Progression with Probabilities in the Situation Calculus: Representation and Succinctness

Daxin Liu
The University of Edinburgh
Edinburgh, United Kingdom
daxin.liu@ed.ac.uk

Vaishak Belle
The University of Edinburgh
Edinburgh, United Kingdom
vbelle@ed.ac.uk

ABSTRACT

Progression in the Situation Calculus is perhaps one of the most extensively studied cases of updating logical theories over a sequence of actions. While it generally requires second-order logic, several useful first-order and tractable cases have been identified. Recently, there has been an interest in studying the progression of probabilistic knowledge bases expressed using degrees of belief on first-order formulas. However, although a few results exist, they do not provide much clarity about how this progression can be computed or represented in a feasible manner.

In this paper, we address this problem for the first time. We first examine the progression of a probabilistic knowledge base (PKB) in a world-level representation; in particular, we show that such a representation is closed under progression for any local-effect actions with quantifier-free contexts. We also propose a more succinct representation of the probabilistic knowledge base, i.e. factored-representation PKB. For this type of PKB, we study the conditions for progression to remain succinct.

KEYWORDS

knowledge representation; probabilistic progression; reasoning about action

ACM Reference Format:
Daxin Liu and Vaishak Belle. 2024. Progression with Probabilities in the Situation Calculus: Representation and Succinctness. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

Reiter's [27] reconsideration of the situation calculus [23] has proven enormously useful for the design of logical agents, essentially paving the way for cognitive robotics [13]. Among other things, it incorporates a simple monotonic solution to the frame problem, leading Reiter to define the notion of regression for basic action theories [38]. But for long-lived agents like robots, Lin and Reiter [15] argue that the notion of progression, that of continually updating the current view of the state of the world, is perhaps better suited. They show that progression is always second-order definable, and in general, it appears that second-order logic is unavoidable [36]. However, Lin and Reiter also identify some first-order definable cases by syntactically restricting situation calculus basic action theories, and since then, many other special cases have been studied [21].

While Lin and Reiter intended to use their work on robots, one criticism leveled at their work, and indeed at much of the work in cognitive robotics, is that the theory is far removed from the kind of probabilistic uncertainty and noise seen in typical robotic applications. What exactly filtering mechanisms (such as Kalman filters [35]) have to do with Lin and Reiter's progression was unclear, until recently. Using extensions to the situation calculus for degrees of belief, likelihood, and noise, a couple of results have been established on studying the progression of situation calculus theories in the presence of probabilities. An early result by Belle and Lakemeyer [3] considered probabilistic progression under the assumption that there is full knowledge about the context formulas in successor state axioms and showed that the values of literals can be updated by evaluating this context. In later work, Belle and Levesque [6] showed that progression can be defined for finitely many nullary fluents by directly modifying the weight function on the possible worlds. More recently, Liu and Feng [17, 18] provided a more general account of Belle and Lakemeyer [3] and showed that even when the context formula is not known, a characterization of the updated values of fluents can be provided.

Although these results are interesting and involved, what is missing from this literature is that they do not provide much clarity on how the progression can be computed or represented feasibly. Closest in spirit, Belle and Levesque [5] consider updating world representations directly. That is, for finitely many nullary fluents, they instantiate worlds as vectors of fluent values, and update them after actions. Because they considered uncountably many alternatives for fluent values and noisy outcomes, they provide a *limit*-based semantics for how projection can be evaluated. However, this semantics is embedded in an account of program termination, making it difficult to understand how it works.

In this paper, we comprehensively address this problem for the first time. We first examine the progression of a probabilistic knowledge base (PKB) in a world-level representation that explicitly specifies the probability of each world; in particular, we show that such a representation is closed under progression for any local-effect actions with quantifier-free contexts. This result establishes how one might "implement" probabilistic progression in a direct and simple manner: one simply has to instantiate the worlds (say as a vector of fluent values as seen in [5]) and update them as actions happen in the manner indicated by our result.

However, the world-level representation is clearly exponential in the vocabulary. A more compact and attractive representation is to use graphical models. Graphical models such as Bayesian

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

networks and Markov networks are extremely popular in the machine learning area and their relational extensions continue to be formalisms that are studied in areas such as statistical relational learning and neuro-symbolic AI [1, 9, 11, 25, 28]. In this paper, we consider one such compact factored representation of PKB, i.e. FR-PKB, that factorizes according to a chordal graph. We study how, under noise-free actions, probabilistic progression is closed in this representation. That is, it is possible to get the progressed knowledge base factorized according to another chordal graph, leading to a compact progressed result (note that we can always capture the progressed knowledge base using the world-level representation, but the point is to ensure that the compactness of representation is not lost). This result can be seen to motivate the updating of probabilistic databases [31, 33]. Lastly, we discuss the challenges of extending our approach to handle noisy actions.

2 PRELIMINARY

In this section, we review the logic $\mathcal{DS}_p$, a probabilistic modal logic of actions and beliefs, and the classical categorical progression.

2.1 The Logic $\mathcal{DS}_p$

$\mathcal{DS}_p$ is a many-sorted language and, for simplicity, we only have two sorts: *object* and *action*. We assume rational numbers are in a sub-sort of object. The language features a fixed countable domain with the unique name assumption: the set of *standard names* $\mathcal{N} = \mathcal{N}_{obj} \cup \mathcal{N}_{act}$ where $\mathcal{N}_{obj}$ are object standard names and $\mathcal{N}_{act}$ are action standard names.

Syntax Formally, the vocabulary consists of

- first-order (FO) variables: $\{x, y, \ldots, u, v \ldots\}$;
- rigid function symbols such as *move*, *put*, etc.;
- fluent predicates symbols $F, F', \ldots$ like *On*, *Clear* including
  - a binary predicate *oi* to denote that two actions are indistinguishable from the agent's viewpoint; and
  - a special binary predicate l that takes an action as its first argument and the action's likelihood as its second argument.[1]
- connectives and other symbols: $=, \wedge, \neg, \forall, \boldsymbol{B}, \Box$, round and square parentheses, period, comma.

For simplicity, we do not include rigid predicates and fluent functions. We treat $\{\vee, \exists, \equiv, \supset\}$ as syntactic abbreviations and use $\Leftrightarrow$ to mean the equivalence of formulas. The *terms* of the languages are the least set of expressions such that (1) every FO variable and standard name (or constant) is a term; (2) if $t_1, \ldots t_k$ are terms and f a k-ary function, then $f(t_1, \ldots t_k)$ is a term. *Ground terms* are terms without free variables and *primitive terms* $\mathcal{P}_t$ are terms of the form $f(n_1 \ldots n_m)$, where $n_i \in \mathcal{N}$. Additionally, we assume $\mathcal{N}_a$ is just the set of action primitive terms. We denote $\mathcal{P}$ the set of all ground atoms. *Well-formed formulas* are constructed as usual in FOL logic with equality. They can further be in the context of modalities. The logic has an epistemic modality: $\boldsymbol{B}(\alpha\colon x)$ is to be read as "α is believed with a probability x" where x is a rational number. We use $\boldsymbol{K}(\alpha)$ as abbreviation for $\boldsymbol{B}(\alpha\colon 1)$, read as "$\alpha$ is known". There are two action modalities $[a]$, $\Box$ in that if α is a formula, then so are $[a]\alpha$ (read:"α holds after action a") and $\Box\alpha$ (read: "α holds after any sequence of actions"). For $z = a_1 \cdots a_k$, we write $[z]\alpha$ to mean $[a_1] \cdots [a_k]\alpha$. We use True to denote truth, which is taken as an abbreviation for, say, $\forall x(x = x)$, and False for its negation. For α, we use α^x_t to denote the formula obtained by substituting free variable x in α with term t. A formula without $\boldsymbol{B}$ is called *objective*; a formula without $[a], \Box$ is called *static*; a formula without fluents, $[a], \Box$ outside $\boldsymbol{B}$ is called *subjective*; and a formula without modalities is called *fluent formula* or *fact*.

