(3) *Deciding whether* $KM_1 \backslash KM \subseteq KM_2 \backslash KM$ *has time complexity* $\mathcal{O}(|W_1| \times |W| \times |W_2|)$.

Proof: Result 1 is equivalently to deciding whether $W \not\subset W_2$ and $W_2 \not\subset W$ (proper set inclusion). Obviously, this can be verified in $\mathcal{O}(|W| \times |W_2|)$ time.

Now we prove Result 2. From set inclusion and intersection properties, it is easy to see that $KM \setminus KM_1 \subseteq KM \setminus KM_2$ iff $KM_2 \cap KM \subseteq KM_1 \cap KM$. Then from Proposition 1 (Results 2 and 4) of [1], it follows that $KM_2 \cap KM \subseteq KM_1 \cap KM$ iff $W_1 \cup W \subseteq W_2 \cup W$. Obviously, checking whether $W_1 \cup W \subseteq W_2 \cup W$ can be done in $\mathcal{O}(|W_1| \times (|W| + |W_2|))$ time.

Finally we prove Result 3. We first prove the following result:

> $KM_1 \not\subseteq KM_2 \cup KM$ if and only if for some $w \in W$ and $w_2 \in W_2$, such that $w_1 \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$ [5] for all $w_1 \in W_1$.

Firstly, if for all $w_1 \in W_1$ we have $w_1 \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$. Then let $\phi = \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$. Clearly, $\phi \in KM_1$, but $\phi \notin KM$ and $\phi \notin KM_2$. That is, $KM_1 \not\subseteq KM_2 \cup KM$

Now we suppose $KM_1 \not\subseteq KM_2 \cup KM$. Then there exists some $\phi \in KM_1$ but $\phi \notin KM$ and $\phi \notin KM_2$. We first show that in this case, $\phi \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$ for some $w \in W$ and $w_2 \in W_2$ respectively. Since $\phi \notin KM$, there exists some $w \in W$ such that $w \models \neg\phi$, and so $\phi \models \neg(\bigwedge w)$. Similarly, from $\phi \notin KM_2$, there exists some $w_2 \in W_2$ such that $w_2 \models \neg\phi$. That is $\phi \models \neg(\bigwedge w_2)$. Combine these two cases, we have $\phi \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$. On the other hand, since $\phi \in KM_1$, we have that for all $w_1 \in W_1$, $w_1 \models \phi$ and then $w_1 \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$.

From the above result, it observed that to decide whether $KM_1 \not\subseteq KM_2 \cup KM$, we only need to check whether there exist some $w \in W$ and $w_2 \in W_2$ such that for all $w_1 \in W_1$, $w_1 \models \neg(\bigwedge w) \wedge \neg(\bigwedge w_2)$. Obviously, we need to check at most $|W| \times |W_2|$ formulas for each $w_1 \in W_1$. So all checks can be done in time $\mathcal{O}(|W_1| \times |W| \times |W_2|)$. ■

Lemma 2 *Let* M, M_1 *and* M_2 *be three k-models. Deciding whether* $M_1 \leq_M M_2$ *can be achieved in polynomial time.*

Proof: According to Definition 2, if $w_1 \neq w_2$, then $M_1 \leq_M M_2$ iff $(w_1 \setminus w \cup w \setminus w_1) \subseteq (w_2 \setminus w \cup w \setminus w_2)$. Clearly, this can be verified in polynomial time.

[5] Note that we use notion $\bigwedge w$ to denote the conjunction of all propositional atoms that occur in w. If an atom is not in w, its negation will be in $\bigwedge w$. For instance, if $w = \{a, c\}$, then $\bigwedge w = a \wedge \neg b \wedge c$ considering that a, b and c are the only propositional atoms in the language.

