

On the Size Complexity and Decidability of First-Order Progression*

Jens Claßen^{1†}, Daxin Liu^{2†}

¹Department of People and Technology, Roskilde University, Denmark

²State Key Laboratory for Novel Software Technology, Nanjing University, China
classen@ruc.dk, daxin.liu@nju.edu.cn

Abstract

Progression, the task of updating a knowledge base to reflect action effects, generally requires second-order logic. Identifying first-order special cases, by restricting either the knowledge base or action effects, has long been a central topic in reasoning about actions. It is known that local-effect, normal, and acyclic actions, three increasingly expressive classes, admit first-order progression. However, a systematic analysis of the size of such progressions, crucial for practical applications, has been missing. In this paper, using the framework of Situation Calculus, we show that under reasonable assumptions, first-order progression for these action classes grows only polynomially. Moreover, we show that when the KB belongs to decidable fragments such as two-variable first-order logic or universal theories with constants, the progression remains within the same fragment, ensuring decidability and practical applicability.

1 Introduction

Our motivation is derived from the aim of enabling intelligent agents to act in real-world scenarios, typically involving incomplete information about the state of the world and unbounded domains of objects. For this purpose, we use the framework of Situation Calculus [McCarthy and Hayes, 1969; Reiter, 2001], a widely studied first-order (FO) formalism for reasoning about action and change. A central problem for various reasoning tasks such as planning, verification, and synthesis, is that of *projection*, i.e., to determine if a query formula ϕ will hold after the execution of an action α , given a theory $\mathcal{D}$ consisting of a knowledge base (KB) describing the current state of affairs and axioms encoding the effects of actions. The two most common approaches are those of *regressing* ϕ through α and then testing the result against the original KB, and that of *progressing* the KB through α and then testing the original ϕ against this updated KB. Arguably, regression is less suited for practical applications because it

needs to be performed for every query formula anew, while it is sufficient to determine a progression once for each action.

We thus focus our attention on progression in this paper. Unfortunately, there is a strong negative result due to Lin and Reiter [1997], which says that representing the progression of an FO theory requires second-order (SO) logic in general. Identifying FO special cases, by restricting the KB or action effects, has therefore long been a central topic in reasoning about actions and change. Three increasingly expressive¹ such classes are the local-effect, normal [Liu and Lakemeyer, 2009], and acyclic [Liu and Claßen, 2024] ones, all of which were shown to be effectively FO progressable.

However, for practical feasibility, at least two additional aspects are crucial that we address in this paper. First, a formal analysis of the size complexity of such progressions has been missing. After formal preliminaries in Sec. 2, we show in Sec. 3 that under certain assumptions, the FO progression for the aforementioned classes grows only polynomially, or even linearly. Besides, we want to guarantee that queries can be effectively evaluated against the progression result, for instance by resorting to decidable fragments of first-order logic (FOL) such as the two-variable fragment FO^2 [Grädel *et al.*, 1997] or universal theories with constants UTC [Arenas *et al.*, 2018]. We hence show these fragments to be closed under progression in case of FO^2 (Sec. 4) and UTC (Sec. 5). We discuss related work and conclude in Sec. 6.

2 Preliminaries

We assume familiarity with FOL including equality, here denoted by $\mathcal{L}$. We will use “dot” notation to indicate that the quantifier preceding the dot has maximum scope, e.g., $\forall x.\phi(x) \supset \psi(x)$ stands for $\forall x(\phi(x) \supset \psi(x))$. We omit leading universal quantifiers and assume free variables are implicitly $\forall$ -quantified, e.g., we identify $\phi(x)$ with $\forall x.\phi(x)$. We use $\psi \Leftrightarrow \phi$ to mean ϕ and ψ are logically equivalent. We denote by $\phi(\mu/\mu')$ the formula obtained by simultaneously substituting every occurrence of expression (term or formula) μ in ϕ with μ' . Let $\mathcal{L}^2$ denote the SO extension of $\mathcal{L}$.

2.1 Situation Calculus

The Situation Calculus [Reiter, 2001] $\mathcal{L}_{sc}$ is a three-sorted language with some SO features suitable for describing dy-

*Extended version of identically-titled paper at IJCAI 2026.

[†]Corresponding Author.

¹The inclusion is not exact, as we discuss in Sec. 3.

dynamic worlds. The sorts are used to distinguish *action*, *situation*, and *object*. $\mathcal{L}_{sc}$ uses a distinct constant S_0 to denote the initial situation; a binary function $do(a, s)$ to generate the successor situation of situation s from doing action a ; a binary relation $Poss(a, s)$ to express action a being executable in situation s . We use a and s for variables of sort action and situation, respectively. Given a ground action α , we use S_α to refer to the situation $do(\alpha, S_0)$. *Fluents* are predicates whose last argument is a situation term; we omit function fluents for simplicity. A (possibly SO) formula ϕ is *uniform* in a situation term s if ϕ mentions no situation terms except s , quantifies no situation variables, and mentions no $Poss$.

Definition 1 (Basic Action Theory). A *basic action theory* (BAT) in $\mathcal{L}_{sc}$ is a set of axioms

$$\mathcal{D} = \Sigma_{ind} \cup \mathcal{D}_{ap} \cup \mathcal{D}_{ss} \cup \mathcal{D}_{una} \cup \mathcal{D}_{S_0}, \text{ where}$$

1. Σ_{ind} is a set of domain-independent axioms [Levesque et al., 1998] that ensure situations are well-structured;
2. $\mathcal{D}_{ap}$ is a set of action precondition axioms;
3. $\mathcal{D}_{ss}$ is a set of successor state axioms (SSAs), one for each fluent predicate F , of the form
$$F(\vec{x}, do(a, s)) \equiv \gamma_F^+(\vec{x}, a, s) \vee \neg \gamma_F^-(\vec{x}, a, s) \wedge F(\vec{x}, s),$$
where γ_F^+ and γ_F^- , the conditions under which action a causes F to become true and false, respectively, have to be uniform in s and satisfy $\mathcal{D} \models \neg(\gamma_F^+ \wedge \gamma_F^-)$;
4. $\mathcal{D}_{una}$ is the set of unique names axioms for actions: $A(\vec{x}) \neq A'(\vec{y})$, and $A(\vec{x}) = A(\vec{y}) \supset \vec{x} = \vec{y}$, where A, A' are distinct action symbols;
5. $\mathcal{D}_{S_0}$, the initial database (or initial KB), is a finite set of sentences uniform in S_0 .

2.2 Progression

A progression should retain all proper information, i.e., logical entailment in terms of the future of the initial KB. Lin and Reiter [1997] provide a model-theoretical definition of progression and prove that progression is SO definable. Vassos and Levesque [2013] prove that any correct progression is equivalent to the SO representation of [Lin and Reiter, 1997]. Here we use the definition from [Vassos and Levesque, 2013].

Let $\mathcal{D}_{ss}[\alpha, S_0]$ be the instantiation of $\mathcal{D}_{ss}$ wrt α and S_0 , i.e. $\mathcal{D}_{ss}[\alpha, S_0] \Leftrightarrow \bigwedge_F \forall \vec{x}. F(\vec{x}, do(\alpha, S_0)) \equiv \Phi_F(\vec{x}, \alpha, S_0)$, where Φ_F denotes the RHS of the SSA for fluent F . Let $F_1, \dots, F_n$ be the set of all fluents. For each fluent F_i , we introduce a new predicate symbol P_i . Let $\phi \uparrow S_0$ be the result of replacing every $F_i(\vec{t}, S_0)$ in ϕ by $P_i(\vec{t})$ and call P_i the *lifting predicate* for F_i . For a finite set of formulas Σ , we also use Σ to denote the conjunctions of its elements.

Definition 2. Let $\mathcal{D}$ be a basic action theory, α an executable ground action (under $\mathcal{D}_{ap}$), and $\mathcal{D}_{S_\alpha}$ a set of (FO or SO) sentences uniform in S_α . We say that $\mathcal{D}_{S_\alpha}$ is a *progression* of $\mathcal{D}_{S_0}$ wrt $\alpha, \mathcal{D}$ iff $\mathcal{D}_{una} \cup \mathcal{D}_{S_\alpha} \Leftrightarrow \mathcal{D}_{una} \wedge$

$$\exists \vec{R}. \{(\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]) \uparrow S_0\}(\vec{P}/\vec{R}) \quad (1)$$

where $\vec{R} = \{R_1, \dots, R_n\}$ are SO predicate variables.

Def. 2 is correct in the sense that for any ϕ uniform in S_α , if $\mathcal{D}_{S_\alpha}$ is a progression of $\mathcal{D}_{S_0}$ wrt $\alpha, \mathcal{D}$, then $\mathcal{D} \models \phi$ iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{D}_{S_\alpha} \models \phi$ [Vassos and Levesque, 2013, Thm. 1].