Semantics The semantics is given in terms of possible worlds. In a dynamic setting, such worlds are defined to interpret not only the current state of affairs but also how that changes with actions. Let $\mathcal{Z} = (\mathcal{N}_a)^*$ be the set of all finite sequences of actions including $\langle\rangle$, the empty sequence. Then a *world* maps $\mathcal{P} \times \mathcal{Z}$ to $\{0, 1\}$. We denote $\mathcal{W}$ the set of all possible worlds. We require that $\forall w \in \mathcal{W}$

- *oi* is an equivalence relation (reflexive, symmetric, and transitive) for all z and *oi* is rigid; [2]
- $l(a, u)$ behaves like a function in all worlds: for all action a, action sequence z, there exists exactly one number n s.t. $w[l(a, n), z] = 1$ and for all $n' \neq n$, $w[l(a, n'), z] = 0$.

An *epistemic state* e is a set of distributions d (weighted functions) that maps $\mathcal{W}$ to $\mathbb{R}^{\geq 0}$. By a model, we mean a triple $\langle e, w, z\rangle$.

Truth for objective sentences is given as:

- $e, w, z \models t_1 = t_2$ iff t_1, t_2 are identical;
- $e, w, z \models p$ iff p is an atom and $w[p, z] = 1$;
- $e, w, z \models \neg\alpha$ iff $e, w, z \not\models \alpha$;
- $e, w, z \models \alpha \wedge \beta$ iff $e, w, z \models \alpha$ and $e, w, z \models \beta$;
- $e, w, z \models \forall x.\alpha$ iff $e, w, z \models \alpha^x_n$ for all $n \in \mathcal{N}$ for the right sort;
- $e, w, z \models [r]\alpha$ iff $e, w, z \cdot r \models \alpha$;
- $e, w, z \models \Box\alpha$ iff $e, w, z \cdot z' \models \alpha$ for all $z' \in \mathcal{Z}$.

To account for stochastic actions, $\mathcal{DS}_p$ uses a notion called *observational indistinguishability* among actions. The idea is that instead of saying stochastic actions have non-deterministic effects, $\mathcal{DS}_p$ says stochastic actions have non-deterministic alternatives which are mutually observationally indistinguishable from the agent's perspective and each of which has a deterministic effect. E.g., to express that the robot's forward action might forward successfully or unsuccessfully, $\mathcal{DS}_p$ uses a formula $oi(fwd(1), fwd(0))$, i.e. the outcomes $\{1, 0\}$ are observationally indistinguishable to the agent (unless sensing is performed). For action sequences, we have:

Definition 2.1. Given a world w, we define $z \approx_w z'$:

(1) $\langle\rangle \approx_w z'$ iff $z' = \langle\rangle$;
(2) $z \cdot a \approx_w z'$ iff $z' = z^* \cdot a^*$, $z \approx_w z^*$ and $w[oi(a, a^*), z] = 1$.

For example, we have $fwd(0) \cdot fwd(1) \approx_w fwd(1) \cdot fwd(0)$. Lastly, we define the likelihood of action sequences in a world:

Definition 2.2. We define $l^* : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}^{\geq 0}$:

- $l^*(w, \langle\rangle) = 1$ for all $w \in \mathcal{W}$;
- $l^*(w, z \cdot a) = l^*(w, z) \times n$ where $w[l(a, n), z] = 1$.

Given e, w and formula α, let $\|\alpha\|_e := \{w' : e, w', \langle\rangle \models \alpha\}$. Intuitively, $\|\alpha\|_e$ is the set of all alternative worlds that might result in α under e. For a distribution d, we define Norm$(d, \|\alpha\|_{\{d\}}, n)$ if $n = \frac{1}{\eta} \times \sum_{\{w' : \|\alpha\|_{\{d\}}\}} d(w')$, where η is a normalizer with the same expression as the numerator but replacing α to True, i.e. the

[1] For simplicity, we omit the usual $Poss(a)$ fluent for action preconditions. Meanwhile, impossible actions are treated as actions with 0 likelihood.

[2] Allowing truth of oi to vary might cause counter-intuitive result, see [20].

set of worlds where α holds has proportioned summed weights (or probability) n in distribution d. Essentially, although $\mathcal{W}$ is uncountable, the NORM here requires d to be a well-defined discrete distribution over possible worlds. See [4] for details. A distribution is called *regular* if $\textsc{Norm}(d, \|\textsc{True}\|_{\{d\}}, 1)$. We denote $\mathcal{D}$ as the set of all regular distributions. Henceforth, we restrict ourselves to regular distributions.

Definition 2.3. Given $w \in \mathcal{W}, z \in \mathcal{Z}, d \in \mathcal{D}$, we define

- w_z as a world s.t. for each atom p and $z' \in \mathcal{Z}$, $w_z[p, z'] = w[p, z \cdot z']$;
- d_z as a distribution s.t. for all $w \in \mathcal{W}$, $d_z(w) = \sum_{\{w':d(w')>0\}} \sum_{\{z':z' \approx_{w'} z, w'_{z'}=w\}} d(w') \times l^*(w', z')$.

w_z is called the *progressed world* of w wrt z. Since worlds are tree-structured, the progressed world w_z of a world w wrt an action sequence z is just a copy of the sub-tree of w starting after the action sequence z, i.e. *forgetting* the past. d_z is called the *progressed distribution* of d wrt z and it is obtained by shifting the weight of worlds according to the actions' likelihood.

Given e, z, let $e_z = \{d_z \in \mathcal{D} : d \in e\}$. Now, we are ready to give truth for $\boldsymbol{B}$. Supposing r is a number,

- $e, w, z \models \boldsymbol{B}(\alpha : r)$ iff for all $d \in e_z$, $\textsc{Norm}(d, \|\alpha\|_{\{d\}}, r)$;

For a sentence α, we write $e, w \models \alpha$ to mean $e, w, \langle\rangle \models \alpha$. If Σ is a set of sentences and α is a sentence, we write $\Sigma \models \alpha$ (read: Σ logically entails α) to mean that for all e and w, if $e, w \models \alpha'$ for every $\alpha' \in \Sigma$, then $e, w \models \alpha$. Satisfiability and validity are defined in the usual way. If α is objective, we write $w \models \alpha$ instead of $e, w \models \alpha$. Similarly, we write $e \models \alpha$ instead of $e, w \models \alpha$ if α is subjective.

2.2 Classical Progression

To express the dynamics of a domain, $\mathcal{DS}_p$ uses a notion of basic action theories (BATs) like the situation calculus [27].

Definition 2.4 (Basic Action Theory). Given a set of fluents $\mathcal{F}$, a BAT Σ consists of:

- The initial theory Σ_0, which is a set of fluent sentences;
- Successor-state axioms (SSAs) Σ_{post}, which is a set of axioms of the form $\square[a]F(\vec{x}) \equiv \gamma_F^+(\vec{x}, a) \lor F(\vec{x}) \land \neg\gamma_F^-(\vec{x}, a)$, one for each fluent $F \in \mathcal{F}$.[3]

γ_F^+ and γ_F^- are the *positive and negative effect conditions* respectively. Intuitively, the initial theory expresses what holds initially, while successor-state axioms describe how actions affect fluents.

Definition 2.5 (Progression). Let $\Sigma_0 \cup \Sigma_{post}$ be a basic action theory. A set of fluent sentences Σ_0' is a progression of Σ_0 via a ground action t wrt Σ_{post} iff for every sentence ϕ

$$\Sigma_0 \cup \Sigma_{post} \models [t]\phi \quad \text{iff} \quad \Sigma_0' \cup \Sigma_{post} \models \phi$$

The definition follows from [37]. Unfortunately, progression is in general only second-order (SO) definable [15, 36], i.e. Σ_0' might not always be first-order. $\mathcal{DS}_p$ can be easily extended to include SO quantifiers and variables in the language which is then interpreted by variable maps. This means that the model has to be expanded with variable maps and the satisfaction relationship needs to be modified accordingly. We omit it here and refer the interested readers to [17, 18] on how this can be achieved.

THEOREM 2.6 ([15]). *The following is a progression of Σ_0 wrt a ground action t and Σ_{post}:*

$$\exists\vec{V}.(\Sigma_0)^{\vec{F}}_{\vec{V}} \land \bigwedge_i \forall\vec{x}.F_i(\vec{x}) \equiv (\gamma_{F_i}(\vec{x}, t))^{\vec{F}}_{\vec{V}} \tag{1}$$

where $\vec{V}$ are SO variables and γ_{F_i} is the RHS of the SSA of fluent F_i.

