Towards Defeasible Reasoning about Actions and Their Effects

Jonas Haldimann[1,*], Mantas Šimkus[2] and Thomas Meyer[1]

[1]University of Cape Town and CAIR, Cape Town, South Africa
[2]TU Wien, Vienna, Austria

Abstract
Intelligent agents interacting with their environment need to reason about the effects their actions will have on their surroundings. For actions with deterministic outcome this problem is well investigated in the context of planning problems. But many real-world actions may not always result in their typical outcome, and exceptional outcomes of an action have to be taken into account. In this work we present a framework for reasoning about actions with uncertain outcomes. Inspired by conditionals, we introduce *defeasible action rules* that capture the typical outcome of action, without committing to what happens in exceptional circumstances. For the semantics of these rules we introduce *ranked transition schemas*: graph-based representations of transitions between possible worlds in which edges between states are assigned a *rank* indicating their typicality. We define an inference operator for defeasible action rules inspired by rational closure, adapting its intuitive semantics to the dynamic setting. Furthermore, we explore a generalization of this framework that uses partial orderings over transitions instead of ranks. This generalization strictly increases expressive power and enables more refined reasoning operators. Finally, we extend the language of defeasible action rules with a notation to express that certain properties are not affected by the execution of an action, thus allowing to explicitly formulate frame axioms in this defeasible setting.

Keywords
Actions, Defeasible rules, Time, Rational Closure, Planning

1. Introduction

Intelligent agents operating in the real world must reason about the effects their actions will have on their environment. Whether it is a robot navigating a building, a planning system scheduling logistics, or a household assistant operating appliances, the ability to model action-effect relationships formally and to draw reliable inferences from them is a foundational requirement of autonomous behaviour. A number of approaches based on classical logic, like the STRIPS formalism of Fikes and Nilsson [1] and its successors, have been proposed and were proven enormously successful for a broad class of planning tasks. In these frameworks, each action is described by a set of preconditions and a pair of add and delete lists that deterministically specify how the world changes upon execution. The challenge of specifying, for any action, all the things that it does not change is called the *frame problem* [2]. It is commonly tackled by the closed-world assumption, stating that everything not mentioned in an action's effects is presumed to persist unchanged. The deterministic, fully observable assumptions underlying classical planning are, however, a significant idealization. Many real-world actions do not reliably produce a single deterministic outcome. For example, a robot pressing a light switch will typically turn the light on, but may fail to do so if the fuse is blown or the light is broken. In many cases the outcome of an action depends on aspects of the world that may be incompletely known, requiring an agent to tolerate exceptions to the rules in her knowledge. Additionally, even if known it is often unfeasible to list all preconditions of an effect in such detail that they cover every thinkable exception, a problem also known as *qualification problem* [3, 4] and a well-recognised obstacle to the use of classical planning in realistic settings.

24th International Workshop on Nonmonotonic Reasoning
*Corresponding author.
jonas@haldimann.de (J. Haldimann); mantas.simkus@tuwien.ac.at (M. Šimkus); tmeyer@cair.org.za (T. Meyer)
0000-0002-2618-8721 (J. Haldimann); 0000-0003-0632-0294 (M. Šimkus); 0000-0003-2204-6969 (T. Meyer)

One approach to handling uncertainty are probabilistic planning frameworks such as Markov Decision Processes (MDPs) [5] and their partially observable generalizations (POMDPs) [6] which address uncertain outcomes by assigning explicit probability distributions over successor states. While these give a principled formalization of uncertainty, they require precise numerical parameters that may be difficult or impossible to obtain in practice. For many practical purposes, an agent needs only a qualitative account of which outcomes are typical and which are exceptional and not exact probabilities and a mechanism to use them for inference about her environment.