If $w_1 = w_2$, then we need to check the following conditions: (i) If $W \subseteq W_1$, then $M_1 \leq_M M_2$ iff condition (a) or (b) in Definition 2 is satisfied. From Lemma 1, we know that deciding whether (a) and (b) to be true can be done in polynomial time. (ii) If $W_1 \subseteq W$, then $M_1 \leq_M M_2$ iff condition (a) or (c) in Definition 2 is satisfied. From Lemma 1, deciding whether (c) to be true is in P. (iii) If $W \not\subset W_1$ and $W_1 \not\subset W$, then $M_1 \leq_M M_2$ iff conditions (b) and (c) should be satisfied. Again, deciding whether condition (c) to be true is in polynomial time. So, the problem is in P. ■

Lemma 3 *Let M, M' be two k-models and μ a S5 formula. Deciding whether $M' \in Res(M, \mu)$ is in co-NP.*

Proof: According to Theorem 1, if $M' \notin Res(M, \mu)$, there must exist another k-model M'' such that $M'' <_M M'$. A guess of a k-model M'' can be done in polynomial time. From Lemma 2, deciding whether $M'' \leq_M M'$ is in P. Since $M'' <_M M'$ iff $M'' \leq_M M'$ and $M' \not\leq_M M''$, this follows that checking whether $M'' <_M M'$ can be decided in polynomial time. So the problem is in co-NP. ■

Theorem 5 *Model checking for knowledge update is in Σ_2^P.*

Proof: From Definition 4, $M \in Mod(T \diamond \mu)$ iff for some $M' \in Mod(T)$, $M \in Res(M', \mu)$. A guess of M' and check whether $M' \in Mod(T)$, i.e. $M' \models T$, can be achieved in polynomial time. According to Lemma 3, deciding whether $M \in Res(M', \mu)$ can be solved with one call to a co-NP oracle. So the problem is in Σ_2^P. ■

7.1 Knowledge gradual update

To prove the hardness, we consider a special form of knowledge update and prove its model checking complexity is Σ_2^P-hard.

Given T and μ, we say the update of T with μ is *knowledge gradual* if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that either $W \subseteq W'$ or $W' \subset W$. Note that, after performing a knowledge gradual update, the agent's knowledge may be decreased or increased (or without change), and the agent's actual world may be changed as well.

Example 7 *Let $T = a \wedge \neg Ka$ and $\mu = K\neg a$. Obviously, T has a unique k-model $M = (\{\{a\}, \emptyset\}, \{a\})$. Then updating M with μ generates a unique k-model of $T \diamond \mu$: $M' = (\{\emptyset\}, \emptyset)$. Obviously, M' has increased knowledge from M and the actual world of M' is also different from M's.*

Lemma 4 *Let $X, Y, \hat{X}, \hat{Y}$ be sets of propositional atoms and a be a propositional atom, where $|X| = |\hat{X}|$, $|Y| = |\hat{Y}|$ and any two sets of $X, Y, \hat{X}, \hat{Y}$ and $\{a\}$ are disjoint. Suppose ϕ is an objective formula formed based on $X \cup Y$. Let T and μ be the following two S5 formulas respectively:*

$$
\begin{aligned}
&T = \gamma_1 \vee \gamma_2, \text{ where} \\
&\gamma_1 = (X \equiv \hat{X}) \equiv \phi \equiv a \wedge (Y \wedge \neg\hat{Y}) \wedge \\
&\qquad \neg K \neg(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}), \\
&\gamma_2 = ((X \wedge \neg\hat{X}) \equiv \neg\phi \equiv \neg a) \wedge \hat{Y}, \text{ and} \\
&\mu = K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}) \vee \\
&\qquad K(\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}) \vee \\
&\qquad K(\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})^{6}.
\end{aligned}
$$

Then $T \diamond \mu$ is knowledge gradual.

Proof: To prove $T \diamond \mu$ to be knowledge gradual, we need to show that for any k-model $M = (W, w) \in Mod(T)$, if $M' = (W', w') \in Res(M, \mu)$, then either $W \subseteq W'$ or $W' \subseteq W$. From the construction of T, it is easy to see that if $M \in Mod(T)$, then either $M \models \gamma_1$ or $M \models \gamma_2$, but $M \not\models \gamma_1 \wedge \gamma_2$. Based on this observation, our proof consists of two cases.

Case 1. Let $M = (W, w) \in Mod(\gamma_1)$. Since $M \models \gamma_1$, we have

$w = X_1 \cup \hat{X}_1 \cup Y \cup \{a\}$, where $X_1 \cup Y \models \phi$ for some $X_1 \subseteq X$.

Furthermore, since $M \models \neg K \neg(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$, there exists a world $w^* \in W$ such that

$w^* = X \cup Y \cup \hat{Y} \cup \{a\}$.