Here the super- and subscript stands for substitution, just like the first-order case.

Progression is closely related to the concept of *forgetting*. 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). [14] provided a semantic characterization of forgetting, which is essentially captured by our notion of progressed worlds. Syntactically, forgetting a predicate P in a sentence ϕ, i.e. *forget*(ϕ, P), is equivalent to $\exists V.\phi^P_V$ while forgetting a ground atom $P(\vec{t})$ in ϕ, i.e. *forget*$(\phi, P(\vec{t}))$ is more complex. Let $\phi[P(\vec{t})]$ be the formula obtained by replacing every occurrence of the form $P(\vec{t}')$ in ϕ with $[\vec{t} = \vec{t}' \land P(\vec{t})] \lor [\vec{t} \neq \vec{t}' \land P(\vec{t}')]$, $\phi_+^{P(\vec{t})}$ and $\phi_-^{P(\vec{t})}$ be formulas obtained by replacing $P(\vec{t})$ in $\phi[P(\vec{t})]$ with TRUE and FALSE respectively. Then we have *forget*$(\phi, P(\vec{t})) = \phi_+^{P(\vec{t})} \lor \phi_-^{P(\vec{t})}$. E.g., let $\phi := eh(A) \land \neg eh(C)$, then we have *forget*$(\phi, eh(A)) = \neg eh(C)$. This result naturally extends to forgetting a finite set of atoms Ω, that is, assume $p \in \Omega$, we define *forget*$(\phi, \Omega) :=$ *forget*$($*forget*$(\phi, p), \Omega\backslash\{p\})$. With these notions, the progression of an initial theory Σ_0 wrt t and Σ_{post} is just to add the effects of t (the big conjunction in Eq. (1)) and forget the past (i.e. the SO existential quantifiers in the head of Eq. (1)).

Since then, efforts have been made to find fragments where progression is first-order. Lin and Reiter [15] showed that progression is FO definable for *strongly context-free* SSAs which essentially capture the STRIPS operators. This is further extended [21] to *local-effect* SSAs. Intuitively, local-effect means that fluents affected by the action are determined by the action's parameters. More formally,

Definition 2.7. A SSA is called

- *strongly context-free* if for all fluent F, γ_F^+ and γ_F^- are disjunctions of formulas of the form $\exists\vec{u}[a = A(\vec{v})]$ where $\vec{v}$ contains $\vec{x}$ and $\vec{u}$;
- *local-effect* if for all fluent F, γ_F^+ and γ_F^- are disjunctions of formulas of the form $\exists\vec{u}[a = A(\vec{v}) \land \phi(\vec{v})]$ where $\vec{v}$ contains $\vec{x}$ and $\vec{u}$ is the set of remaining variables, $\phi(\vec{v})$ is the context formula;

Clearly, strongly context-free SSAs (SCF-SSAs) are special cases of local-effect SSAs (LE-SSAs) without context formulas. The result in [21] to progress local-effect action theories is based on two observations. First, if a sentence ϕ entails that two predicates P, Q agree on everything except on a finite set of instances Δ, then forgetting Q in ϕ amounts to forgetting the set of Q atoms instantiated by Δ in ϕ and then replacing Q by P. That is:

THEOREM 2.8 ([21]). *Let $\Delta = \{\vec{t}_1, \ldots, \vec{t}_n\}$ be a set of vectors of ground terms, P, Q two predicates, $Q(\Delta)$ be the set $\{Q(\vec{t}) \mid t \in \Delta\}$, and $P \approx_\Delta Q$ denote the formula $\forall\vec{x}. \notin \Delta \supset P(\vec{x}) \equiv Q(\vec{x})$. Then* *forget*$(\phi \land P \approx_\Delta Q, Q) \Leftrightarrow$ *forget*$(\phi, Q(\Delta))^Q_P$.

[3] Free variables are implicitly universally quantified from the outside. The $\square$ modality has lower syntactic precedence than the connectives, and $[\cdot]$ has the highest priority.

Second, after instantiating a LE-SSA of fluent F with a ground action $A(\vec{n})$, γ_F^+ (likewise for γ_F^-) is equivalent to disjunctions of the form $\vec{x} = \vec{m} \wedge \phi(\vec{n})$ where $\vec{m} \subseteq \vec{n}$. That is only finitely many instances of $F(\vec{x})$ are affected (those $\vec{x} = \vec{m}$). More formally, let $\Delta_F = \{\vec{m} : \vec{x} = \vec{m} \text{ appears in } \gamma_F(\vec{x}, A(\vec{n}))\}$, then the big conjunction part in Eq.(1) (for fluent F_i) amounts to the conjunction of

$$\forall \vec{x}.\vec{x} \notin \Delta_{F_i} \supset F_i(\vec{x}) \equiv V_i(\vec{x}) \tag{2a}$$

$$\bigwedge_{\vec{m} \in \Delta_{F_i}} F_i(\vec{m}) \equiv (\gamma_{F_i}(\vec{m}, t))^{\vec{F}}_{\vec{V}} \tag{2b}$$

Hence one could use Eq.(2a) and Theorem 2.8 to eliminate the second-order quantifier in the result of Theorem 2.6. Formally, let Ω be the *characteristic set* of action $A(\vec{n})$, i.e. the set of atoms given by: $\Omega = \{F(\vec{t}) : F \text{ is a fluent, and } \vec{t} \in \Delta_F\}$. For each fluent F we introduce a new predicate Q.[4]

Theorem 2.9 ([21]). *If Σ_{post} is local-effect, then the following FO sentence is a progression of Σ_0 wrt t and Σ_{post},*

$$forget(\Sigma_0 \wedge \bigwedge_{F_i(\vec{m}) \in \Omega} Q_i(\vec{m}) \equiv \gamma_{F_i}(\vec{m}, t), \Omega)^{\vec{Q}}_{\vec{F}}. \tag{3}$$

Example 2.10. Consider the following blocks world. A single action $move(x, y, z)$ can move a block x from block y to block z. We use a fluent $eh(x)$ to indicate if the height of block x is an even number. Every block is assumed to have a unit height. $move(x, y, z)$ causes the height of x to be even if and only if the height of z is odd. The following LE-SSA can express this:

$$\begin{aligned}\Box[a]eh(x) \equiv\ & \exists y, z.[a = move(x, y, z) \wedge \neg eh(z)] \vee \\ & \neg[\exists y, z.a = move(x, y, z) \wedge eh(z)] \wedge eh(x)\end{aligned} \tag{4a}$$

Let $\Sigma_0 = \{\neg eh(A)\}$. For the ground action $move(A, B, C)$, we have $\Omega = \{eh(A)\}$. Let Q be a new predicate for eh. Then Eq. (3) equals to $forget(\neg eh(A) \wedge Q(A) \equiv eh(c), \Omega)^{Q}_{eh}$, which is equivalent to $eh(A) \equiv \neg eh(C)$. That is, the height of A is even if and only the height of C is not even.

3 PROGRESSING WR-PKB

To perform probabilistic progression, one needs to specify the likelihood of each outcome of stochastic action. $\mathcal{DS}_p$ achieves this by including an axiom of observational indistinguishability among actions and a likelihood axiom in the BATs. The observational indistinguishability axiom Σ_{oi} specifies the equivalence relations among actions. In this paper, we focus on Σ_{oi} where each action has finitely many alternatives, that is, Σ_{oi} is of the form $\Box oi(a, a') \equiv \bigvee_i a \in \mathcal{N}_i \wedge a' \in \mathcal{N}_i$ where $\mathcal{N}_i \subset \mathcal{N}_a$ are finite. Accordingly, the likelihood axiom (LA) Σ_l is as $\Box l(a, u) \equiv \bigvee_i \bigvee_{t_i \in \mathcal{N}_i} a = t_i \wedge u = q_{t_i}$ subject to $\sum_{t_i \in \mathcal{N}_i} q_{t_i} = 1$ for all i.

3.1 Representing Probabilistic Knowledge Base

Like classical progression where one needs to use a theory to describe the world state, we need to specify what is believed by the agent, namely the probabilistic knowledge base (PKB). By a *probabilistic knowledge base*, we mean a set of static subjective fluent formulas without nested $\boldsymbol{B}$. Clearly, different syntactic constraints can be imposed on formulas to obtain different types of PKB.