Defeasible reasoning with conditionals, developed in the context of nonmonotonic reasoning (NMR), provides exactly this kind of qualitative, exception-tolerant inference. Conditionals $a \rightsquigarrow b$ formalize rules of the form "if a then typically b" allowing for exceptions; and to solve the qualification problem more specific conditionals may override general conditionals. The framework of Kraus, Lehmann, and Magidor [7] establishes foundational postulates and semantic models for well-behaved defeasible consequence relations, and Lehmann and Magidor's rational closure [8] offers a concrete, algorithmically tractable inference operator. Other examples for inference operators for conditionals are Pearl's system Z [9] which coincides with rational closure and the more involved system W [10] which refines rational closure and allows for more inferences. Common semantic models for conditionals are Spohn's ranking functions [11] or models based on typicality orderings [7].

So far, classical planning and the non-monotonic tradition have been investigated largely in isolation. On the planning side, the semantics are typically tied to transition systems without a notion of typicality. On the NMR side, defeasible conditionals have been studied extensively for static knowledge bases, where rules describe properties of the world at a single point in time, but comparatively little work has been done on extending this machinery to the dynamic setting. Especially, KLM-style reasoning with defeasible conditionals has not yet been combined with planning. To the best of our knowledge, in the planning context defeasibility has only been considered for frame rules, i.e., rules that describe how properties do not change unless there is a reason for them to change. These have been realized for example by using circumscription [12], or by treating them as defeasible rules in the semantic [13]. Other nonmonotonic planning mechanisms used for planning, such as Answer Set Programming (ASP) [14], realize exceptions are explicitely by using default negation.

In this paper we bridge conditional reasoning and planning by developing a framework for defeasible reasoning about actions and their effects. Inspired by conditionals, we introduce *defeasible action rules* of the form $p, \alpha \rightsquigarrow E$, expressing that if precondition p holds and action α is performed, then typically the formulas in E hold afterwards. As semantic for such rules we introduce *ranked transition schemas*: graph-structured representations of transitions between possible worlds in which each edge is assigned a rank indicating its typicality. This adapts the notion of a ranking function from the static reasoning to the dynamic setting, with ranks now placed on transitions rather than worlds. Using these ranked transition schemes we define an inference operator for defeasible action rules inspired by rational closure, adapting the Z-partition construction to stratify rules and rank transitions accordingly.

We then propose a more expressive generalization of ranked transition schemas that uses partial orderings over transitions rather than integer ranks. Using these *ordered transition schemas* we define a finer-grained inference operator analogous to system W.

Finally, we extend the language of defeasible action rules with *frame rules* that allow to state that the valuation of a formula is not changed in the effects of an action. These frame rules allow frame conditions to be stated explicitly and to be taken into account in defeasible reasoning about an actions effects.

In summary, the contributions of this paper are the following:

- Definition of defeasible action rules and ranked transition models as their semantic models
- Development of rational closure for action rules
- Generalization of ranked transition models to ordered transition models
- Development of w-closure for defeasible action rules
- Introduction of frame rules expressing that something is not changed by an action.

In Section 2 we briefly recall the background on conditional reasoning and ranking functions before defining action rules and ranked transition schemas in Section 3. Then, we develop rational closure for defeasible action rules in Section 4. Ordered transition schemas and w-closure are introduced in Section 5 before we consider frame rules in Section 6.

2. Background

A *(propositional) signature* is a finite set Σ of identifiers or *variables*. For a signature Σ, we denote the propositional language over Σ by $\mathcal{L}_\Sigma$. As usual, $\top$ denotes a tautology and $\bot$ an unsatisfiable formula. The set of interpretations over a signature Σ is denoted as Ω_Σ. Interpretations are also called *worlds*. An interpretation $\omega \in \Omega_\Sigma$ is a *model* of a formula f if A holds in ω. This is denoted as $\omega \models A$. The set of models of a formula (over a signature Σ) is denoted as $Mod(A) = \{\omega \in \Omega_\Sigma \mid \omega \models A\}$. A formula A *entails* a formula B, denoted by $A \models B$, if $Mod(A) \subseteq Mod(B)$. We will represent interpretations (or worlds) by complete conjunctions, e.g., the interpretation over $\Sigma_{abc} = \{a, b, c\}$ that maps a and c to *true* and b to *false* is represented by $a \wedge \neg b \wedge c$, or just $a\bar{b}c$ where overlining represents negation of atoms. Thus, every world $\omega \in \Omega_\Sigma$ is also a formula in $\mathcal{L}_\Sigma$.