2.3 Forgetting

There are two approaches to identifying fragments of the Situation Calculus where progression is FO definable: one that restricts the structure of the initial theory, and one that restricts action effects. We focus on the latter in this paper, in particular, we consider *local-effect action theories* (BAT-LE), *normal action theories* (BAT-NR), and *acyclic action theories* (BAT-AC). The next sections give definitions for these classes, presents the associated progression methods, and provides size complexity results as new contribution.

The progression of a theory through an action can be interpreted in terms of *forgetting* all information about the original state of the world, and only keeping information relevant to the new state. Intuitively, forgetting a ground atom (or predicate) in a theory leads to a weaker theory that entails the same set of sentences that are “irrelevant” to the atom (or predicate). For a sentence ϕ and ground atom $P(\vec{t})$, let $\phi[P(\vec{t})]$ be the formula obtained by replacing every occurrence of the form $P(\vec{t})$ in ϕ with $[\vec{t} = \vec{t}' \wedge P(\vec{t})] \vee [\vec{t} \neq \vec{t}' \wedge \neg P(\vec{t})]$. Clearly, $\phi \Leftrightarrow \phi[P(\vec{t})]$. Let $\phi_+^{P(\vec{t})}$ and $\phi_-^{P(\vec{t})}$ be formulas obtained by replacing $P(\vec{t})$ in $\phi[P(\vec{t})]$ with TRUE and FALSE, respectively.

Theorem 3 ([Lin and Reiter, 1994]). Let $P(\vec{t})$ be a ground atom, P a predicate symbol, and ϕ a sentence. Then

- $forget(\phi, P(\vec{t})) \Leftrightarrow \phi_+^{P(\vec{t})} \vee \phi_-^{P(\vec{t})}$;
- $forget(\phi, P) \Leftrightarrow \exists R. \phi(P/R)$, where R is a SO variable.

Consequently, progression in Eq. (1) amounts to adding the effects of the action ($\mathcal{D}_{ss}[\alpha, S_0]$) and forgetting the truth of fluents at situation S_0 ($\exists$ -quantified from the outside). Thus, progression is FO definable in cases where the SO quantifiers in (1) can be eliminated.

W.l.o.g., we assume $\gamma_F^\pm(\vec{x}, a, s)$ in SSAs are disjunctions of formulas of the form (one for each action symbol A)

$$\exists \vec{z}. (a = A(\vec{v}) \wedge \phi_A^\pm(\vec{x}, \vec{z}, s))$$

where $A(\vec{v})$ is an action term, $\vec{v}$ may contain some variables from $\vec{x}$, and the remaining variables are $\vec{z}$; $\phi_A^\pm(\vec{x}, \vec{z}, s)$ is a formula uniform in s with no action terms, and where all free variables symbols are among $\vec{x}, \vec{z}$, and s . Under this assumption and $\mathcal{D}_{una}$, for every ground action $\alpha = A(\vec{t})$,

$$\gamma_F^\pm(\vec{x}, \alpha, S_0) \Leftrightarrow \exists \vec{z}. (\vec{t} = \vec{v} \wedge \phi_A^\pm(\vec{x}, \vec{z}, S_0)). \quad (2)$$

Henceforth, whenever we refer to $\gamma_F^\pm(\vec{x}, \alpha, S_0)$, we mean the RHS of Eq. (2) and this applies to $\mathcal{D}_{ss}[\alpha, S_0]$ as well.

3 Size Complexity of Progression

Previous work on the FO progression of acyclic action theories [Liu and Claßen, 2024] identified some complexity results in terms of computation time. However, the resulting theory’s size is equally crucial, since it directly affects the computational costs of checking queries against the theory. Here, we identify upper bounds for the size of FO progressions for BAT-LE, BAT-NR and BAT-AC theories.

Formally, the size $|\phi|$ of a formula ϕ is the number of atoms (including equalities), boolean connectives from $\{\wedge, \vee, \neg\}$,

and quantifiers from $\{\exists, \forall\}$ it contains. As mentioned, a formula with free variables is understood as its universal closure. Where possible, leading universal quantifiers are omitted, and do not count towards the size. Also, as usual, we identify finite theories with the conjunction of formulas they contain. To avoid confusion, when we talk about the size of such a finite theory, we always mean the size of the formula it presents, and *not* the number of formulas in the set.

The results in this section will present bounds for the size of the progression, which is understood as the set of sentences $\mathcal{D}_{S_\alpha}$ that replaces $\mathcal{D}_{S_0}$ in the BAT to reflect the state of the world in the new situation after executing ground action α . In particular (cf. Def. 2), the axioms encoding unique names for actions, $\mathcal{D}_{una}$, are not considered part of the progression, but may be used during the construction of a progression for simplification purposes. We further assume that the vocabulary of action and fluent symbols is fixed and finite.

3.1 BAT-LE

The idea behind local-effect action theories [Liu and Lakemeyer, 2009] is to impose constraints on SSAs such that any ground action only affects finitely many objects, namely those explicitly mentioned in the action's arguments. In consequence, only finitely many fluent instances are affected.

Formally, a ground action $\alpha = A(\vec{t})$ is said to have local effects on a fluent F if $\gamma_F^\pm(\vec{x}, \alpha, S_0)$ in Eq. (2) is of the form $\exists \vec{z}. [\vec{t} = \vec{v} \wedge \phi_A^\pm(\vec{z}, S_0)]$. We call a BAT a BAT-LE (wrt α) if α is local-effect to all fluents. For such a BAT, $\mathcal{D}_{una}$ implies $\gamma_F^\pm(\vec{x}, \alpha, S_0) \equiv \bigvee_i (\vec{x} = \vec{t}_i \wedge \phi_A^\pm(\vec{x}, S_0))$ (as all possible occurrences of $\vec{x}$ on the RHS of Eq. (2) are among $\vec{v}$). Hence, there are only finitely many affected fluent instances, i.e., the *characteristic set* $\Omega(s)$ given by

$$\{F(\vec{t}, s) \mid (\vec{x} = \vec{t} \text{ appears in } \gamma_F^+(\vec{x}, \alpha, s) \text{ or } \gamma_F^-(\vec{x}, \alpha, s))\}.$$

Let $\mathcal{D}_{ss}[\Omega]$ then be the set of instantiations of $\mathcal{D}_{ss}$ wrt Ω :

$$\mathcal{D}_{ss}[\Omega] = \{F(\vec{t}, S_\alpha) \equiv \Phi_F(\vec{t}, \alpha, S_0) \mid F(\vec{t}, s) \in \Omega(s)\},$$

where $\Phi_F(\vec{t}, \alpha, S_0)$ is the right-hand side of the SSA for F . Liu and Lakemeyer [2009] observe that for BAT-LE, forgetting predicates in the progression as in Eq. (1) can be done via forgetting the ground atoms in Ω : If θ is a consistent, finite set of ground literals, let $\phi[\theta]$ denote the result of replacing every occurrence of an atom $P(\vec{t})$ in ϕ by the formula

$$\bigvee_{j=1}^k (\vec{t} = \vec{t}_j \wedge v_j) \vee \left(\bigwedge_{j=1}^k \vec{t} \neq \vec{t}_j \right) \wedge P(\vec{t}), \quad (3)$$

where for $j = 1, \dots, k$, θ specifies truth value v_j for $P(\vec{t}_j)$. Theorem 3.6 of [Liu and Lakemeyer, 2009] states that the following is then a progression of $\mathcal{D}_{S_0}$ wrt α , $\mathcal{D}$:

$$\bigwedge \mathcal{D}_{una} \wedge \bigvee_{\theta \in \mathcal{M}(\Omega(S_0))} (\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\Omega])[\theta](S_0/S_\alpha) \quad (4)$$

Here, $\mathcal{M}(\Omega(S_0))$ denotes the set of all maximal consistent sets of literals over the atoms in $\Omega(S_0)$. The big disjunction in Eq. (4) essentially amounts to *forget*($\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\Omega], \Omega(S_0)$) (which means iteratively forgetting atoms in $\Omega(S_0)$). The claim in the following theorem was already stated in [Liu and Lakemeyer, 2009], however, without a proof.

Theorem 4. *If $\mathcal{D}$ is BAT-LE wrt ground action α , then its progression can be represented by a FO theory of size $O(2^c(n + m))$, where c is the size of the characteristic set, n is the size of $\mathcal{D}_{S_0}$, and m is the size of $\mathcal{D}_{ss}$.*

Proof. The size of $\mathcal{D}_{S_0}$ is n by assumption, and the size of $\mathcal{D}_{ss}[\Omega]$ is bounded by $cm + c - 1$, since there are c many formulas in $\mathcal{D}_{ss}[\Omega]$, each of size less or equal to m , with $(c - 1)$ conjunction symbols between them. If ϕ denotes $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\Omega]$, we hence have $|\phi| \leq n + cm + c$.