While $\mathcal{DS}_p$ allows PKBs with incomplete beliefs like $\boldsymbol{B}(p_1 : r_1) \vee \boldsymbol{B}(p_2 : r_2)$ or beliefs involving infinitely many random variables like $\forall x.\boldsymbol{B}(\phi(x) : 0.5)$, we only consider PKB that can be characterized by a joint distribution of finitely many binary random variables and leave PKBs with unknown distributions for future study.

Fact-Independent PKB. A natural and also compact way to represent a PKB is using a theory $T = \{\boldsymbol{B}(\alpha_1 : n_1), \ldots, \boldsymbol{B}(\alpha_k : n_k)\}$ to express that fact α_i hold with probability n_i, together with a theory T_{ind} saying that α_i are mutually independent. We called such PKBs *fact-independent* PKBs (FI-PKB for short). In this paper, we only focus on PKB over atomic facts, i.e. α_i are ground atoms. Although the size of T_{ind} might be exponential in k, in practice T_{ind} is usually omitted and is only implicitly assumed in reasoning about probability. Let X_{α_i} be a binary random variables for fact α_i, clearly, $T \cup T_{ind}$ defines a joint distribution $Pr(\cdot)$ over $X_{\alpha_1}, \ldots X_{\alpha_k}$. In fact, FI-PKBs correspond to the notion of the *tuple-independent* probabilistic database [33]. Unfortunately, this representation is not closed under progression: the progression of a fact-independent PKB might no longer be a fact-independent PKB as actions would introduce new dependencies among facts.

World-Representation PKB. Alternatively, one could list out the probabilities of all boolean combinations of facts α_i by using formulas of the form $\boldsymbol{B}(\beta_1 \wedge \beta_2 \wedge \ldots \beta_k : r)$ where $\beta_i \in \{\alpha_i, \neg\alpha_i\}$. Such a boolean combination essentially represents a truth assignment θ for atomic facts AF, i.e. $\theta = \beta_1 \wedge \beta_2 \wedge \ldots \beta_k$. We denote the set of all assignments as Θ and call such a PKB *world-representation* PKB, (WR-PKB for short). That is, a WR-PKB T is a set $\{\boldsymbol{B}(\theta : r_\theta) | \theta \in \Theta\}$. Clearly, the set of formulas has to be consistent: $\sum_{\theta \in \Theta} r_\theta = 1$. WR-PKB is not compact as the size of Θ is exponential to the number of facts. Below is a WR-PKB over three facts $\{eh(A), eh(B), eh(C)\}$.

	$eh(A)$	$eh(B)$	$eh(C)$	r_θ
θ_1	0	0	0	0.378
θ_2	0	0	1	0.162
θ_3	0	1	0	0.018
θ_4	0	1	1	0.042
θ_5	1	0	0	0.028
θ_6	1	0	1	0.012
θ_7	1	1	0	0.108
θ_8	1	1	1	0.252

(5)

3.2 Progressing WR-PKB

Definition 3.1. Given a WR-PKB T over facts AF $= \{\alpha_1, \ldots, \alpha_k\}$, a BATs $\Sigma = \Sigma_{post} \cup \Sigma_l \cup \Sigma_{oi}$, and a stochastic action t, another WR-PKB T' (over AF) is said to be a progression of T wrt Σ and t if

$$T \wedge \boldsymbol{K}\Sigma \models [t]T'.$$

Note that, this definition coincides with Def. 2.5 as T' completely specifies the joint distribution. Intuitively, given a WR-PKB that represents an initial joint distribution, the progression computes another one according to the effects and likelihoods of the action.

Unfortunately, given a PKB, the effects of a stochastic action might not always be captured by another PKB (over finitely many random variables). This is because PKB is essentially propositional,

[4] The purpose of new predicates is to ensure the LHS of Eq. (2b), i.e. $F_i(\vec{m})$ won't be affected when forgetting the characteristic set Ω.

yet the progression of a propositional theory Σ_0 might contain quantifiers even second-order quantifiers as shown by Theorem 2.6. Intuitively, given an initial PKB over a finite set of facts, an action might change the truth of infinitely many facts according to the SSAs, hence the result cannot be captured by another finite PKB. An interesting question is then what conditions are required to ensure that the progression of a PKB is another PKB? Perhaps, a naive but also natural way is to require the domain to be finite by using an axiom like $\forall x.x = n_1 \vee x = n_2 \ldots \vee x = n_k$. Yet such a formula is unsatisfiable in $\mathcal{DS}_p$ as the domain is infinite. Of course, one could fake such a finite domain in $\mathcal{DS}_p$, for example, by assuming that a predicate P has finitely many instances and all quantifiers are quantifying over its finite instance, yet again, this corresponds to include the *closed-worlds assumption* [26] as a premise. What we are after here are conditions that the progression of a PKB is *closed* in an open-world setting. Below, we show that WR-PKB is closed under progression for *quantifier-free* LE-SSAs (QF-LE-SSAs for short).

Definition 3.2. A set of LE-SSAs is said to be essentially quantifier-free if all the context formulas $\phi(\vec{v})$ are quantifier-free.

SCF-SSAs are all essentially quantifier-free as their contexts amount to TRUE. We denote the progression of a initial theory Σ_0 wrt a LE-SSAs Σ_{post} and a ground action t_a given by Theorem 2.9 as $Prog(\Sigma_0, \Sigma_{post}, t_a)$. Clearly, $Prog(\Sigma_0, \Sigma_{post}, t_a)$ is propositional if Σ_0 is propositional and Σ_{post} is quantifier-free.

Given a BAT Σ with QF-LE-SSAs, let $t = A(\vec{n})$ be a stochastic action, w.l.o.g, we assume the set of alternatives of t under Σ_{oi} is given by $\mathcal{N}_t = \{A_1(\vec{n}_1), \ldots, A_k(\vec{n}_k)\}$ (also $t \in \mathcal{N}_t$). Let Ω_{t_i} be the characteristic sets of action $t_i \in \mathcal{N}_t$, Φ_{t_i} be the set of atoms that appears in the context formulas $\phi(\vec{n}_i)$. We call the stochastic action t *progressable* wrt a WR-PKB T and BAT Σ if $\Omega_{t_i} \cup \Phi_{t_i} \subseteq \mathsf{AF}$ for all $t_i \in \mathcal{N}_t$, that is, facts that the stochastic actions t_a will have effects on and condition on are believed with a certain degree in the PKB T.

In Example 2.10, suppose the ground action $t = move(A, B, C)$ has an alternative $t' = moveFailed$ that affects nothing, additionally, the two actions have equal likelihoods (let $\mathcal{N}_{move} = \{t, t'\}$, then the following Σ_{oi} ad Σ_l can express this), we have $\Omega_t = \{eh(A)\}$ and $\Phi_t = \{eh(C)\}$.

$$\Box oi(a, a') \equiv a \in \mathcal{N}_{move} \wedge a' \in \mathcal{N}_{move} \tag{6a}$$

$$\Box l(a, u) \equiv a \in \mathcal{N}_{move} \wedge u = 0.5 \tag{6b}$$

Therefore, t is progressable for a WR-PKB over facts containing $\{eh(A), eh(C)\}$.