A *conditional* $(B|A)$ connects two formulas A, B and represents the rule "If A then usually B". Conditionals are commonly used with a three-valued semantics with respect to worlds [15]: for a world ω, a conditional $(B|A)$ is either *verified* by ω if $\omega \models A \wedge B$, *falsified* by ω if $\omega \models A \wedge \neg B$, or *not applicable* to ω if $\omega \models \neg A$. The set of conditional beliefs accepted by an agent is modelled by an inference relation. An *inference relation* is a binary relation $\mid\!\sim$ on propositional formulas with $A \mid\!\sim B$ representing that A (defeasibly) entails B.

Ranking functions are a common model for conditionals. A *ranking function*, also called *ordinal conditional function* (OCF), is a function $\kappa : \Omega_\Sigma \to \mathbb{N}_0 \cup \{\infty\}$ such that $\kappa^{-1}(0) \neq \emptyset$; ranking functions were first introduced (in a more general form) by Spohn [11]. The intuition of a ranking function is that the rank of a world is lower if the world is more plausible. Therefore, ranking functions can be seen as some kind of "implausibility measure". Ranking functions are extended to formulas by $\kappa(f) = \min_{\omega \in Mod(f)} \kappa(\omega)$. A ranking function κ models a conditional $(B|A)$, denoted as $\kappa \models (B|A)$, if $\kappa(A \wedge B) < \kappa(A \wedge \neg B)$, i.e., if the verification of the conditional is strictly more plausible than its falsification. The *inference relation* $\mid\!\sim_\kappa$ *induced by a ranking function* κ is defined by $A \mid\!\sim_\kappa B$ iff $\kappa(A) = \infty$ or $\kappa(A \wedge B) < \kappa(A \wedge \neg B)$.

Ranking functions can be restricted to a subset of Ω_Σ to capture the plausibility of worlds in a given setting. The *conditionalization* of ranking functions κ to the models of a formula $f \in \mathcal{L}_\Sigma$ is the function $\kappa|_f : Mod(f) \to \mathbb{N}_0 \cup \{\infty\},\ \omega \mapsto \kappa(\omega) - \kappa(f)$.

3. Modelling Defeasible Effects of Actions

For this paper we consider the setting where an agent uses a propositional signature Σ to describe the state of the environment. The set A contains all the actions available to the agent. With this we define defeasible action rules that describe changes in the environment in response to an agent's actions.

Definition 1. *A* defeasible action rule *is a rule of the form*

$$p, \alpha \rightsquigarrow E$$

where p is a formula, $\alpha \in A$, and E is a set of formulas. The formula p is called the precondition*, and formulas in E are called the* effects *of the defeasible action rule.*

Intuitively, $p, \alpha \rightsquigarrow E$ encodes that if p holds and action α is taken, then in the next step typically the formulas in E hold. The following example illustrates the forms such rules can take. Observe how the rules represent typical or expected outcomes of actions – while at the same time allowing for exceptions. This allows more specific rules to "override" more general ones.

Example 1. *Assume an agent knows the following about the light switch in her room:*

- *If she pushes the button while the light is off, typically the light is on afterwards.*
- *But if she pushes the button with the light off and while the fuse is blown, typically the light is not on afterwards.*
- *If she pushes the button with dirty hands, the button is usually dirty afterwards.*

Using the signature $\Sigma = \{light,\ fuse,\ dirty\_hand,\ dirty\_button\}$ *and the actions* $A = \{do\_nothing,\ push\_button\}$ *this can be modelled by the following action rules:*

$$\neg light, push\_button \rightsquigarrow \{light\} \qquad (r_1)$$

$$\neg light \wedge \neg fuse, push\_button \rightsquigarrow \{\neg light\} \qquad (r_2)$$

$$dirty\_hand \wedge \neg dirty\_button, push\_button \rightsquigarrow \{dirty\_button\}. \qquad (r_3)$$

Observe, how (r_2) *overrides* (r_1) *for the special case that the light is off and the fuse is blown.*

An interesting special case arises when the set of effects E consists entirely of literals. In this case, the rule carries not only a descriptive reading, specifying what typically holds after the action, but can also be interpreted operationally: the effect in the form of a literal a or $\neg a$ is most naturally reached by directly flipping the truth value a in the initial world. This gives rules with literal effects a flavour reminiscent of STRIPS operators, where the add and delete lists of an action determine precisely how the state is updated. While this observation could allow us to connect defeasible action rules to established planning methods in future work, for now we will focus on the general case where E can contain arbitrary formulas.