Observe that for any $\theta \in \mathcal{M}(\Omega(S_0))$, the number of elements in θ is the same as the number of elements in Ω , i.e., c . In particular, in (3), $k \leq c$. If a denotes the maximal arity of any predicate P in ϕ , $\vec{t} = \vec{t}_j$ abbreviates a formula of size less than $2a$, and so the size of the formula in (3) is bounded by $4ca + 4c + 3$. Therefore, for any ϕ , $|\phi[\theta]| \leq (4ca + 4c + 3)|\phi|$.

Finally, observe that the step (S_0/S_α) substitutes one constant by another, and hence does not change the size of the formula. Also, there are 2^c many θ 's in $\mathcal{M}(\Omega(S_0))$. As argued above, the size of $\mathcal{D}_{una}$ can be ignored. In summary, we can estimate the size of the progression to be no more than $2^c \cdot (n + cm + c) \cdot (4ca + 4c + 3)$. Since a is constant under the assumption of a fixed vocabulary, we get that the size of the progression is in $O(2^c(n + m))$. $\square$

We remark that c can be bounded by a constant. Note that the elements of $\Omega[S_0]$ are of the form $F(\vec{t}, S_0)$, where F is a local-effect fluent, and the $\vec{t}$ have to be chosen from the arguments of ground action α . Therefore, if l is the number of local-effect fluents, b is the arity of α , and a is again the maximal arity among predicates, $c \leq l \cdot b^a$, where all of l , a , and b are constant under the assumption of a fixed vocabulary. It can hence also be argued that the size of the progression of a BAT-LE is linear in the size of the original theory.

3.2 BAT-NR

To facilitate global effects, BAT-NR relaxes the restriction on SSAs but imposes additional assumptions on $\mathcal{D}_{S_0}$ so that one can use a variant Ackermann's Lemma [Ackermann, 1935] to eliminate SO quantifiers. Formally, a finite theory T is *semi-definitional* wrt a predicate P if the only occurrence of P in T is of the form $P(\vec{x}) \supset \phi(\vec{x})$ or $\psi(\vec{x}) \supset P(\vec{x})$. ϕ (respectively ψ) is called a necessary (sufficient) condition of P . Let WSC_P be the set formulas $\psi(\vec{x})$ such that $\psi(\vec{x}) \supset P(\vec{x})$ is in T , and SNC_P be the set of formulas $\phi(\vec{x})$ with $P(\vec{x}) \supset \phi(\vec{x})$ in T . Then the following holds:

Theorem 5 ([Liu and Lakemeyer, 2009]). *Let T be finite and semi-definitional wrt P , and T' the set of sentences in T not mentioning P . Then $\text{forget}(T, P) \Leftrightarrow T' \wedge \bigwedge_{\psi \in \text{WSC}_P, \phi \in \text{SNC}_P} \forall \vec{x}. \psi(\vec{x}) \supset \phi(\vec{x})$.*

Proposition 6. *In any model of $\mathcal{D}$, the SSA for F instantiated by ground action α , $\mathcal{D}_{ss}^F[\alpha, S_0]$, is equivalent to:*

$$\neg \gamma_F^+(\vec{x}, \alpha, S_0) \wedge F(\vec{x}, S_\alpha) \supset F(\vec{x}, S_0) \quad (5a)$$

$$F(\vec{x}, S_0) \supset \gamma_F^-(\vec{x}, \alpha, S_0) \vee F(\vec{x}, S_\alpha) \quad (5b)$$

$$\gamma_F^+(\vec{x}, \alpha, S_0) \supset F(\vec{x}, S_\alpha) \quad (5c)$$

$$\gamma_F^-(\vec{x}, \alpha, S_0) \supset \neg F(\vec{x}, S_\alpha) \quad (5d)$$

We say a formula ϕ is semi-definitional wrt a fluent F if $\phi \uparrow S_0$ is semi-definitional wrt F 's lifting predicate. The proposition suggests that $\mathcal{D}_{ss}^F[\alpha, S_0]$ for every fluent F is semi-definitional wrt F . Hence, if $\mathcal{D}_{S_0}$ is also semi-definitional wrt fluents, one can use Thm. 5 to forget their lifting predicates in Eq. (1), which is the idea behind normal actions and their progression. For an action α , we distinguish local-effect fluents $\text{LE}(\alpha)$ on the one hand (whose SSAs are as in BAT-LE), and non-local effect fluents $\text{NLE}(\alpha)$ on the other hand. Normal actions then require that for all $F \in \text{NLE}(\alpha)$, fluents occurring in $\gamma_F^\pm$ can only be in $\text{LE}(\alpha)$. We call a BAT $\mathcal{D}$ a BAT-NR (wrt action α) if α is normal and $\mathcal{D}_{S_0}$ is semi-definitional wrt fluents in $\text{NLE}(\alpha)$. For such a BAT, one can forget lifting predicates of fluents in $\text{NLE}(\alpha)$ by Thm. 5 and thereafter fluents in $\text{LE}(\alpha)$ as in Eq. (4).

Theorem 7. *If $\mathcal{D}$ is BAT-NR wrt ground action α , the result of forgetting the lifting predicates for all fluents in $\text{NLE}(\alpha)$ in $(\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]) \uparrow S_0$ can be represented by a FO theory whose size is in $O((n+m)^2)$, where n is the size of $\mathcal{D}_{S_0}$, and m is the size of $\mathcal{D}_{ss}$.*

Proof. First, note that we can ignore the lifting operator $\uparrow S_0$ here, since it only replaces atoms by other atoms, and thus does not affect the size of a formula. Now, assume that for fluents $F \in \text{NLE}(\alpha)$, axioms in $\mathcal{D}_{ss}[\alpha, S_0]$ are represented according to Prop. 6. Note that (5a) and (5b) are semi-definitional wrt $F(\vec{x}, S_0)$, and (5c) and (5d) do not mention $F(\vec{x}, S_0)$. Theorem 5 is hence applicable to $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]$. Intuitively, the set of semi-definitional formulas in $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]$ can be represented as a directed graph, where any formula $\psi(\vec{x}) \supset \phi(\vec{x})$ corresponds to an edge from $\psi(\vec{x})$ to $\phi(\vec{x})$. Forgetting the lifting predicate for a fluent $F \in \text{NLE}(\alpha)$ then amounts to eliminating the node for $F(\vec{x}, S_0)$, “bypassing” all involved edges, i.e., every pair of ingoing edge $\psi(\vec{x}) \supset F(\vec{x}, S_0)$ and outgoing edge $F(\vec{x}, S_0) \supset \phi(\vec{x})$ is replaced by a new edge $\psi(\vec{x}) \supset \phi(\vec{x})$. This elimination operation is performed iteratively for all $F \in \text{NLE}(\alpha)$ (in any order), leaving only nodes that correspond to conditions that do *not* mention any of the non-local-effect fluents. In the worst case, the resulting graph is total, i.e., every condition is connected to every other condition through an edge. Since the size of all conditions together is bounded by $n + m$, the size of the theory representing all combinations of implications between them is in $O((n+m)^2)$. Moreover, the size of the original formulas in $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]$ that do not mention any of the NLE fluents (including (5c) and (5d)) is bounded by $n + m$ as well. All in all, the size of the resulting theory is thus in $O((n+m)^2)$. $\square$

According to [Liu and Lakemeyer, 2009], the progression of a BAT-NR through an action α is obtained by first forgetting the (lifting predicates) of the fluents in $\text{NLE}(\alpha)$, and afterwards forgetting the remaining atoms over the fluents in $\text{LE}(\alpha)$. We can estimate the size of the resulting theory by combining the results from Theorems 4 and 7, giving:

Corollary 8. *If $\mathcal{D}$ is BAT-NR wrt ground action α , then its progression can be represented by a FO theory of size $O(2^c(n+m)^2)$, where c is the size of the characteristic set, n is the size of $\mathcal{D}_{S_0}$, and m is the size of $\mathcal{D}_{ss}$.*

3.3 BAT-AC

While normal action theories allow for non-local effects of a certain form, they include the requirement that for BAT-NR, the effect conditions $\gamma_F^\pm$ for a fluent F may only mention local-effect fluents. This assumption ensures that forgetting the lifting predicate for F will not generate formulas that mention the lifting predicate of another NLE fluent F' , which might not have the right (semi-definitional) form to proceed with forgetting F' in the same manner. However, this excludes domains where actions' effects for a NLE fluent might depend on the state of another NLE fluent. BAT-AC, the class of acyclic action theories [Liu and Claßen, 2024] overcomes this restriction by allowing NLE fluents to depend on one another, as long the dependencies do not form any cycles. Here, additional restrictions on the form of $\gamma_F^\pm$ guarantee that the lifting predicates for NLE fluents can be forgotten in a recursive fashion when following the order of dependencies among fluents. Formally, a formula $\phi(\vec{x})$ is said to be in *good form* wrt a predicate P if it is of the form

$$[\psi(\vec{x}) \vee P(\vec{t})] \wedge [\psi'(\vec{x}) \vee \neg P(\vec{t})] \wedge \psi''(\vec{x}) \quad (6)$$

where ψ, ψ', ψ'' contain no P , and terms in $\vec{t}$ are either ground terms or free variables among $\vec{x}$.