Now we specify a k-model of μ as follows:

$M^* = (\{w^*\}, w^*)$.

It is easy to see that $M^* \models K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$. So $M^* \models \mu$. Furthermore, M^* is the unique k-model of $K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$. We prove M^* is the unique k-model in $Res(M, \mu)$.

Note $Diff(w, w^*) = (X \setminus X_1) \cup \hat{X}_1 \cup \hat{Y}$. Besides M^*, μ has other two types of k-models:

$M_1 = (W_1, w_1)$, where $w_1 = X \cup Y_1 \cup \hat{Y}$, where $Y_1 \subset Y$ $(Y_1 \neq Y)$, and
$M_2 = (W_2, w_2)$, where $w_2 = X \cup Y \cup \hat{Y}$.

Note that

---

[6] Recall that $\bigvee \neg Y = \bigvee_{y_i \in Y} \neg y_i$.

$M_1 \models K(\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y})$, and
$M_2 \models K(\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$.

Consider

$$Diff(w, w_1) = (X \setminus X_1) \cup \hat{X}_1 \cup (Y \setminus Y_1) \cup \hat{Y} \cup \{a\},$$
$$Diff(w, w_2) = (X \setminus X_1) \cup \hat{X}_1 \cup \hat{Y} \cup \{a\}.$$

Clearly, we have

$$Diff(w, w^*) \subset Diff(w, w_1), \text{ and}$$
$$Diff(w, w^*) \subset Diff(w, w_1).$$

So $Res(M, \mu) = \{M^*\}$. Also observe that
$M^* = (\{w^*\}, w^*)$, $\{w^*\} \subset W$.

<u>*Case 2*</u>. Let $M = (W, w) \in Mod(\gamma_2)$. We have

$$w = X \cup Y_1 \cup \hat{Y}, \text{ where } X \cup Y_1 \models \neg\phi \text{ for some } Y_1 \subseteq Y.$$

If $Y_1 \neq Y$, then we have

$$w \models \neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}),$$

this implies that there exists a subset W_1 of W where for each $w_i \in W_1$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

By specifying W_1 to be the maximal such subset of W, it is easy to note that $M_1 = (W_1, w)$ is a k-model in $Res(M, \mu)$.

On the other hand, if $Y_1 = Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}),$$

this implies that there exists a subset W_2 of W where for each $w_i \in W_2$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}).$$

Similarly, by specifying W_2 to be the maximal such subset of W, it is easy to note that $M_2 = (W_2, w)$ is a k-model in $Res(M, \mu)$.

Since in both cases, we have $W_1 \subseteq W$ and $W_2 \subseteq W$, this follows that for any k-model M of T where $M \models \gamma_2$, every resulting k-model after updating M with μ only increases the knowledge from M. ∎

Theorem 6 *Model checking for knowledge update is Σ_2^P-complete. The hardness holds even if the update is knowledge gradual.*

Proof: We only need to prove the hardness part. This part is based on a variation of the proof of Lemma 10. We prove the hardness by giving a polynomial transformation from deciding the validity of $\exists X \forall Y E$, where E is a Boolean expression using propositional atoms over $X \cup Y$. We construct T, μ and a k-model M^* over propositional atoms $X \cup Y \cup \hat{X} \cup \hat{Y} \cup \{a\}$, where $|\hat{X}| = |X|$ and $|\hat{Y}| = |Y|$, and any two sets among $X, Y, \hat{X}, \hat{Y}$ and $\{a\}$ are disjoint.

$$
\begin{aligned}
&T = \gamma_1 \vee \gamma_2, \text{ where} \\
&\gamma_1 = (X \equiv \hat{X}) \equiv E \equiv a \wedge (Y \wedge \neg \hat{Y}) \wedge \\
&\qquad \neg K \neg (a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}), \\
&\gamma_2 = ((X \wedge \neg \hat{X}) \equiv \neg E \equiv \neg a) \wedge \hat{Y}, \\
&\mu = K(a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}) \vee \\
&\qquad K(\neg a \wedge X \wedge \neg \hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}) \vee \\
&\qquad K(\neg a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}), \\
&M^* = (W^*, w^*), \text{ where} \\
&W^* = \{w^*\},\ w^* = X \cup Y \cup \hat{Y} \cup \{a\}.
\end{aligned}
$$

Note that

$$M^* \models K(a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}).$$

So M^* is a k-model of μ. Furthermore, it is the unique k-model of $K(a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y})$. From Lemma 10, we know that $T \diamond \mu$ is knowledge gradual. Now we will show that M^* is a k-model of $T \diamond \mu$ if and only if $\exists X \forall Y E$ is valid.