As a semantic for the new rules we propose the following structure, which represents which transitions between worlds are possible outcomes of an action and how typical they are.

Definition 2. *An* ranked transition schema *is a tuple* $\mathcal{T} = (V, T, \kappa)$ *consisting of a set of worlds* $V = \Omega_\Sigma$*, a set of transitions* $T \subseteq V \times V \times A$*, and a ranking function* $\kappa : T \to \mathbb{N}_0$ *such that at least one transition in* T *is assigned rank* 0*.*

In a ranked transition schema $\mathcal{T} = (V, T, \kappa)$, the set V captures the set of different situations an agent can be in, each represented by a world. A transition $t \in T$ with $t = (\omega, \omega', \alpha)$ represents the transition from ω to ω' after the action α was taken. The ranking function κ captures how typical each transition is, where low ranks represent more typical transitions and high ranks represent less typical transitions. Accordingly, T contains all possible transitions with even the slightest possibility of being effected.

In the context of NMR, ranking functions sometimes include a rank ∞ as the rank of completely impossible worlds. Ranked transition schemas do not use infinite ranks. Instead, it simply does not include impossible transitions in T. This is in line with many NMR formalisms, such as preferential models [7], which only orders the possible states. Our design choice of omitting infinite ranks is therefore not a loss of generality.

The key insight underlying this semantic structure is that a model for defeasible action rules must represent the typicality of transitions between states for given actions, since the uncertainty of defeasible action rules lies in their effects rather than in their preconditions.

Example 2. *Consider the signature* $\Sigma = \{l, f\}$ *and the set of actions* $A = \{n, p\}$*. We can take these to be abbreviations of the atoms* $light$ *and* $fuse$ *as well as the actions* $do\_nothing$ *and* $push\_button$ *from Example 1 (ignoring reasoning about dirt). The tuple* $\mathcal{T} = (V, T, \kappa)$ *with*

- $V = \{lf, l\overline{f}, \overline{l}f, \overline{l}\,\overline{f}\}$,
- $T = \{t_1, t_2, t_3, t_4, t_5, t_6, t_7, t_8, t_9\}$ *with*

$$t_1 = (\overline{l}f, \overline{l}f, p) \quad t_2 = (\overline{l}\,\overline{f}, \overline{l}f, p) \quad t_3 = (lf, lf, p) \quad t_4 = (\overline{l}\,\overline{f}, l\overline{f}, p) \quad t_5 = (l\overline{f}), \overline{l}\,\overline{f}, p)$$

$$t_6 = (lf, lf, n) \quad t_7 = (\overline{l}f, \overline{l}f, n) \quad t_8 = (l\overline{f}, l\overline{f}, n) \quad t_9 = (\overline{l}\,\overline{f}, \overline{l}\,\overline{f}, n)$$

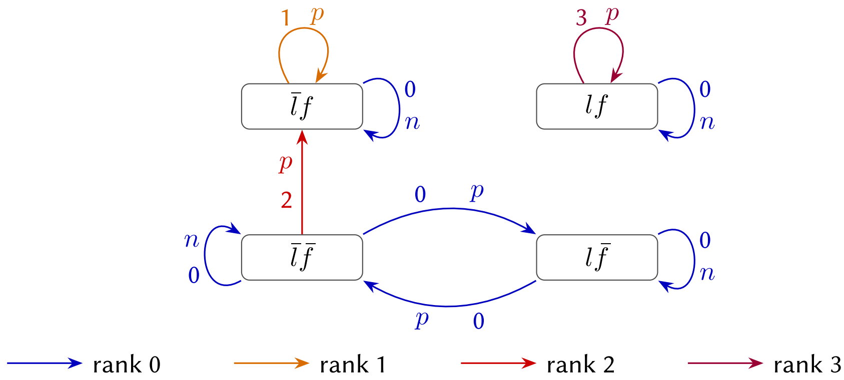

Figure 1: Illustrantion of the ranked transition schema $\mathcal{T}$ from Example 2. Ranks are emphasized by colour.