THEOREM 3.3. *Given a WR-PKB T over facts $\mathsf{AF} = \{\alpha_1, \ldots, \alpha_k\}$, a BAT with QF-LE-SSAs Σ and a progressable stochastic action t (under T and Σ), the following WR-PKB T' over AF is a progression of T wrt Σ and t:*

$$T' = \{\boldsymbol{B}(\theta : r_\theta) | r_\theta = \sum_{\theta' \in \Theta} \sum_{t' \in \mathcal{N}_t} r'_{\theta'} \times q_{t'} \times \mathbb{I}(\theta, Prog(\theta', \Sigma_{post}, t'))\}$$

where $\boldsymbol{B}(\theta' : r'_{\theta'}) \in T$, $q_{t'}$ is the likelihood of t', and $\mathbb{I}(\cdot)$ is an indicator function given by:

$$\mathbb{I}(\theta, Prog(\theta', \Sigma_{post}, t')) = \begin{cases} 1 & \theta \Leftrightarrow Prog(\theta', \Sigma_{post}, t') \\ 0 & \textit{otherwise.} \end{cases}$$

The proof is obtained by the definition of progressed distribution and progressed world. Intuitively, to compute the new degree of belief of an assignment θ, one needs to consider all the assignments θ' and action alternatives t' such that θ' might progress to θ via action t', the result is just to sum the product of the degrees of belief of θ' (in T) and the likelihood of t'. E.g., let $T = \{\boldsymbol{B}(\neg eh(A) \wedge \neg eh(C) : 1)\}$, then $T \wedge \boldsymbol{K}\Sigma \models [move(A, B, C)]T'$ where Σ consists of Eq.(4a),Eq.(6a), and Eq.(6b), and T' is as (ignore the trivial beliefs):

$$T' = \{\boldsymbol{B}(\neg eh(A) \wedge \neg eh(C) : 0.5), \boldsymbol{B}(eh(A) \wedge \neg eh(C) : 0.5)\}$$

This is because progressing $\neg eh(A) \wedge \neg eh(C)$ with the ground action $move(A, B, C)$ we obtain $forget(\neg eh(A) \wedge \neg eh(C) \wedge Q(A) \equiv eh(c), \Omega)^{Q}_{eh}$, which is equivalent to $eh(A) \wedge \neg eh(C)$, additionally, the likelihood of $move(A, B, C)$ is 0.5, therefore $eh(A) \wedge \neg eh(C)$ receives 0.5 degree of belief.

We comment that the notion of progressable action imposes both constraints on the BAT and the PKB, yet the constraints on the PKB are not about the type of the PKB but only instances. In the previous example, the action $move(A, B, C)$ is not progressable wrt a WR-PKB over a single fact $eh(A)$ due to lack of information for fact $eh(C)$, but it is progressable wrt any WR-PKB over facts containing $\{eh(A), eh(C)\}$. In fact, we have that

THEOREM 3.4. *Given a BAT Σ with QF-LE-SSAs and an action t, there is always a WR-PKB T such that t is progressable wrt T, Σ.*

In other words, WR-PKB is closed under progression for BATs with QF-LE-SSAs. The proof is straightforward: since every stochastic action t has finite alternatives N_t and for each $t_i \in N_t$, Ω_{t_i} and Φ_{t_i} are finite, their union is also finite.

4 PROGRESSING FR-PKB

The inability of WR-PKB to compactly represent a PKB motivates us to resort to probabilistic graphical models.

4.1 Factored-Representation PKB

Probabilistic graphical model is a powerful approach that allows compactly representing dependency patterns among random events [24]. The central idea of these approaches is the use of factored representations for modeling correlations. In what follows, we denote a random variable by X (possibly with super or subscripts), a set of random variables by $\mathbf{X}$, and the domain of X by $dom(X)$.

Definition 4.1. A factor $f(\mathbf{X})$ is a function of a (small) set of random variables $\mathbf{X} = \{X_1, \ldots, X_k\}$ s.t. $f(\mathbf{X} = \mathbf{x}) \in [0, 1]$ for all $\mathbf{x} \in dom(X_1) \times \ldots \times dom(X_k)$.

A factored representation of a joint distribution $Pr(\mathbf{X})$ allows it to be represented compactly as a product of factors $Pr(\mathbf{X} = \mathbf{x}) = \prod_{i=1}^{m} f_i(\mathbf{X}_i = \mathbf{x}_i)$ where $\mathbf{X}_i \subseteq \mathbf{X}$ is the set of random variables restricted to factor f_i and $\mathbf{x}_i$ is the corresponding assignment. Moreover, we say a joint distribution $Pr(\mathbf{X})$ *factorizes* (or is *factorizable*) according to an undirected graph $G = (V, E)$, if $V = \mathbf{X}$, and $Pr(\mathbf{X} = \mathbf{x}) = \prod_{c \in C} f_c(\mathbf{X}_c = \mathbf{x}_c)$ where C is the set of (maximal) cliques of G, $\mathbf{X}_c \subseteq \mathbf{X}$ is the set of random variables in clique c. For a clique c, we might use c to refer to the set of vertices in c. Since $\mathbf{X}_c$ is usually smaller than $\mathbf{X}$, representing $f_c(\mathbf{X}_c)$ is easier than $Pr(\mathbf{X})$ leading to a compact representation.

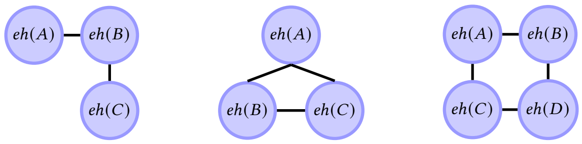

(a) (b) (c)

Figure 1: Three undirected graphs

In general, the factor functions $f_c(\mathbf{X}_c)$ (wrt a graph G) do not carry out explicit meanings and are hard to interpret. For this reason and for the purpose of representing PKB, we focus on the *decomposable graph*, viz *chordal graph*.

Definition 4.2 (Chordal graph). A graph G is chordal if every cycle of length four or more has at least one chord, i.e. an edge that is not part of the cycle but connects two vertices of the cycle.

The chordal graph is also known as *triangulated graph*, this is because every induced cycle in a chordal graph has exactly three vertices (triangle). For example, graphs in Fig. (1a) and (1b) are chordal graphs while the graph in Fig. (1c) is not.

For a chordal graph G and a pair of connected cliques c_i, c_j in G, we denote $s_{i,j}$ the *separator* of c_i and c_j as the maximal set of their common vertices, $\mathcal{S}$ the set of all separators in G, including any repetitions. For example, graph (1a) has two cliques $eh(A)$-$eh(B)$, $eh(B)$-$eh(C)$ and a separator $eh(B)$, graph (1b) has one clique (the whole graph) and no separators.

Theorem 4.3 ([24]). *If a joint distribution $Pr(\mathbf{X})$ is factorizable according to a chordal graph G, then*

$$Pr(\mathbf{X}=\mathbf{x}) = \frac{\prod_{c \in C} Pr(\mathbf{X}_c = \mathbf{x}_c)}{\prod_{s \in \mathcal{S}} Pr(\mathbf{X}_s = \mathbf{x}_s)} \tag{7}$$

where $Pr(\mathbf{X}_c = \mathbf{x}_c)$ is the marginal probability in $Pr(\mathbf{X})$, likewise for $Pr(\mathbf{X}_s = \mathbf{x}_s)$ and the RHS is understood as 0 whenever the denominator is 0.

Theorem 4.3 lies the theoretical foundation of our *factored representation PKB* (FR-PKB).

Definition 4.4. A factored representation PKB over facts AF $= \{\alpha_1, \ldots, \alpha_k\}$ is a pair $\langle G, T\rangle$, where $G = (V, E)$ is an undirected chordal graph, $V =$ AF, $T = \{\boldsymbol{B}(\theta\colon r_\theta) | \theta \in \Theta_c \text{ or } \theta \in \Theta_s, c \in C, s \in \mathcal{S})\}$. Here, C and $\mathcal{S}$ are the set of all cliques and separators of G, Θ_c and Θ_s are the set of all assignment over facts in c and s.

Essentially, T is the set of formulas that specifies the marginal probability of cliques and separators in G, while G specifies the correlations of the binary random variables for facts $\alpha_1, \ldots, \alpha_k$. By Theorem 4.3, a FR-PKB is logically equivalent to a WR-PKB. For example, the WR-PKB in Eq. (5) is equivalent to the FR-PKB $\langle G, T\rangle$ where G is as Fig. (1a) and T is given as Fig. (2).

Moreover, FR-PKB is a *complete representation* in the sense that any WR-PKB has an equivalent FR-PKB. Given a WR-PKB with sentences T, in the extreme case, we can construct an FR-PKB $\langle G, T\rangle$ where G is a complete graph. Clearly, an equivalent FR-PKB will be more compact if the chordal graph G is sparser and there is always a minimal G. In this paper, we will not insist G to be minimal. In this regard, fact-independent PKBs are special cases of FR-PKBs where the graph has no edges, i.e. vertices are mutually independent.

	$eh(B)$	r_θ
θ_0	0	0.58
θ_1	1	0.42

(a)

	$eh(A)$	$eh(B)$	r_θ
θ_2	0	0	0.54
θ_3	0	1	0.06
θ_4	1	0	0.04
θ_5	1	1	0.36

(b)

	$eh(B)$	$eh(C)$	r_θ
θ_6	0	0	0.406
θ_7	0	1	0.174
θ_8	1	0	0.126
θ_9	1	1	0.294

(c)

Figure 2: The theory of a FR-PKB over three facts

4.2 Progressing FR-PKB

The progression of FR-PKB can be defined similarly as Def.3.1.