Proposition 9 ([Liu and Claßen, 2024]). *If $\phi(\vec{x})$ is in good form wrt P , then*

1. $\phi(\vec{x})$ can be rewritten to be semi-definitional wrt P ;
2. $\neg\phi(\vec{x})$ can be rewritten to be in good form wrt P .

Now, given SSAs $\mathcal{D}_{ss}$ and a ground action α , the *dependency graph* G for $\mathcal{D}_{ss}$ wrt α is a directed graph whose vertices are the set of fluents and there is an edge $F \rightarrow F'$ for two fluents F, F' iff $F \in \text{NLE}(\alpha)$ and F' appears in $\gamma_F^\pm$. A ground action α is *acyclic* (or has acyclic effects) if 1) its dependency graph G is acyclic, i.e., a DAG; and 2) for each fluent F , there are *at most two* fluents F', F'' in $\text{NLE}(\alpha)$ such that $F \rightarrow F'$ and $F \rightarrow F''$, where $\gamma_F^\pm$ and $\gamma_{F'}^\pm$ are in good form wrt F' and F'' , respectively.

We call a BAT a BAT-AC (wrt α), if α is acyclic and $\mathcal{D}_{S_0}$ is *separably semi-definitional* wrt $\text{NLE}(\alpha)$, that is, each sentence in $\mathcal{D}_{S_0}$ mentions at most one fluent F in $\text{NLE}(\alpha)$ and is semi-definitional wrt F . [Liu and Claßen, 2024] showed that if a BAT $\mathcal{D}$ is a BAT-AC, the progression of $\mathcal{D}_{S_0}$ wrt $\alpha, \mathcal{D}$ is FO definable. This is because for any $F, F' \in \text{NLE}(\alpha)$, even if forgetting the lifting predicate of F will possibly generate new formulas that mention F' , the assumption of $\gamma_F^\pm$ being in good form wrt F' ensures that the new formulas can be rewritten to be semi-definitional wrt F' again. Hence, one can recursively forget lifting predicates for $F \in \text{NLE}(\alpha)$ (following the order of the DAG) and later forget lifting predicates of $F \in \text{LE}(\alpha)$ as in Eq. (4). For analyzing the size of the resulting theory, we need the following:

Definition 10. If $\phi(\vec{x})$ is in good form wrt P , i.e., $\phi(\vec{x}) := (\psi(\vec{x}) \vee P(\vec{t})) \wedge (\psi'(\vec{x}) \vee \neg P(\vec{t})) \wedge \psi''(\vec{x})$, we define $|\phi|_P := \max\{|\psi|, |\psi'|, |\psi''|\}$ as the *condition size* of ϕ wrt P .

As we will see, the size of conditions within good-form formulas will be the determining factor for the size complexity of a progression for a BAT-AC. Furthermore, we assume

that for each $F \in \text{NLE}(\alpha)$, $\mathcal{D}_{S_0}$ contains *exactly one* axiom $\Delta_F = (NWS_F(\vec{x}) \vee F(\vec{x}, S_0)) \wedge (SNC_F(\vec{x}) \vee \neg F(\vec{x}, S_0))$

This is without loss of generality: If $\mathcal{D}_{S_0}$ contains multiple semi-definitional formulas for F , let NWS_F be the negation of the weakest sufficient condition, and SNC_F the strongest necessary condition. Observe that with the implicit negation in NWS , Δ_F is indeed in semi-definitional form, and such formulas can be assumed to be of good form (6). Also, this assumption has no impact on size or condition size:

Proposition 11. *If $\phi_1(\vec{x}), \phi_2(\vec{x}), \dots, \phi_n(\vec{x})$ are all semi-definitional wrt predicate P , then $\phi_1(\vec{x}) \wedge \phi_2(\vec{x}) \wedge \dots \wedge \phi_n(\vec{x})$ can be rewritten to a single equivalent formula $\varphi(\vec{x})$ such that $|\varphi(\vec{x})| \leq |\bigwedge_i \phi_i(\vec{x})|$ and $|\varphi(\vec{x})|_P \leq |\bigwedge_i \phi_i(\vec{x})|_P$.*

Next, we consider formulas in good form and how rewriting them as in Prop. 9 affects the size of the result:

Proposition 12. *If $\phi(\vec{x})$ is in good form wrt P , then*

1. $\phi(\vec{x})$ can be rewritten to $\varphi(\vec{y})$ that is semi-definitional wrt $P(\vec{y})$, moreover, if $a = \max\{|\vec{x}|, |\vec{y}|\}$, then $|\varphi(\vec{y})| \leq |\phi(\vec{x})| + 7a + 2$ and $|\varphi(\vec{y})|_P \leq |\phi(\vec{x})|_P + 3a + 1$;
2. $\neg\phi(\vec{x})$ can be rewritten to $\varphi(\vec{x})$ that is in good form wrt P with $|\varphi(\vec{x})| \leq 3|\phi(\vec{x})|$ and $|\varphi(\vec{x})|_P \leq 3|\phi(\vec{x})|_P + 2$;
3. if $\psi(\vec{x})$ mentions no P , then $\phi(\vec{x}) \vee \psi(\vec{x})$ can be rewritten to a formula $\varphi(\vec{x})$ in good form wrt P , where $|\varphi(\vec{x})| = |\phi(\vec{x})| + 3|\psi(\vec{x})| + 3$ and $|\varphi(\vec{x})|_P = |\phi(\vec{x})|_P + |\psi(\vec{x})| + 1$.

Then we have:

Theorem 13. *If $\mathcal{D}$ is BAT-AC wrt ground action α , the result of forgetting the lifting predicates for all fluents in $\text{NLE}(\alpha)$ in $(\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]) \uparrow S_0$ can be represented by a FO theory whose size is in $O(2^d(n + m))$, where d is the depth of the dependency graph, n the size of $\mathcal{D}_{S_0}$, and m the size of $\mathcal{D}_{ss}$.*

Proof. The proof for Theorem 3 in [Liu and Claßen, 2024] makes an inductive argument that one can forget fluents $F \in \text{NLE}(\alpha)$ iteratively while maintaining the required form, as long as one forgets all predecessors of F in the dependency graph before forgetting F itself. By Prop. 6 and Thm. 5, forgetting any such F amounts to replacing the corresponding Δ_F by the following (parameters of $\gamma_F^\pm$ are omitted):

$$NWS_F(\vec{x}) \vee SNC_F(\vec{x}) \quad (7a)$$

$$\neg\gamma_F^+ \vee F(\vec{x}, S_\alpha) \quad (3k + 3a + 5) \quad (7b)$$

$$\neg\gamma_F^- \vee \neg F(\vec{x}, S_\alpha) \quad (3k + 3a + 6) \quad (7c)$$

$$\gamma_F^+ \vee \neg F(\vec{x}, S_\alpha) \vee SNC_F(\vec{x}) \quad (w_F + k + 3a + 5) \quad (7d)$$

$$NWS_F(\vec{x}) \vee \gamma_F^- \vee F(\vec{x}, S_\alpha) \quad (w_F + k + 3a + 4) \quad (7e)$$