- $\kappa : T \to \mathcal{N}_0$ *with*

$$\kappa(t_1) = 1 \qquad \kappa(t_2) = 2 \qquad \kappa(t_3) = 3 \qquad \kappa(t_4) = 0 \qquad \kappa(t_5) = 0$$
$$\kappa(t_6) = 0 \qquad \kappa(t_7) = 0 \qquad \kappa(t_8) = 0 \qquad \kappa(t_9) = 0$$

is an example for a ranked transition schema. For world $\bar{l} f$ *and action* p *for example, the plausible (rank 0) transition is to* lf*, while the transition to* $\bar{l}\,\bar{f}$ *is possible but unplausible (rank 2).*

Ranked tansition schemas can be consciesly illustrated by graphs where worlds correspond to nodes and transitions to edges, labelled with their action and rank. $\mathcal{T}$ *is illustrated in Figure 1.*

We can now connect transitions with defeasible action rules. For a given transition $t = (\omega, \omega', \alpha)$ and a rule $r = p, \alpha \rightsquigarrow E$ we say

- t *verifies* r if $\omega \models p$ and $\omega' \models E$,
- t *falsifies* r if $\omega \models p$ and $\omega' \not\models E$, and
- r is *not applicable* to t if $\omega \not\models p$.

This is in line with the three-valued semantics for conditionals outlined in Section 2. We can now state when a ranked transition schema models a defeasible action rule.

Definition 3. *A ranked transition schema* $\mathcal{T} = (V, T, \kappa)$ *models a defeasible action rule* $r = p, \alpha \rightsquigarrow E$ *if there is no transition* $t_f \in T$ *falsifying* r*, or if there is at least one transition* $t_v \in T$ *verifying* r *and*

$$\min_{\substack{t_v \in T,\\ t_v \text{ verifies } r}} \kappa(t_v) < \min_{\substack{t_f \in T,\\ t_f \text{ falsifies } r}} \kappa(t_f).$$

That is, $\mathcal{T}$ models r if the most typical transitions that verify r are strictly more typical than the most typical transitions that falsify it. For a set R of defeasible action rules, we say that $\mathcal{T}$ models R if $\mathcal{T}$ models every rule in R.

Before defining inferences induced by ranked transition schemas, let us introduce some notation. Let $\mathcal{T} = (V, T, \kappa)$ be a ranked transition schema. For a given world ω and action α, let

$$T_{\omega,\alpha} = \{(\omega_e, \omega'_e, \alpha_e) \in T \mid \omega = \omega_e, \alpha = \alpha_e\}$$

denote the set of transitions outgoing from ω with action α. To determine the plausibility of different outcomes of α in ω, we need a form of conditionalization restricting the ranking κ to the situation. For this we write

$$\kappa_{\omega,\alpha} : T_{\omega,\alpha} \to \mathbb{N}_0, \qquad t \mapsto \kappa(t) - \min\{\kappa(t') \mid t' \in T_{\omega,\alpha}\}.$$

Subtracting $\min\{\kappa(t') \mid t' \in T_{\omega,\alpha}\}$ ensures a form of normalization, which means here that at least one transition in $T_{\omega,\alpha}$ has rank 0 with respect to $\kappa_{\omega,\alpha}$. By slight abuse of notation we also write $\kappa_{\omega,\alpha}(\omega')$ to denote $\kappa(t)$ for $t = (\omega, \omega', \alpha)$, allowing to concisely access the plausibility of the world ω' as successor of ω after action α.

With this, ranked transition schemas as models for uncertain action rules allow for a number of interesting inference tasks. The most immediate is using them to answer defeasible queries about what typically holds after executing a certain action.