Definition 4.5. Given a FR-PKB $\langle G, T\rangle$ over facts AF $= \{\alpha_1, \ldots, \alpha_k\}$, a BATs $\Sigma = \Sigma_{post} \cup \Sigma_l \cup \Sigma_{oi}$, and a stochastic action t, another FR-PKB $\langle G', T'\rangle$ (over AF) is a progression of T wrt Σ and t, written as $\langle G, T\rangle \wedge \boldsymbol{K}\Sigma \models [t]\langle G', T'\rangle$, if

$$T_{\text{WR-PKB}} \wedge \boldsymbol{K}\Sigma \models [t]T'_{\text{WR-PKB}}$$

where $T_{\text{WR-PKB}}$ and $T'_{\text{WR-PKB}}$ is the corresponding WR-PKBs of the FR-PKBs $\langle G, T\rangle$ and $\langle G', T'\rangle$.

Since every FR-PKB has a corresponding WR-PKB, we immediately obtain that FR-PKB is closed under progression for BATs with QF-LE-SSAs. The notion of progressable actions carries out naturally to FR-PKBs as well. Of course, for FR-PKBs, one could use Theorem 3.3 to derive a WR-PKB as a progression for QF-LE-SSAs and progressable actions, since all WR-PKBs are FR-PKB as well. Nevertheless, the graph of the corresponding FR-PKB of the progressed result tends to be fully connected, i.e. a complete graph, leading to a less compact representation. Hence, the interest here is to identify conditions where progression remains compact. In what follows, we show that for deterministic actions with SCF-SSAs, such compactness can be retained and the progressed graph can be obtained by some simple graph operations. Thereafter, we discuss challenges in extending the methods to general stochastic actions.

Deterministic Actions with SCF-SSAs. By a deterministic action, we mean a stochastic action that has only one alternative and likelihood 1. After such an action t, facts in Ω_t become either True or False, yet, this does not mean that a PKB over facts AF will degenerate to a deterministic KB as facts in AF$\backslash\Omega_t$ remains uncertain. Nevertheless, the observation that facts in Ω_t become independent of other facts provides a clue to build the new graph G'.

Given a FR-PKB $\langle G, T\rangle$ with $G = \langle V, E\rangle$, for the deterministic action t, we define a new graph $G' = \langle V', E'\rangle$ as $V' = V$ and

$$E' = \{\langle v_1, v_2\rangle \in E | v_1 \notin \Omega_t \text{ and } v_2 \notin \Omega_t.\}$$

Namely, G' is a graph that is generated from G by deleting all edges connecting a vertex in Ω_t. Although provably the graph G' is always chordal, the progressed WR-PKB $T'_{\text{WR-PKB}}$ might not be factorizable according to G'.

Theorem 4.6. *Given a FR-PKB $\langle G, T\rangle$ over facts* $\mathsf{AF} = \{\alpha_1, \ldots, \alpha_k\}$*, a BATs* Σ *with SCF-SSAs, and a progressable deterministic action* t*, let* $T_{\text{WR-PKB}}$ *be the corresponding WR-PKB of* $\langle G, T\rangle$ *and* $T'_{\text{WR-PKB}}$ *be the progression of* $T_{\text{WR-PKB}}$ *wrt* Σ *and* t*, additionally, let* G' *be a graph built from* G *as above, then* $T'_{\text{WR-PKB}}$ *is factorizable according to* G' *if* $\Omega_t \subseteq V\backslash S$*, where* V *is the vertex set of* G*,* S *is the set of all vertices in the separators of* G*.*

That is, if all the affected facts are not in any separator, i.e. $\Omega_t \subseteq V\backslash S$, then the progressed PKB is factorizable according to G'. In fact, under this assumption, the theory of progressed FR-PKB T' can be computed efficiently (linear in the size of T) by 3 steps.

(1) Adding $\boldsymbol{B}(\alpha\colon 1)$ or $\boldsymbol{B}(\alpha\colon 0)$ for $\alpha \in \Omega_t$ to T' according to if α is positively or negatively affected by t;
(2) For each clique c' (and separator s') in G' that was previously a clique c (separator s) in G, retain $\boldsymbol{B}(\theta\colon r_\theta)$ for $\theta \in \Theta_c$ ($\theta \in \Theta_s$) in T to T';
(3) For each clique $c' \in G'$ that can be obtained by deleting edges in a clique $c \in G$, add T' with $\{\boldsymbol{B}(\theta'\colon r_{\theta'}) | \theta' \in \Theta'_c\}$ where $r_{\theta'} = \sum_{\{\theta \in \Theta_c | \theta \text{ implies } \theta'\}} r_\theta$.

We mark the above process to derive G' and T' as DeleteEdges. Intuitively, the degree of belief of facts in Ω_t becomes either 1 or 0 depending on whether they are positively or negatively affected by t (Step 1); for facts in a clique or separator that was not affected by t, the degree of beliefs of any assignment of the facts remains unchanged (Step 2); For the assignment θ' in a clique c' of the new graph G' that is obtained by removing an affected fact in the clique $c \in G$, its degree of belief is obtained by summing over the degree of belief over assignments $\theta \in \Theta_c$ where θ can progress to θ' by t, i.e. θ implies θ' (Step 3).

The proof is straightforward. Intuitively, by Theorem 3.3, the effect of t is simply shifting the degree of beliefs over assignment θ that contains affected facts in Ω_t to a particular assignment (wrt Ω_t). Therefore, the degree of belief of the particular assignment amounts to the sum of the degree of belief of multiple assignments. Moreover, by Theorem 4.3, if the denominators in Eq. (7) are the same, the summation of multiple assignments can be pushed inside to the numerators, i.e. reducing to the summation of cliques that contain the affected facts. Additionally, such summation of cliques yields the factors of cliques in the new graph.

Example 4.7. Consider a variant of the blocks world where the action *putground*(x) can put block x to the ground and therefore set the height of x even, i.e. *eh*(x) is true. This can be specified by

$$\square[a]\,eh(x) \equiv a = putground(x) \vee eh(x) \tag{8}$$

Let $\langle G, T\rangle$ be a FR-PKB where G is given as Fig.(1a) and T as Fig. 2. For the ground action $t = putground(A)$, $\Omega_t = \{ehA\}$. The new graph G' will contain only one edge between the vertices of $eh(B), eh(C)$. Additionally, since the clique $eh(B) - eh(C)$ in G' is also a clique in G, Tuples in Fig. (2c) are retained in T', meanwhile $\{\boldsymbol{B}(eh(A)\colon 1)\}$ is added to T'.

Now, what if the affected facts are on separators? As mentioned before, in such case, the progressed PKB $T'_{\text{WR-PKB}}$ might not be factorizable to the graph G' constructed above. An example would be the block world with the ground action $t = putground(B)$. In this case, the new graph G' will have 0 edges. Any distribution that is factorizable wrt G' implies that $eh(A)$, $eh(B)$, and $eh(C)$ are independent, which is not the case for $T'_{\text{WR-PKB}}$. The reason is that facts that were not directly connected in a Markov net are merely *conditionally independent* but not independent. Moreover, the effect of a deterministic action t is just to make facts Ω_t independent, leaving alone the correlations among other facts. Hence, removing edges in a graph might undermine such correlations.

The idea is to add edges to G to generate a new chordal graph G_{sup} s.t. vertices of Ω_t is no longer in separators of G_{sup}. Thereafter, we could derive a progressed FR-PKB of $\langle G, T\rangle$ according to G_{sup} as before. Nevertheless, the challenge here is how to ensure that PKB $\langle G, T\rangle$ is factorizable according to G_{sup}. The solution is to merge cliques of G (by adding edges) that contain a separator in Ω_t to a bigger clique. More formally, given a FR-PKB $\langle G, T\rangle$ with $G = \langle V, E\rangle$, and a vertex v s.t. $v \in s_v$ and s_v is a separator of G, let $C_{s_v} \subseteq C$ be the set of all cliques that contains s_v, we built a new graph $G_{sup,v} = \langle V, E_{sup,v}\rangle$ as:

$$E_{sup,v} = \{(v_1, v_2) | v_1 \in C_{s_v} \text{ and } v_2 \in C_{s_v}\}.$$

The property of the chordal graph ensures that the graph $G_{sup,v}$ is again chordal. By iterating, this extends naturally to a set of vertices Ω_t in G and we denote the result as G_{sup,Ω_t} or simply G_{sup}.