By definition, there are at most two other NLE fluents F_1, F_2 with $F \rightarrow F_1$ and $F \rightarrow F_2$, where F_1 only occurs in γ_F^+ , and F_2 only in γ_F^- . (7b) and (7d) will subsequently be incorporated into the new axiom Δ_{F_1} , and (7c) and (7e) into the new axiom Δ_{F_2} . If w_F is the condition size of the updated Δ_F and k the maximal size of an SSA for any NLE fluent, then by Prop. 12, upper bounds for the condition sizes of formulas (7b)–(7e) after rewriting them to be in good form wrt F_1, F_2 are as stated above in parentheses. If w is an upper bound for

the condition size of any *original* Δ_F axiom, the condition size for the *updated* $\Delta_{F'}$ is bounded according to

$$w_{F'} \leq w + \sum_{F \rightarrow F'} (w_F + 4k + 6a + 10). \quad (8)$$

The sum of the sizes of (7a)–(7e) over all NLE fluents F gives an upper bound for the size of the progression. Observe that the total size of $\gamma_F^\pm$, $F(\vec{x}, S_\alpha)$ atoms, and boolean connectives is less than $|\mathcal{D}_{ss}^F[\alpha, S_0]|$ as in Prop. 6, i.e., less than m summed over all NLE fluents. This leaves two occurrences of $NWS_F(\vec{x})$ and two of $SNC_F(\vec{x})$, each of size $\leq w_F$, summing up to $S = 4 \sum_{F \in \text{NLE}(\alpha)} w_F$. To determine the sum of all w_F , we use the intuition that the linear recurrence in (8) can be unrolled along all source-to-node paths. Intuitively, each node F starts with a w_F “weight” of w and adds $w_F + 4k + 6a + 10$ to each of its children. This contribution is added multiple times, namely for each path starting in F . As the out-degree of nodes is ≤ 2 , the number of distinct paths starting in F is $\leq 2^{d+1}$. If l is the number of NLE fluents, we get $S \leq 4 \cdot 2^{d+1} \cdot l \cdot (w + 4k + 6a + 10)$. Since a is a constant, $l \cdot w \in O(n)$, and $l \cdot k \in O(m)$, the claim follows. $\square$

Combining the above theorem with Theorem 4 again yields:

Corollary 14. *If $\mathcal{D}$ is BAT-AC wrt ground action α , then its progression can be represented by a FO theory of size $O(2^{c+d}(n + m))$, where c is the size of the characteristic set, d is the depth of the dependency graph, n is the size of $\mathcal{D}_{S_0}$, and m is the size of $\mathcal{D}_{ss}$.*

Observe that the progression for BAT-LE and BAT-AC is linear in the size of the theory, while that for BAT-NR is quadratic. As mentioned previously, the three classes roughly form a hierarchy of increasing expressiveness, but not exactly. This is because BAT-NR allow “mixed” semi-definitional axioms of the form $F(\vec{x}, S_0) \supset G(\vec{x}, S_0)$, where both F and G are NLE fluents, whereas BAT-AC requires a strict separation so that a semi-definitional axiom for F does not mention any NLE fluent G , and vice versa. If we impose this additional restriction on BAT-NR, the resulting class will indeed be a proper subset of BAT-AC, consisting of action theories whose dependency graph has a depth of at most one, and where every edge $F \rightarrow G$ is between an NLE fluent F and a LE fluent G . The size of the progression then is in $O(2^c(n + m))$.

4 The Two-Variable Fragment

For practical applications like FO planning, it is not enough to have an FO progression, but the progressed KB has to admit effective query evaluation. This necessitates a decidable progressed theory. A class of theory $\mathcal{T}$ is called *decidable* if for any $T, \phi \in \mathcal{T}$, it is decidable to check if $T \models \phi$. One popular decidable class of FO theories is the *two-variable fragment* with equality, called FO² [Grädel et al., 1997]. It consists of all formulas with at most 2 variable symbols, which we denote as x and y by convention. By $\vec{x}$ we then mean a vector of variables that either consists only of x or of x and y .

Theorem 15 ([Grädel et al., 1997]). *FO² has the exponential model property: there is a constant c s.t. every satisfiable FO²-sentence ϕ has a model of cardinality at most $2^{c|\phi|}$, where $|\phi|$ is the size of ϕ .*

Consequently, satisfiability of FO^2 is NEXPTIME-complete and entailment checking is CoNEXPTIME-complete.

Typically, FO^2 does not include general function symbols. Yet, in some literature, constant symbols are allowed as they can be simulated through unary predicates whose extensions consist of exactly one element. That is to say, for any constant A , one introduces a predicate $P_A(x)$ and includes an axiom

$$\exists x P_A(x) \wedge \forall x \forall y. P_A(x) \wedge P_A(y) \supset (x = y)$$

saying that there is exactly one domain element in the extension of P_A . It is easy to see that this idea can be extended to also include ground situation terms $do(\alpha, S_0)$.

Moreover, in the following, we assume that every rigid relation symbol has arity at most two. Again, relations of higher arity can be simulated by introducing new predicate symbols, with formulas growing at most linearly in size. For details, we refer the reader to [Grädel *et al.*, 1997]. For fluent predicates, we allow for up to two *object* arguments in addition to the *situation* argument. The idea here is that the situation argument is irrelevant, as it will be instantiated by some ground situation term when evaluating queries.

Definition 16. Given a ground action α , a BAT $\mathcal{D}$ is an FO^2 -BAT wrt α if $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]$ is a finite FO^2 theory.

In what follows, we show that the progressions of FO^2 -BATs wrt local-effect, normal, and acyclic actions are all within the FO^2 fragment, hence decidable.

FO^2 -BAT-LE To begin with, we consider FO^2 -BAT where the action α has local effects, i.e., FO^2 -BAT-LE.

Theorem 17. *Given a BAT $\mathcal{D}$ that is BAT-LE wrt ground action α , if $\mathcal{D}$ is also an FO^2 -BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in FO^2 .*

It is not hard to check that $(\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\Omega])[\theta]$ in Eq. (4) is in FO^2 by assumption. Besides, a disjunction of FO^2 theories is still in FO^2 , so the progression as in Eq. (4) is in FO^2 .

FO^2 -BAT-NR For FO^2 -BAT with normal actions, we have

Theorem 18. *Given a BAT $\mathcal{D}$ that is BAT-NR wrt ground action α , if $\mathcal{D}$ is an FO^2 -BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in FO^2 .*

For BAT-NR, forgetting NLE fluents as in Thm. 5 involves combining sufficient conditions and necessary conditions of semi-definitional formulas, which clearly does not introduce new variables. Hence, the result is still in FO^2 . Besides, for LE fluents, as mentioned before, forgetting them only involves disjoining two FO^2 theories, which is still in FO^2 . Hence, the resulting progression is in FO^2 .

FO^2 -BAT-AC The situation for FO^2 -BAT with acyclic actions is slightly more complicated, as we need to ensure that the generated formulas when forgetting an NLE fluent are semi-definitional wrt another NLE fluent and the rewrite has to be in FO^2 . The following lemma justifies this.

Lemma 19. *If $\phi(\vec{x}) \in \text{FO}^2$ is in good form wrt P , then*

1. $\phi(\vec{x})$ can be rewritten to a formula in FO^2 that is semi-definitional wrt P ;
2. $\neg\phi(\vec{x})$ can be rewritten to a formula in FO^2 that is in good form wrt P .

The proof of Item 2 is simple, as the rewrite only involves pushing negation inward and distributing conjunctions over disjunctions, operations that preserve membership in FO^2 . For Item 1, let $\phi(\vec{x}) = \{\forall \vec{x}(\psi(\vec{x}) \vee P(\vec{t})), \forall \vec{x}(\psi'(\vec{x}) \vee \neg P(\vec{t})), \forall \vec{x}\psi''(\vec{x})\}$. By assumption, $\vec{x}$ denotes either only x or x and y . Since P is unary or binary, $\vec{t}$ has length one or two, and since functions are excluded, $\vec{t}$ can only consist of a combination of x, y , and some constants A and B . Hence, for $\forall \vec{x}(\psi(\vec{x}) \vee P(\vec{t}))$ (likewise for the $\neg P$ case), we have cases:

- $\forall x(\psi(x) \vee P(A)) \Leftrightarrow \forall x((x = A) \wedge \exists x \neg \psi(x)) \supset P(x)$.
- $\forall x(\psi(x) \vee P(x))$ has the right form already.
- $\forall x \forall y(\psi(x, y) \vee P(A, B))$ can be rewritten as $\forall x \forall y((x = A \wedge y = B) \wedge \exists x \exists y \neg \psi(x, y)) \supset P(x, y)$.
- $\forall x \forall y(\psi(x, y) \vee P(x, B))$ can be rewritten as $\forall x \forall y((y = B) \wedge \exists y \neg \psi(x, y)) \supset P(x, y)$.
- $\forall x \forall y(\psi(x, y) \vee P(A, y))$ can be rewritten as $\forall x \forall y((x = A) \wedge \exists x \neg \psi(x, y)) \supset P(x, y)$.
- $\forall x \forall y(\psi(x, y) \vee P(x, y))$ has the right form already.

Hence $\forall \vec{x}(\psi(\vec{x}) \vee P(\vec{t}))$, thus $\phi(\vec{x})$, can be expressed in FO^2 .