($\Rightarrow$) Suppose $\exists X \forall Y E$ is valid. Then for some $X_1 \subseteq X$, $X_1 \cup Y \models E$. We specify a k-model of γ_1 as follows:

$$M = (W, w), \text{ where } w = X_1 \cup \hat{X}_1 \cup Y \cup \{a\}.$$

Since $M \models \neg K \neg (a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y})$, it is clear that the world w^* must be in W, i.e. $w^* \in W$. With the same justification as described in the proof of Lemma 10, we conclude that M^* is a k-model of updating M with μ.

($\Leftarrow$) Suppose $\exists X \forall Y E$ is not valid. That is, $\forall X \exists Y \neg E$ is valid. Then $X \cup Y_1 \models \neg E$ for some $Y_1 \subseteq Y$. In this case, T has the following type of k-models:

$$M = (W, w), \text{ where } w = X \cup Y_1 \cup \hat{Y}.$$

Note that $w \models ((X \wedge \neg \hat{X} \equiv \neg E \equiv \neg a) \wedge \hat{Y})$. That is, M is a k-model of γ_2.

If $Y_1 \neq Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg \hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

We now specify a k-model of μ as follows: $M_1 = (W_1, w_1)$, where $w_1 = w$ and W_1 is the maximal subset of W such that for each $w_i \in W_1$,

$$w_i \models (\neg a \wedge X \wedge \neg \hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

Since

$$Diff(w, w^*) = (X \setminus X_1) \cup \hat{X_1} \cup \hat{Y}, \text{ and}$$
$$Diff(w, w_1) = \emptyset \subset Diff(w, w^*),$$

M^* is not a k-model in $Res(M, \mu)$.

If $Y_1 = Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}).$$

Again, we can specify a k-model of μ as follows: $M_2 = (W_2, w_2)$, where $w_2 = w$ and W_2 is maximal subset of W such that for each $w_i \in W_2$,

$$w_i \models (\neg a \wedge X \wedge \neg \hat{X} \wedge Y \wedge \hat{Y}).$$

Since $Diff(w, w_2) = \emptyset \subset Diff(w, w^*)$, M^* is not a k-model in $Res(M, \mu)$ in this case either.

Finally, suppose for some $X_1 \subseteq X$ and $Y_1 \subseteq Y$, E is evaluated to be true on $X_1 \cup Y_1$, i.e. $X_1 \cup Y_1 \models E$. Without loss of generality, we can assume $Y_1 \neq Y$ [7]. This implies that γ_1 does not have a k-model under this situation. Therefore, if $\exists X \forall Y E$ is not valid, all k-models of T must be k-models of γ_2. ■

8 A tractable subclass - knowledge decreased update

In this section, we identify a subclass of knowledge update problems for which model checking can be achieved in polynomial time. We first introduce a useful notation. Let α be a S5 formula and ϕ^α be an objective formula (i.e. no K occurs in it) occurring in α. We then say ϕ^α is an *objective sub-formula* of α. We denote the set of all objective sub-formulas of α as $Sub^o(\alpha)$. For instance, given $\alpha = Ka \vee K\neg b$, $Sub^o(\alpha) = \{a, b, \neg b\}$.

[7] Note that this assumption is always feasible. For instance, if $X_1 \cup Y \models E$, we can expand Y to be Y' by adding a new atom y' into Y to make $Y \neq Y'$, i.e. $Y' = Y \cup \{y'\}$, and modify E to be $E' = E \wedge \neg y'$ such that $X_1 \cup Y \models E'$ but $X_1 \cup Y' \not\models E'$.

Definition 9 *Given S5 formulas T and μ, updating T with μ is called* knowledge decreased *if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that (i) $W \subseteq W'$ and $w = w'$; and (ii) for any $w^* \in W'$, $w^* \in W$ iff either $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some ϕ^μ in $Sub^o(\mu)$.*

From the above definition, it is easy to see that if an update is knowledge decreased, then the actual world of the agent's state will not change, and the agent's knowledge can only be decreased. Furthermore, the set of posible worlds in the agent's resulting state can be specifically computed from her previous state. We have the following important result on the model checking for knowledge decreased update.