Theorem 4.8. *Given a FR-PKB* $\langle G, T\rangle$ *and a set of vertices* $\Omega_t \subseteq S$*, let* $\mathrm{T}_{\text{WR-PKB}}$ *be the corresponding WR-PKB of* $\langle G, T\rangle$ *and* G_{sup} *a supergraph of* G *constructed as above,* $\mathrm{T}_{\text{WR-PKB}}$ *is factorizable according to* G_{sup}*.*

The proof is by induction on the size of Ω_t and the key in the induction step is that by Theorem 4.3, marginal probabilities over assignment of the newly merged clique can be expressed again by Eq. (7) via the factors of the original cliques and their separators. This also provides a way to derive the theory T_{sup} of the new FR-PKB. That is, T_{sup} can be obtained by the following steps:

(1) for all unaffected cliques c and separators s, retain $\boldsymbol{B}(\theta\colon r_\theta)$ from T to T_{sup} where $\theta \in \Theta_c$ or $\theta \in \Theta_s$;
(2) for all newly merged cliques $c^\star \in G_{sup}$, assuming w.l.o.g $c_1, \ldots c_k$ are the original cliques and $s = \bigcap_i c_i$ is the separator in G that generating $c^\star$, adding $\boldsymbol{B}(\theta\colon r_\theta)$ to T_{sup} where $\theta \in \Theta_{c^\star}$ and $r_\theta = \prod_{c_i} r_{i,\theta} / r_{s,\theta}$, here $r_{i,\theta}$ ($r_{s,\theta}$) is degree of belief of the resp. assignment over clique c_i (separator s) from θ.

We mark the above process to derive G_{sup} and T_{sup} as InsertEdges. Hence, the overall procedure to progress a FR-PKB $\langle G, T\rangle$ wrt a deterministic action t in a SCF-SSAs is as follows (mark as Progress):

(1) partitioning Ω_t into $\Omega_s \cup \Omega_{\bar{s}}$ where Ω_s are vertices in separators and $\Omega_{\bar{s}} = \Omega_t \backslash \Omega_s$;
(2) deriving a new FR-PKB $\langle G_{sup}, T_{sup}\rangle$ from $\langle G, T\rangle$ and Ω_s via the procedure InsertEdges;
(3) deriving the progressed FR-PKB $\langle G'_{sup}, T'_{sup}\rangle$ from $\langle G_{sup}, T_{sup}\rangle$ and $\Omega_{\bar{s}}$ via the procedure DeleteEdges.

Theorem 4.9. *Given a FR-PKB* $\langle G, T\rangle$ *over facts* $\mathsf{AF} = \{\alpha_1, \ldots, \alpha_k\}$*, a BATs* $\Sigma = \Sigma_{post} \cup \Sigma_I \cup \Sigma_{oi}$*, and an action* t *s.t.* t *is deterministic in* Σ *and* Σ_{post} *is strongly context-free, let* Ω_t *be the characteristic set of* t *in* Σ *and* $\langle G'_{sup}, T'_{sup}\rangle$ *the FR-PKB obtained from the procedure* Progress *via the input* $\langle G, T\rangle$ *and* Ω_t*, then*

$$\langle G, T\rangle \wedge \boldsymbol{K}\Sigma \models [t]\langle G'_{sup}, T'_{sup}\rangle.$$

The proof is by Theorem 4.8 and Theorem 4.6.

Example 4.10. Consider the variant of the blocks world in Example 4.7 and the ground action $t = putground(B)$. Let $\langle G, T\rangle$ be a FR-PKB where G is given as Fig.(1a) and T as Fig. 2 again. We have $\Omega_t = eh(B)$, which is the separator of G. By the procedure InsertEdges, we merge the cliques $eh(A)$ - $eh(B)$ and $eh(B)$ - $eh(C)$ resulting a graph G_{sup} as in Fig. (1b) and theory T_{sup} as Eq. (5). After the procedure DeleteEdges, we obtain a graph G'_{sup} with only one edge among $eh(A)$ and $eh(C)$. In addition, T'_{sup} is as $T'_{sup} = \{\boldsymbol{B}(eh(B)\colon 1), \boldsymbol{B}(eh(A) \land eh(C)\colon 0.264), \boldsymbol{B}(\neg eh(A) \land eh(C)\colon 0.204), \boldsymbol{B}(eh(A) \land \neg eh(C)\colon 0.136), \boldsymbol{B}(\neg eh(A) \land \neg eh(C)\colon 0.396)\}$.

We comment that although the procedure InsertEdges tends to add edges to the graph of a FR-PKB, the procedure DeleteEdges tends to simplify the graph. As a result, the progressed FR-PKB via Progress tends to be compact compared to the WR-PKB.

Challenge in Extending to General Stochastic Actions. Perhaps, the most appealing point of the Progress procedure is that progression can be obtained by some simple graph operators. A natural question is if we can generalize the approach to deterministic action with QF-LE-SSAs and even general stochastic actions. As aforementioned, an important feature of stochastic action is that they will introduce new correlations among facts. Hence, the challenges lie in how to modify the graph to reflect such a change in correlations.

In fact, two types of correlations may be introduced: *logical correlations* and *numerical correlations*. Logical correlations are in the form of logical disjunctions. In the block world domain (Example 2.10), the ground action $move(A, B, C)$ introduces a logical correlation among the facts $eh(A)$ and $eh(C)$: whatever the initial PKB is, after the deterministic action $move(A, B, C)$, $eh(A) \equiv \neg eh(C)$ is believed with full degree. This is in contrast to the deterministic actions with SCF-SSAs where classical progression will only yield negations and conjunctions. Therefore, for the progression of deterministic actions with QF-LE-SSAs, there are two sub-questions: 1) how to identify such logical correlations and 2) how to edit the graph to insert such correlations. For 1), transferring the result of progression into conjunctive normal form (CNF) might not be satisfactory as CNF might contain redundant logical correlations. E.g. the two clauses $(P \lor Q \lor R) \land (P \lor Q \lor \neg R)$ in CNF expressed two pairs of correlations, yet it could also be simplified to one pair, i.e. $P \lor Q$. Hence, perhaps what we are interested in is the minimal CNF formula. Yet, the minimal circuit problem is proven to be Σ_2^P-complete in [7]. For 2) simply adding edges among facts for new correlations does not work, as the resulting graph might not be chordal. E.g. adding an edge between $eh(A)$ and $eh(D)$ to the path $eh(A)$-$eh(B)$-$eh(C)$-$eh(D)$ yields a circle of length 4 (see Fig. (1c)).

What makes the situation worse is the *numerical correlations*. Numerical correlations exist in the form of the numerical likelihood of stochastic actions. Simply changing the structure of the graph will not help to capture such correlations. For instance, in Example 4.7, the procedure Progress yields a graph with a single edge $eh(B)$-$eh(C)$ for the ground action $putground(A)$. Likewise, for the ground action $putground(C)$, procedure Progress generates a graph with one edge $eh(A)$-$eh(B)$. Imagine a stochastic action that includes both $putground(A)$ and $putground(C)$ as alternatives with, say half-and-half likelihood, a natural thought of updating the graph according to the stochastic action perhaps is to merge the two graphs resulting a graph as in Fig. (1a), nevertheless, the progression of the WR-PKB in Eq (5) computed as Theorem 3.3 is not factorizable according to it. In fact, simply merging the graph does not capture the half-and-half numerical correlation among $putground(A)$ and $putground(C)$. This observation makes the graph editing approach to probabilistic progression less promising.

5 RELATED WORK AND CONCLUSION

The situation calculus is perhaps one of the most popular languages for actions in KR, although there are others. In the *fluent calculus* [34] has explored the progression of categorical knowledge with noisy actions. Classical categorical progression focuses on succinctness of progression includes [32]. In [22], using Bayesian networks, an account of progression is provided, but no characterization is given as to under what conditions the representation is guaranteed to remain factorizable. In dynamic epistemic logic, there are accounts of probabilities [12], but the progression on a knowledge basis is not explicitly considered. There are various formalisms for capturing probabilities and logic, but not explicitly considering actions; e.g. [2] and [25], amongst others.