Theorem 20. *Given a BAT $\mathcal{D}$ that is BAT-AC wrt ground action α , if $\mathcal{D}$ is a FO^2 -BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in FO^2 .*

We remark that the results in this section are not obvious in the following sense. FO^2 is a fragment that fails to have the Craig interpolation and Beth definability properties [Jung and Wolter, 2021]. Among other things, this means that while forgetting a predicate P from a FO^2 theory might be possible (i.e., yields an equivalent FO theory), it can happen that the resulting theory is not expressible in FO^2 . Hence, not every FO^2 -BAT has a progression that is in FO^2 , but for the fragments we consider here, this can be guaranteed.

5 Universal Theories with Constants

Another decidable fragment of FOL is given by universal theories with constants (UTC). In fact, it is shown that Bernays-Schönfinkel [1928] (BS) theories are decidable, i.e., theories with sentences of the form $\exists x_1 \dots \exists x_k \forall y_1 \dots \forall y_m \phi$, where ϕ is a quantifier-free formula (potentially with equalities) mentioning no functions except constants. Satisfiability checking of BS theories is NEXPTIME-complete. UTC theories are a special case of BS theories where all sentences are universally quantified, i.e., of the form $\forall y_1 \dots \forall y_m \phi$. For UTC theory T and query ϕ , deciding if $T \models \phi$ can be converted to check if the BS theory $T \cup \{\neg\phi\}$ is unsatisfiable, hence the complexity is CoNEXPTIME-complete.

Definition 21. Given a ground action α , a BAT $\mathcal{D}$ is a UTC-BAT wrt α if $\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]$ is a finite UTC theory.

In what follows, we show that the progressions of UTC-BATs wrt local-effect, normal, and acyclic actions are all within the UTC fragment, hence decidable.

UTC-BAT-LE First, for local effect actions, we have that

Theorem 22. *Given a BAT $\mathcal{D}$ that is BAT-LE wrt ground action α , if $\mathcal{D}$ is an UTC-BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in UTC.*

Again, the key here is to show that the disjunction of UTC theories can be expressed in UTC. This is indeed the case by applying the following rule to properly rename variables $\forall \vec{x} \phi(\vec{x}) \vee \forall \vec{x} \psi(\vec{x}) \Leftrightarrow \forall \vec{x} \forall \vec{y} (\phi(\vec{x}) \vee \psi(\vec{y}))$.

UTC-BAT-NR Likewise, for normal actions, we have

Theorem 23. *Given a BAT $\mathcal{D}$ that is BAT-NR wrt ground action α , if $\mathcal{D}$ is an UTC-BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in UTC.*

The proof is like its counterpart in FO^2 , it suffices to show that forgetting NLE fluents preserves membership in UTC. This is indeed the case as the sufficient and necessary conditions to be combined when forgetting NLE fluents, as in Thm. 5, are all quantifier-free by assumption; hence, the result is still quantifier-free, leading to a UTC FO progression.

UTC-BAT-AC Lastly, for acyclic actions, we have that

Lemma 24. *If $\phi(\vec{x}) \in \text{UTC}$ is in good form wrt P , then*

1. $\phi(\vec{x})$ can be rewritten to a formula in UTC that is semi-definitional wrt P ;
2. $\neg\phi(\vec{x})$ can be rewritten to a formula in UTC that is in good form wrt P .

Again, Item 2 is trivial as the rewrite only involves pushing negation inward and distributing conjunctions over disjunctions, operations that preserve membership in UTC. We prove the first item. For Item 1, let $\phi(\vec{x}) = \{\forall \vec{t} (\psi(\vec{x}) \vee P(\vec{t})), \forall \vec{t} (\psi'(\vec{x}) \vee \neg P(\vec{t})), \forall \vec{t} \psi''(\vec{x})\}$. $\psi(\vec{x}), \psi'(\vec{x}), \psi''(\vec{x})$ are all quantifier-free and $\vec{t}$ can only contain variables in $\vec{x}$ by assumption. $\forall \vec{t} (\psi(\vec{x}) \vee P(\vec{t}))$ (likewise for the $\neg P$ case) can be rewritten as $\forall \vec{x} \forall \vec{y}. \vec{y} = \vec{t} \supset (\psi(\vec{x}) \vee P(\vec{y}))$, which is in UTC and semi-definitional wrt P , and hence so can $\phi(\vec{x})$.

Theorem 25. *Given a BAT $\mathcal{D}$ that is BAT-AC wrt ground action α , if $\mathcal{D}$ is an UTC-BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}, \alpha$ can be expressed in UTC.*

We remark that the fact that the progression of a UTC-BAT can be expressed by another UTC theory here is not so surprising, as [Arenas *et al.*, 2018] showed that the progression of a UTC-BAT wrt *any actions* can be expressed by another UTC theory. In fact, a resolution-like algorithm is proposed there to compute such a progression. Nevertheless, the progression might not be finite, and deciding if a finite progression exists is undecidable; namely, the proposed algorithm might not terminate. In this regard, our contribution here is that we show that if the actions are furthermore local-effect, normal, or acyclic, finite progressions of the respective UTC-BATs actually exist (and can be computed accordingly). In addition, due to our results in Section 3, the size of such finite progressions grows only polynomially, even linearly.

6 Related Work and Discussion

In this paper, we studied the size complexity and decidability of FO progressions in terms of local-effect, normal, and acyclic actions. Previous contributions along similar lines include those due to Liu and Lakemeyer [2009] and due to Liu and Claßen [2024], who showed that a progression can be efficiently computed for normal and acyclic theories, respectively, under the restrictions that the instantiated SSAs

are quantifier-free and that the initial theory is in so-called *proper*⁺ form. Intuitively, *proper*⁺ KBs [Lakemeyer and Levesque, 2002] correspond to consistent, potentially infinite sets of ground clauses. What makes them particularly interesting is that certain queries can be computed in polynomial time [Liu and Levesque, 2005]. De Giacomo *et al.* [2016] showed that progression is FO definable for Situation Calculus theories where the number of fluent instances affected by an action is *bounded*, in which case verification (and hence projection) is decidable. Arenas *et al.* [2018] investigate the feasibility and complexity of progression for UTC-based action theories (without the additional restriction of SSAs to local-effect, normal or acyclic form as done here). Eiter and Soldà [2024] studied computational aspects of progression in the context of temporal equilibrium logic.

More broadly, related work includes that by Gu and Soutchanski [2010] who studied the complexity of *regression* for action theories over the two-variable fragment of FOL with counting quantifiers as well as description logics $\mathcal{ALCCO}(U)$ and $\mathcal{ALCCQO}(U)$. Lakemeyer and Levesque [2014] presented a variant of the modal epistemic Situation Calculus based on a model of limited belief that admits a decidable, yet incomplete form of reasoning. Calvanese *et al.* [2015] proved that Situation Calculus theories extended with description logic TBoxes acting as state constraints lead to undecidability of the satisfiability problem, even for the simplest kinds of description logics, and the simplest kind of action theories. When TBoxes are merely viewed as part of the initial theory rather than state constraints, Zarriß [2018] showed that the verification problem (which subsumes the projection problem) for action formalisms based on various description logics becomes decidable, providing exact complexities. Note that while many description logics can be viewed as fragments of FO^2 , related results on projection are not directly applicable to a FO^2 -based Situation Calculus since translating FO^2 into description logic leads to a (provably unavoidable) exponential blow-up [Lutz *et al.*, 2001].

Possible directions for future work include exploring whether more expressive classes, or less expressive but more tractable classes, admit similar closure properties as presented here. Our results also suggest the possibility to use progression instead of regression for the purpose of decidable verification of Golog as in [Zarriß and Claßen, 2016]. Our main motivation however is in practical applications, i.e., equipping real-world agents with the capacity to plan and act in the presence of incomplete information and unbounded domains. There is a recent trend in AI planning to work directly on FO representations [Corrêa and De Giacomo, 2024], but merely as a way to avoid the bottleneck of grounding over the given set of objects, which is still assumed to be finite. When it comes to the full expressiveness of Situation Calculus (within decidability bounds), there is currently a lack of concrete implementations and corresponding benchmarks. We believe that our work here lays the foundations towards such planning systems that work in truly open worlds.

Appendix

Proof of Theorem 13

First, we provide a detailed proof for Proposition 12.