Theorem 7 *Model checking for knowledge decreased update can be achieved in polynomial time.*

Proof: Given T, μ and a k-mode $M' = (W', w')$. Suppose $T \diamond \mu$ be knowledge decreased. To check whether $M' \in Mod(T \diamond \mu)$, we need to do the following things:

(1) Check whether $M' \models \mu$,
(2) Compute a subset W of W' such that for any $w^* \in W'$, $w^* \in W$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some $\phi^\mu \in Sub^o(\mu)$,
(3) Check whether $(W, w) \models T$.

Clearly, Steps 1 and 3 can be done in polynomial time. As $|Sub^o(\mu)| \leq |\mu|$, it follows that Step 2 can be also done in polynomial time. ■

It is worthwhile to mention some concrete forms of knowledge decreased update which, as we have presented earlier, have important applications in practical domains.

Theorem 8 *Ignorance and forgetting updates are knowledge decreased.*

Proof: The proof directly follows from Propositions 4 and 6 respectively. ■

Corollary 1 *Model checking for ignorance and forgetting updates can be achieved in polynomial time.*

9 An intractable subclass - Knowledge increased update

In this section, we address another subclass of knowledge update problems whose model checking complexity are intractable but lower than the general case. Such investigation will be useful for us to design more optimal model checking algorithms for these subclasses of update problems.

As a contrary case to the knowledge decreased update, the knowledge increased update is defined as follows.

Definition 10 *Given T and μ, updating T with μ is called* knowledge increased *if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that (i) $W' \subset W$, and $w = w'$; and (ii) for any $w^* \in W$, $w^* \in W'$ iff either $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some ϕ^μ in $Sub^o(\mu)$.*

It is clear that if a knowledge increased update is performed to an agent's knowledge set, it only increases the agent's knowledge and does not change the agent's actual world. Unfortunately, different from the knowledge decreased update, the model checking problem for knowledge increased update is not tractable.

Theorem 9 *Model checking for knowledge increased update is NP-complete.*

Proof: Membership proof. Given T, μ and $M' = (W', w')$. To deciding whether $M' \in Mod(T \diamond \mu)$, we only need to show that for some $M \in Mod(T)$, $M' \in Res(M, \mu)$. A guess of $M = (W, w)$ and verifying $M \models T$ can be done in polynomial time. Since $T \diamond \mu$ is knowledge increased, to decide $M' \in Res(M, \mu)$, we only need to check: (1) $w = w'$, and (2) for any $w^* \in W$, $w^* \in W'$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some $\phi^\mu \in Sub^o(\mu)$. Obviously, both (1) and (2) can be checked in polynomial time. So the problem is in NP.

Hardness proof. The hardness is proved by transforming the NP-complete SAT problem to a gaining knowledge update that has been showed to be knowledge increased. Let E be a CNF on the set of propositional atoms X. We construct formulas T, μ and a k-model M' over two disjoint sets X and $\hat{X}$ where $|X| = |\hat{X}|$.

$T = (X \equiv \hat{X}) \wedge \neg K(X \equiv \hat{X})$,
$\mu \equiv K(X \equiv \hat{X} \vee \neg E)$, and
$M' = (W', w')$, where
$W' = \{w'\}$, $w' = X \cup \hat{X}$.

Clearly, $M \models \mu$. We will show that E is satisfiable iff $M' \in Mod(T \diamond \mu)$. Note that since $T \models X \equiv \hat{X} \vee \neg E$ and $\mu = K(X \equiv \hat{X} \vee \neg E)$, $T \diamond \mu$ is a gaining knowledge update that is knowledge increased according to Theorem 6.

($\Rightarrow$) Suppose E is satisfiable. Let $X_1 \subseteq X$ such that $X_1 \models E$. We specify a k-model as follows:

$M^* = (W^*, w^*)$,

$W^* = \{w^*, w''\}$,
$w^* = w' = X \cup \hat{X}$, and
$w'' = X_1 \cup \hat{X}_1$, where $\hat{X}_1 = \{\hat{x_i} \mid \hat{x_i} \in \hat{X} \text{ and } x_i \notin X_1\}$.