In machine learning, formalisms for representing probabilities and updating them over time exist, like Dynamic Bayesian Networks [8] and Kalman filters [35]. Our result on factored representations can be seen as a version of dynamic Bayesian networks but while allowing for distinct actions (and not simply time steps). Our world-representations result can be seen as a richer version of Kalman filters which allowing for local effect action theories. The nature of our inquiry in terms of the underlying representation language makes our results largely orthogonal to probabilistic planning languages [29], including probabilistic STRIPS [39]. There, the emphasis is on states and their transitions after noisy actions, but often under the assumption that the successor state is immediately visible. Our formalism, however, is a general-purpose epistemic language for degrees of belief where the effects of actions may not be observable. However, it is clearly possible that with some effort one could implement a fragment of our language using probabilistic planning languages, or vice versa: the semantics for some probabilistic planning languages could be given using the underlying belief logic. Likewise, relating our results to POMDPs [10] and relational POMDPs [30] might allow us to reason about projection using one of these tools.

We study probabilistic progression for two representations of PKBs and show that the world-level representation is the most obvious choice if noisy actions are considered whereas the factored representation is advantageous if only deterministic actions are considered. For future work, leveraging progression to implement agents executing belief-based programs [5, 16, 19] is interesting. Besides, looking into ways to connect the situation calculus theory with probabilistic planning [39] and probabilistic decision theory [30] will allow for better cross-fertilization of formalisms, and lay the foundations for cognitive robotics [13].

ACKNOWLEDGMENTS

This work has been supportedp by a Royal Society University Research Fellowship.

REFERENCES

[1] Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. 2022. Logic tensor networks. *Artificial Intelligence* 303 (2022), 103649.

[2] Chitta Baral, Michael Gelfond, and J. Nelson Rushton. 2009. Probabilistic reasoning with answer sets. *TPLP* 9, 1 (2009), 57–144.

[3] Vaishak Belle and Gerhard Lakemeyer. 2011. A Semantical Account of Progression in the Presence of Uncertainty. In *Proc. AAAI*. AAAI Press, San Francisco, USA, 165–170.

[4] Vaishak Belle and Gerhard Lakemeyer. 2017. Reasoning about Probabilities in Unbounded First-Order Dynamical Domains. In *IJCAI*. ijcai.org, Melbourne, Australia, 828–836.

[5] Vaishak Belle and Hector Levesque. 2015. Allegro: Belief-based programming in stochastic dynamical domains. In *IJCAI*. AAAI Press, Buenos Aires, Argentina, 2762–2769.

[6] Vaishak Belle and Hector J. Levesque. 2014. How to progress beliefs in continuous domains. In *Proc. KR*. AAAI Press, Vienna, Austria, 438–447.

[7] David Buchfuhrer and Christopher Umans. 2011. The complexity of boolean formula minimization. *J. Comput. System Sci.* 77, 1 (2011), 142–153.

[8] Thomas Dean and Keiji Kanazawa. 1988. Probabilistic temporal reasoning. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*. AAAI Press / The MIT Press, St. Paul, MN, USA, 524–528.

[9] Lise Getoor and Ben Taskar. 2007. *Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning)*. The MIT Press, Cambridge, MA, USA.

[10] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101, 1–2 (1998), 99 – 134. https://doi.org/10.1016/S0004-3702(98)00023-X

[11] Kristian Kersting and Luc De Raedt. 2000. Bayesian Logic Programs. In *ILP Work-in-progress reports*. CEUR-WS.org, London, UK, 8.

[12] B.P. Kooi. 2003. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information* 12, 4 (2003), 381–408.

[13] Gerhard Lakemeyer and Hector J. Levesque. 2007. Cognitive robotics. In *Handbook of Knowledge Representation*. Elsevier, Dagstuhl, Germany, 869–886.

[14] Fangzhen Lin and Ray Reiter. 1994. Forget it. In *Working Notes of AAAI Fall Symposium on Relevance*. AAAI Press, New Orleans, Louisiana, 154–159.

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

[16] Daxin Liu. 2023. *Projection in a probabilistic epistemic logic and its application to belief-based program verification*. PhD dissertation. RWTH Aachen University.

[17] Daxin Liu and Qihui Feng. 2021. On the Progression of Belief. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, Vol. 18. ijcai.org, Online event, 465–474.

[18] Daxin Liu and Qihui Feng. 2023. On the progression of belief. *Artificial Intelligence* 322 (2023), 103947.

[19] Daxin Liu, Qinfei Huang, Vaishak Belle, and Gerhard Lakemeyer. 2023. Verifying Belief-Based Programs via Symbolic Dynamic Programming. In *ECAI 2023*. IOS Press, Kraków, Poland, 1497–1504.

[20] Daxin Liu and Gerhard Lakemeyer. 2021. Reasoning about Beliefs and Meta-Beliefs by Regression in an Expressive Probabilistic Action Logic. In *IJCAI*. ijcai.org, Online event / Montreal, Canada, 1951–1958.

[21] Yongmei Liu and Gerhard Lakemeyer. 2009. On first-order definability and computability of progression for local-effect actions and beyond. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*. ijcai.org, Pasadena, California, USA, 860–866.

[22] Yves Martin and Michael Thielscher. 2009. Integrating Reasoning About Actions and Bayesian Networks. In *International Conference on Agents and Artificial Intelligence*. INSTICC Press, Valencia, Spain, 298–304.

[23] John McCarthy and Patrick J Hayes. 1981. Some philosophical problems from the standpoint of artificial intelligence. In *Readings in artificial intelligence*. Morgan Kaufmann, San Francisco, USA, 431–450. https://doi.org/10.1016/B978-0-934613-03-3.50033-7

[24] Judea Pearl. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Francisco, USA.

[25] Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole. 2016. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 10, 2 (2016), 1–189.

[26] Raymond Reiter. 1981. On closed world data bases. In *Readings in artificial intelligence*. Elsevier, Amsterdam, The Netherlands, 119–140.

[27] Raymond Reiter. 2001. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT press, Cambridge, MA, USA.

[28] Matthew Richardson and Pedro Domingos. 2006. Markov logic networks. *Machine learning* 62, 1 (2006), 107–136.

[29] Scott Sanner. 2011. *Relational Dynamic Influence Diagram Language (RDDL): Language Description*. Technical Report. Australian National University.

[30] Scott Sanner and Kristian Kersting. 2010. Symbolic Dynamic Programming for First-order POMDPs. In *Proc. AAAI*. AAAI Press, Atlanta, Georgia, USA, 1140–1146.

[31] Prithviraj Sen and Amol Deshpande. 2006. Representing and querying correlated tuples in probabilistic databases. In *2007 IEEE 23rd international conference on data engineering*. IEEE, IEEE Computer Society, Istanbul, Turkey, 596–605.

[32] Afsaneh Shirazi and Eyal Amir. 2005. First-Order Logical Filtering. In *Proc. IJCAI*. Professional Book Center, Edinburgh, Scotland, UK, 589–595.

[33] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. 2022. *Probabilistic databases*. Springer Nature, Berlin, Germany.

[34] Michael Thielscher. 2001. Planning with Noisy Actions (Preliminary Report). In *Proc. Australian Joint Conference on Artificial Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg, 27–45.

[35] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA.

[36] Stavros Vassos and Hector J Levesque. 2008. On the Progression of Situation Calculus Basic Action Theories: Resolving a 10-year-old Conjecture.. In *AAAI*, Vol. 8. AAAI Press, Chicago, Illinois, USA, 1004–1009.

[37] Stavros Vassos and Hector J Levesque. 2013. How to progress a database III. *Artificial Intelligence* 195 (2013), 203–221.

[38] Richard Waldinger. 1981. Achieving several goals simultaneously. In *Readings in artificial intelligence*. Morgan Kaufmann, San Francisco, USA, 250–271.

[39] H. Younes and M. Littman. 2004. *PPDDL 1. 0: An extension to PDDL for expressing planning domains with probabilistic effects*. Technical Report. Carnegie Mellon University.