Proposition (Restatement of Proposition 12). If $\phi(\vec{x})$ is in good form wrt P , then

1. $\phi(\vec{x})$ can be rewritten to $\varphi(\vec{y})$ that is semi-definitional wrt $P(\vec{y})$, moreover, if $a = \max\{|\vec{x}|, |\vec{y}|\}$, then $|\varphi(\vec{y})| \leq |\phi(\vec{x})| + 7a + 2$ and $|\varphi(\vec{y})|_P \leq |\phi(\vec{x})|_P + 3a + 1$;
2. $\neg\phi(\vec{x})$ can be rewritten to $\varphi(\vec{x})$ that is in good form wrt P with $|\varphi(\vec{x})| \leq 3|\phi(\vec{x})|$ and $|\varphi(\vec{x})|_P \leq 3|\phi(\vec{x})|_P + 2$;
3. if $\psi(\vec{x})$ mentions no P , then $\phi(\vec{x}) \vee \psi(\vec{x})$ can be rewritten to a formula $\varphi(\vec{x})$ in good form wrt P , where $|\varphi(\vec{x})| = |\phi(\vec{x})| + 3|\psi(\vec{x})| + 3$ and $|\varphi(\vec{x})|_P = |\phi(\vec{x})|_P + |\psi(\vec{x})| + 1$.

Proof. Let $\phi(\vec{x}) = (\psi(\vec{x}) \vee P(\vec{t})) \wedge (\psi'(\vec{x}) \vee \neg P(\vec{t})) \wedge \psi''(\vec{x})$,

1. $\phi(\vec{x})$ is equivalent to $\varphi(\vec{y})$ given by

$$\begin{aligned} & ((\forall \vec{x}. \vec{y} \neq \vec{t} \vee \psi(\vec{x})) \vee P(\vec{y})) \wedge \\ & ((\forall \vec{x}. \vec{y} \neq \vec{t} \vee \psi'(\vec{x})) \vee \neg P(\vec{y})) \wedge (\forall \vec{x}. \psi''(\vec{x})) \end{aligned} \quad (9)$$

Hence, the size increases by 3 occurrences of $\forall \vec{x}$ (each size l), 2 occurrences of $(\vec{y} \neq \vec{t})$ (each size $2l$), and two additional disjunction symbols;

2. $\neg\phi(\vec{x})$ is equivalent to a formula $\varphi(\vec{x})$ given by
 $(\neg\psi(\vec{x}) \vee \neg\psi''(\vec{x}) \vee P(\vec{t})) \wedge (\neg\psi'(\vec{x}) \vee \neg\psi''(\vec{x}) \vee \neg P(\vec{t})) \wedge (\neg\psi(\vec{x}) \vee \neg\psi'(\vec{x}) \vee \neg\psi''(\vec{x}))$
 Let $l_1 = |\psi(\vec{x})|$, $l_2 = |\psi'(\vec{x})|$, and $l_3 = |\psi''(\vec{x})|$. Then $|\phi(\vec{x})| = l_1 + l_2 + l_3 + 7$, and $|\varphi(\vec{x})| = 2l_1 + 2l_2 + 3l_3 + 18$. Also, $|\phi(\vec{x})|_P = \max\{l_1, l_2, l_3\}$, and every condition in $\varphi(\vec{x})$ has size $\leq l_1 + l_2 + l_3 + 2 \leq |\varphi(\vec{x})|_P + 2$;
3. $\phi(\vec{x}) \vee \psi(\vec{x})$ is equivalent to $\varphi(\vec{x})$ given by
 $(\psi(\vec{x}) \vee \psi'(\vec{x}) \vee P(\vec{t})) \wedge (\psi(\vec{x}) \vee \psi''(\vec{x}) \vee \neg P(\vec{t})) \wedge (\psi(\vec{x}) \vee \psi''(\vec{x}))$. Observe that the new formula adds 3 occurrences of $\psi(\vec{x})$ and 3 disjunction symbols to $\phi(\vec{x})$.

□

Theorem (Restatement of Theorem 13). If $\mathcal{D}$ is BAT-AC wrt ground action α , the result of forgetting the lifting predicates for all fluents in $\text{NLE}(\alpha)$ in $(\mathcal{D}_{S_0} \cup \mathcal{D}_{ss}[\alpha, S_0]) \uparrow S_0$ can be represented by a FO theory whose size is in $O(2^d(n + m))$, where d is the depth of the dependency graph, n the size of $\mathcal{D}_{S_0}$, and m the size of $\mathcal{D}_{ss}$.

Proof. The proof for Theorem 3 in [Liu and Claßen, 2024] makes an inductive argument that it is possible to iteratively forget fluents in $\text{NLE}(\alpha)$ of decreasing fluent depth while maintaining the required form. Looking at an arbitrary F , this means that we forget all its predecessors F^* before forgetting F itself. By induction assumption, none of the F^* have any predecessors to consider, and there is a single semi-definitional axiom of the form

$$\Delta_F = (NWS_F(\vec{x}) \vee F(\vec{x}, S_0)) \wedge (SNC_F(\vec{x}) \vee \neg F(\vec{x}, S_0))$$

in the initial KB for each of them (let Δ_F denote the axiom for F henceforth). Hence, to forget F , one only needs to consider $\Delta_F \wedge \mathcal{D}_{ss}^F[\alpha, S_0]$, which by Prop. 6 amounts to replacing

the original Δ_F by the following (parameters of $\gamma_F^\pm$ are omitted, expressions in parentheses are explained further below):

$$NWS_F(\vec{x}) \vee SNC_F(\vec{x}) \quad (7a)$$

$$\neg\gamma_F^+ \vee F(\vec{x}, S_\alpha) \quad (3k + 4) \quad (7b)$$

$$\neg\gamma_F^- \vee \neg F(\vec{x}, S_\alpha) \quad (3k + 5) \quad (7c)$$

$$\gamma_F^+ \vee \neg F(\vec{x}, S_\alpha) \vee SNC_F(\vec{x}) \quad (w_F + k + 4) \quad (7d)$$

$$NWS_F(\vec{x}) \vee \gamma_F^- \vee F(\vec{x}, S_\alpha) \quad (w_F + k + 3) \quad (7e)$$

Formula (7a) will remain unchanged when forgetting subsequent fluents (recall that NWS_F and SNC_F do not mention any NLE fluents). By definition, there are at most two other NLE fluents F'_+ , F'_- in the dependency graph such that $F \rightarrow F'_+$ and $F \rightarrow F'_-$, where F'_+ only occurs in γ_F^+ , and F'_- only in γ_F^- . Formulas (7b) and (7d) will subsequently be incorporated into the new KB axiom $\Delta_{F'_+}$, and formulas (7a) and (7e) into the new axiom $\Delta_{F'_-}$.

We now estimate the condition size of $\Delta_{F'}$ in terms of the conditions size of Δ_F for $F \rightarrow F'$. For this purpose, let w_F denote the condition size of the *updated* Δ_F for any F , after its predecessors have been forgotten, but before forgetting F itself. Furthermore, let k be the maximal size of an SSA for any NLE fluent, and w be the maximal condition size of any Δ_F in the original KB. Using Prop. 12, this gives us upper bounds for the condition sizes of formulas (7b)–(7e) as stated above in parentheses. Observe that the sum of condition sizes of (7b) and (7d) for F'_+ is the same as the sum of condition sizes for (7c) and (7e) for F'_- , namely $w_F + 4k + 8$. Also by Prop. 12, converting each of them to semi-definitional form adds at most $(3a + 1)$ to their individual condition size, where a is the maximal arity among all fluents. The original $\Delta_{F'}$ already have the right form, and their condition size is bounded by w . Putting these together, by Prop. 11 we have that the condition size for the updated $\Delta_{F'}$ is bounded according to

$$w_{F'} \leq w + \sum_{F \rightarrow F'} (w_F + 4k + 6a + 10). \quad (8)$$

Note that this bound also holds for fluents without predecessors in the dependency graph; in this case, the sum is empty and only w remains.

For determining the size of the resulting KB, note that formulas of the form (7a) are added for all NLE fluents. Formulas of the form (7c)–(7e) are deleted in case that the $\gamma_F^\pm$ mention other NLE fluents F' , but kept if they only mention LE fluents. An upper bound is hence given by the sum of sizes of (7a)–(7e) over *all* NLE fluents F . Observe that the total size of occurrences of $\gamma_F^\pm$, $F(\vec{x}, S_\alpha)$ atoms, and boolean connectives is less than the size of $\mathcal{D}_{ss}^F[\alpha, S_0]$ when using the representation of Prop. 6, summing up to less than $|\mathcal{D}_{ss}[\alpha, S_0]|$ over all NLE fluents. This leaves two occurrences of $NWS_F(\vec{x})$ and two occurrences of $SNC_F(\vec{x})$, each of size $\leq w_F$, summing up to

$$S = 4 \sum_{F \in \text{NLE}(\alpha)} w_F.$$

To determine the sum of all w_F , we use the intuition that the linear recurrence in (8) can be unrolled along all source-to-node paths. Intuitively, each node starts with a w_F “weight” of w . For each child, a node F adds $w_F + 4k + 6a + 10$ to each of its children. This contribution of F is added multiple times, namely for each path starting in F . Because the out-degree is bounded by 2, the number of distinct paths starting in F is $\leq 2^{d+1}$, where d is the depth of the dependency graph. If the number of NLE fluents is l , we hence get $S \leq 4 \cdot 2^{d+1} \cdot l \cdot (w + 4k + 6a + 10)$. The claim follows from the facts that a is a constant, $l \cdot w \in O(|\mathcal{D}_{S_0}|)$, and $l \cdot k \in O(|\mathcal{D}_{ss}|)$. $\square$

Proof of Theorem 18

To prove Theorem 18, we need the following lemma.