Since $w'' \not\models X \equiv \hat{X}$, it is easy to see that $M^* \models \neg K(X \equiv \hat{X})$. Therefore, M^* is a k-model of T. On the other hand, since $w'' \models E$ and $w'' \not\models X \equiv \hat{X}$, it follows that $W' = \{w'\} = \{w^*\} = W^{*(w^*,\phi)}$, where $\phi = (X \equiv \hat{X}) \vee \neg E$. From Lemma 6, $M' \in Res(M^*, \mu)$, so $M' \in Mod(T \diamond \mu)$.

($\Leftarrow$) Now suppose E is not satisfiable. That is, for any $X_1 \subseteq X$, $X_1 \models \neg E$. Then from Lemma 6, for any k-model of T of the form $M = (W, w)$, where $w \neq w'$, $M' \notin Res(M, \mu)$. We consider k-models of T of the form $M = (W, w)$ where $w = w'$ (note $w' \in W$). Without loss of generality, we assume that there is one world $w^* \in W$ such that $w^* \not\models X \equiv \hat{X}$, otherwise $M \models K(X \equiv \hat{X})$ and M cannot be a k-model of T. On the other hand, since E is not satisfiable, $\neg E$ must be true in each world in W. So $M \models K\neg E$ and hence $M \models K(X \equiv \hat{X} \vee \neg E)$. This implies that $W' \neq W^{(w,\phi)}$, where $\phi = (X \equiv \hat{X}) \vee \neg E$. So M' is not a k-model of $T \diamond \mu$. ■

It is interesting to note that some specific forms of knowledge update we discussed earlier are knowledge increased.

Theorem 10 *Gaining knowledge and sensing updates are knowledge increased.*

Proof: The proof directly follows from Propositions 3 and 5 respectively. ■

Corollary 2 *Model checking for gaining knowledge and sensing updates are NP-complete.*

10 Concluding Remarks

While research on reasoning about knowledge has made significant progress in the last decade, e.g. [4,8,14,17,21], the problem of modeling the dynamics of knowledge has only received attention in recent years and mainly been motivated from the study of belief revision and update. Van der Meyden recently studied the computational aspect of knowledge modeling in distributed systems [19] where the issue of knowledge update was discussed. Although van der Meyden showed that his knowledge update presented a generalization of certain aspects of standard knowledge base update, he did not explore knowledge update from a more semantical perspective because the notion of knowledge update was only used for the purpose of efficiently implementing model checking. On the other hand, the dynamic semantics for epistemic logic was considered by Groeneveld recently. In [6], updates on Kripke models were specified. In that update semantics, an update is defined as *eliminative* or *conscious* and knowledge is represented in a many-order setting under the possible worlds semantics. Under the eliminative update notion, an agent changes her knowledge *minimally* according to the new knowledge, while under the conscious update notion, an agent not only changes her knowledge by combining the new knowledge into the state, but also reflects the new knowledge at a higher-order level. While Groeneveld's work provides an initial account of knowledge update, it, however, did not examine its minimal change characterization in detail and its relationship to the traditional belief update (i.e. Whether Katsuno and Mendelzon's update postulates were satisfied). Furthermore, its other semantic and computational properties also remain unclear.

In this paper we developed an explicit notion of knowledge update as an analogous notion to belief update and illustrated its usefulness in characterizing the knowledge change of an agent in presence of new knowledge. In our formulation, knowledge update is particularly relevant in reasoning about actions and plan verifications when there are sensing or forgetting actions. We presented simpler alternative characterization of knowledge update for particular cases, and showed its equivalence to the original characterization. We discussed when particular knowledge (or ignorance) persists with respect to a knowledge update. We also undertook a further study about the complexity issue of knowledge update. In particular, we analyzed the complexity of model checking for knowledge update in the general case and in special cases. We identify special subcases where the model checking is either tractable or its complexity is lower than the general case. We expect that these results would be useful for designing more optimal model checking algorithms in the implementation of knowledge update.

We believe our work here to be a starting point on knowledge update, and as evident from the research in belief update and revision in the past decade. A lot remains to be done in knowledge update. For example, issues such as multiagent knowledge update, iterative knowledge update, abductive knowledge update, minimal knowledge in knowledge update, etc. remain to be explored. Similarly, in regards to reasoning about actions, additional specific cases of knowledge update need to be identified and simpler alternative characterization for them will be needed to be developed.