Lemma 26. *If T is an FO^2 theory and semi-definitional wrt P , then $\text{forget}(T, P)$ can be expressed in FO^2 .*

The proof is directly by Thm. 5. $\text{forget}(T, P) \Leftrightarrow T' \wedge \bigwedge_{\psi \in \text{WSC}_P, \phi \in \text{SNC}_P} \forall \vec{x}. \psi(\vec{x}) \supset \phi(\vec{x})$. Since T' , ψ , and ϕ can mention at most two variables x, y , $\text{forget}(T, P) \in \text{FO}^2$.

Theorem (Restatement of Theorem 18). Given a BAT $\mathcal{D}$ that is BAT-NR wrt ground action α , if $\mathcal{D}$ is an FO^2 -BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}$, α can be expressed in FO^2 .

Proof. The result is a consequence of Lemma 26. Since for normal actions, their FO progression is computed by applying Thm. 5 when forgetting non-local-effect fluents, the FO progression is also within FO^2 if we apply Lemma 26. Besides, for local-effect fluents, forgetting them via Eq. (4) leads to a theory that is still in FO^2 by Thm. 17; this completes the proof. $\square$

Proof of Theorem 20

Theorem (Restatement of Theorem 20). Given a BAT $\mathcal{D}$ that is BAT-AC wrt ground action α , if $\mathcal{D}$ is an FO^2 -BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}$, α can be expressed in FO^2 .

Proof. The proof is pretty much like the proof of Thm. 13 except that we can now use Lemma 19 to show that the formulas generated when forgetting non-local-effect fluents can be rewritten to a FO^2 formula that is semi-definitional wrt another non-local-effect fluent. This process can be iterated over all non-local-effect fluents. For local-effect fluents, one can use Thm. 17 to ensure that forgetting them leads to an FO^2 theory as well. $\square$

Proof of Theorem 23

To prove Theorem 23, we need the following lemma.

Lemma 27. *If T is an UTC theory and semi-definitional wrt P , then $\text{forget}(T, P)$ can be expressed in UTC.*

The proof is directly from Thm. 5. $\text{forget}(T, P) \Leftrightarrow T' \wedge \bigwedge_{\psi \in \text{WSC}_P, \phi \in \text{SNC}_P} \forall \vec{x}. \psi(\vec{x}) \supset \phi(\vec{x})$. Since $T' \in \text{UTC}$ and ψ, ϕ are quantifier-free, $\text{forget}(T, P) \in \text{UTC}$.

Theorem (Restatement of Theorem 23). Given a BAT $\mathcal{D}$ that is BAT-NR wrt ground action α , if $\mathcal{D}$ is an UTC-BAT wrt α , then the progression of $\mathcal{D}_0$ wrt $\mathcal{D}$, α can be expressed in UTC.

Proof. It is a consequence of Lemma 27. Since for normal actions, their FO progression is computed by applying Thm. 5 when forgetting non-local-effect fluents, the FO progression is also within UTC if we apply Lemma 27. Besides, for local-effect fluents, forgetting them by Eq. (4) leads to a theory that is still in UTC by Thm. 22; this completes the proof. $\square$

Acknowledgements

Daxin is funded by the National Natural Science Foundation of China (NSFC No. 62506152) and the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM118).

References

- [Ackermann, 1935] Wilhelm Ackermann. Untersuchungen über das Eliminationsproblem der mathematischen Logik. *Mathematische Annalen*, 110(1):390–413, 1935.
- [Arenas et al., 2018] Marcelo Arenas, Jorge A. Baier, Juan S. Navarro, and Sebastian Sardiña. On the progression of situation calculus universal theories with constants. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Proceedings of the Sixteenth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2018)*, pages 484–493. AAAI Press, 2018.
- [Bernays and Schönfinkel, 1928] Paul Bernays and Moses Schönfinkel. Zum Entscheidungsproblem der mathematischen Logik. *Mathematische Annalen*, 99:342–372, 1928.
- [Calvanese et al., 2015] Diego Calvanese, Giuseppe De Giacomo, and Mikhail Soutchanski. On the undecidability of the situation calculus extended with description logic ontologies. In Qiang Yang and Michael Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2840–2846. AAAI Press, 2015.
- [Corrêa and De Giacomo, 2024] Augusto B. Corrêa and Giuseppe De Giacomo. Lifted planning: Recent advances in planning using first-order representations. In Kate Larson, editor, *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 8010–8019. ijcai.org, 2024.
- [De Giacomo et al., 2016] Giuseppe De Giacomo, Yves Lespérance, Fabio Patrizi, and Stavros Vassos. Progression and verification of situation calculus agents with bounded beliefs. *Studia Logica*, 104:705–739, 2016.
- [Eiter and Soldà, 2024] Thomas Eiter and Davide Soldà. Computational aspects of progression for temporal equilibrium logic. In Kate Larson, editor, *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 3342–3350. ijcai.org, 2024.
- [Grädel et al., 1997] Erich Grädel, Phokion G. Kolaitis, and Moshe Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

- [Gu and Soutchanski, 2010] Yilan Gu and Mikhail Soutchanski. A description logic based situation calculus. *Annals of Mathematics and Artificial Intelligence*, 58(1–2):3–83, 2010.
- [Jung and Wolter, 2021] Jean Christoph Jung and Frank Wolter. Living without beth and craig: Definitions and interpolants in the guarded and two-variable fragments. In *LICS 2021*, pages 1–14. IEEE, 2021.
- [Lakemeyer and Levesque, 2002] Gerhard Lakemeyer and Hector J. Levesque. Evaluation-based reasoning with disjunctive information in first-order knowledge bases. In *Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2002)*, pages 73–81. AAAI Press, 2002.
- [Lakemeyer and Levesque, 2014] Gerhard Lakemeyer and Hector J. Levesque. Decidable reasoning in a fragment of the epistemic situation calculus. In Thomas Eiter, Chitta Baral, and Giuseppe De Giacomo, editors, *Proceedings of the Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2014)*, pages 468–477. AAAI Press, 2014.
- [Levesque *et al.*, 1998] Hector J. Levesque, Fiora Pirri, and Raymond Reiter. Foundations for the situation calculus. *Electron. Trans. Artif. Intell.*, 2:159–178, 1998.
- [Lin and Reiter, 1994] Fangzhen Lin and Ray Reiter. Forget it! In *Working Notes 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.
- [Liu and Claßen, 2024] Daxin Liu and Jens Claßen. First-order progression beyond local-effect and normal actions. In Kate Larson, editor, *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 3475–3483. ijcai.org, 2024.
- [Liu and Lakemeyer, 2009] Yongmei Liu and Gerhard Lakemeyer. On first-order definability and computability of progression for local-effect actions and beyond. In Craig Boutilier, editor, *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 860–866. AAAI Press, 2009.
- [Liu and Levesque, 2005] Yongmei Liu and Hector J. Levesque. Tractable reasoning in first-order knowledge bases with disjunctive information. In Manuela M. Veloso and Subbarao Kambhampati, editors, *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI 2005)*, pages 639–644. AAAI Press, 2005.
- [Lutz *et al.*, 2001] Carsten Lutz, Ulrike Sattler, and Frank Wolter. Description logics and the two-variable fragment. In Carole A. Goble, Deborah L. McGuinness, Ralf Möller, and Peter F. Patel-Schneider, editors, *Working Notes of the 2001 International Description Logics Workshop (DL-2001)*, volume 49 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2001.
- [McCarthy and Hayes, 1969] John McCarthy and Patrick Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. American Elsevier, New York, 1969.
- [Reiter, 2001] Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, 2001.
- [Vassos and Levesque, 2013] Stavros Vassos and Hector J. Levesque. How to progress a database III. *Artificial Intelligence*, 195:203–221, 2013.
- [Zarrieß and Claßen, 2016] Benjamin Zarrieß and Jens Claßen. Decidable verification of Golog programs over non-local effect actions. In Dale Schuurmans and Michael Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016)*, pages 1109–1115. AAAI Press, 2016.
- [Zarrieß, 2018] Benjamin Zarrieß. *Verification of Golog Programs over Description Logic Actions*. PhD thesis, TU Dresden, 2018.