Acknowledgement

The authors thank Norman Foo and Abhaya Nayak for useful discussions on this topic. The research of the firat author was supported by NSF under grants IRI 9501577 and 0070463, and part of this research was performed while he was visiting the University of Western Sydney. The research of the second author was supported in part by Australian Research Council under grant A49803542.

References

[1] C. Baral and Y. Zhang, On the semantics of knowledge update. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pp 97-102. Morgan Kaufmann Publishers, Inc., 2001.

[2] C. Baral and Y. Zhang, The complexity of model checking for knowledge update. In *Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning (KR-02)*, pp??, Morgan Kaufmann Publishers, Inc., 2002.

[3] J Engelfreit. Monotonicity and persistence in preferential logics. *Journal of Artificial Intelligence Research*, 8:1-12, 1998.

[4] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge.* MIT press, 1995.

[5] P. Gardenfors, Knowledge in Flux. MIT Press, 1988.

[6] W. Groeneveld, *Logical Investigations into Dynamic Semantics*, PhD Thesis, Universiteit van Amsterdam, 1995.

[7] M. Garey and D. Johnson, *Computers and Intractability*, Freeman, San Franciso, 1979.

[8] J.Y. Halpern and Y. Moses, A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence* **54** (1992) 311-379.

[9] A. Herzigi and O. Rifi, Propositional belief base update and minimal change. *Artificial Intelligence* **115** 107-138, 1999.

[10] W. van der Hoek, J. Jaspars and E. Thijsse, Persistence and minimality in epistemic logic. *Annals of Mathematics and Artificial Intelligence* **27** (1999) 25-47.

[11] W. van der Hoek and E. Thijsse, A general approach to multi-agent minimal knowledge. *Studia Logica*, 2002 (to appear).

[12] G. Kartha and V. Lifschitz. Actions with indirect effects: Preliminary report. In *KR 94*, pages 341–350, 1994.

[13] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. of KR 91*, pages 387–394, 1991.

[14] H. Levesque, All I know: A study in autoepistemic logic. *Artificial Intelligence* **42** (1990) 263-309.

[15] P. Liberatore, The complexity of belief update. *Artificial Intelligence* **119** (2000) 141-190.

[16] P. Liberatore and M. Schaerf, Belief revision and update: complexity of model checking. *Journal of Computer and System Sciences* **62**, 43-72 (2001).

[17] A. Lomuscio and M. Ryan, Ideal agents sharing (some!) knowledge. In *Proceedings of ECAI-98*, pp 557-561, 1998.

[18] P. Marquis, Possible models approach via independency. In *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI'94)*, pp 336-340, 1994.

[19] R. van der Meyden, Common knowledge and update in finite environment. *Information and Computation* **140** (1998) 115-157.

[20] J-J. Meyer and W. van der Hoek. *Epistemic logics for computer science and artificial intelligence.* Cambridge Tracts in Theoretical Computer Science, 41, Cambridge University Press, 1995.

[21] G. Schwarz and M. Truszczyński, Minimal knowledge problem: A new approach. *Artificial Intelligence* **67** (1994) 193-215.

[22] R. Scherl and H. Levesque. The frame problem and knowledge producing actions. In *AAAI 93*, pages 689-695, 1993.

[23] T. Son and C. Baral. Formalizing sensing actions: a transition function based approach. *Artificial Intelligence* **125**, 19-91, 2001.

[24] H. Turner. Representing actions in logic programs and default theories. *Journal of Logic Programming* **31(1-3)**, 245-298, 1997.

[25] M. Winslett. Reasoning about action using a possible models approach. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-88)*, pages 89-93, 1988.

[26] Y. Zhang. Monotonicity in rule based update. In *Proceedings of the 1999 International Conference on Logic Programming (ICLP'99)*, pages 471-485. MIT Press, 1999.

[27] Y. Zhang and N.Y. Foo, Updates with disjunctive information: From syntactical abd semantical perspectives. *Computational Intelligence* **16** (2000) 29-52.

[28] Y. Zhang. The complexity of logic program updates. In *Proceedings of Australian Joint Conference on Artificial Intelligence (AI-2001). Lecture Notes in Artificial Intelligence*, Vol. 2256, pp 631-642, 2001.