Definition 4. *In the context of a ranked transition schema $\mathcal{T} = (V, T, \kappa)$, a* query $\omega, \alpha \rightsquigarrow e$ *consisting of a world ω, an action α and a formula e is answered positively if*

$$\{\omega' \mid \kappa_{\omega,\alpha}(\omega') = 0\} \subseteq Mod(e).$$

That is, the query is answered positively if all most typical successor states of ω after executing α satisfy e.

Beyond single-step queries, ranked transition schemas also support reasoning about traces (sequences of actions). Given a initial world ω_0 and a sequence of actions $(\alpha_1, \dots, \alpha_n)$ the rank of a trace $(\omega_0, \omega_1, \dots, \omega_n)$ is given by

$$\kappa((\omega_0, \dots, \omega_n), (\alpha_1, \dots, \alpha_n)) = \kappa_{\omega_0,\alpha_1}(\omega_1) + \dots + \kappa_{\omega_{n-1},\alpha_n}(\omega_n).$$

These ranks can be used to answer more involved questions about multi-step behaviour. For example, an agent that has executed a sequence of actions without observing the intermediate states may use trace ranks to reason retrospectively about the most plausible sequence of states leading to the current situation. Trace ranks also support more refined planning: rather than simply finding a sequence of actions that typically achieves a goal, an agent may seek a particular safe plan, in which unintended outcomes are highly atypical. Or conversely, an agent may be interested in risky plans if it offers an unlikely but significant reward. Connecting these models to planning problems in greater depth is a promising research path for future work.

What is left open so far is the question of how to obtain an adequate ranked transition schema from a given set of defeasible action rules. In the next section we will develop a method for this, enabling reasoning from defeasible action rules.

4. Reasoning with Defeasible Action Rules

The modelling relation introduced in the previous section is deliberately permissive: a ranked transition schema models a defeasible action rule as soon as a single transition verifying the rule is ranked more typical than the falsifying ones. For reasoning applications, however, the challenge is to construct a suitable ranked transition schema from a given set of defeasible action rules. While many different approaches inspired by inference operators for conditionals seem feasible, here we use a method inspired by rational closure [8] and the equivalent system Z [9], adapting its construction to the setting of action rules and uncertain transitions.

The first key step in rational closure is assessing the exceptionality of rules. Intuitively, a rule is exceptional in a set of rules if there is no situation in which it can be verified without some other rule being falsified. We carry this idea over to the action setting as follows.

Definition 5. *Let R be a set of defeasible action rules. A rule $(p, \alpha \rightsquigarrow E) \in R$ is* exceptional *in R if there is no world $\omega \in \Omega_\Sigma$ and action $\alpha \in A$ such that $\omega \models p$ and $E \cup \bigcup_{(p',\alpha \rightsquigarrow E') \in R,\ \omega \models p'} E'$ is consistent. $\mathcal{E}(R)$ denotes the set of rules in R that are exceptional in R.*

Equivalently, a rule r is exceptional for a set of rules R if there is no transition that verifies r without simultaneously falsifying some other rule in R. This captures the idea that the rule applies only in situations that are already atypical with respect to the broader rule set.

Using this notion of exceptionality, we can group the rules in R into a sequence of sets of increasing specificity, analogous to the Z-partition construction for conditional knowledge bases.

Definition 6. *Let R be a set of defeasible action rules and*

$$\begin{array}{ll} E_0 = R & \\ E_1 = \mathcal{E}(E_0) & R_0 = E_0 \setminus E_1 \\ E_2 = \mathcal{E}(E_1) & R_1 = E_1 \setminus E_2 \\ \dots & \dots \end{array}$$

Let n be the highest number with $R_n \neq \emptyset$. Let $R_\infty = E_{n+1}$. The tuple $(R_1, \dots, R_n, R_\infty)$ is called the Z-partition *of R.*

The sets $R_0, \dots, R_n$ contain rules of increasing exceptionality: R_0 contains the most general rules in R, while sets contain progressively more specific rules that can only be verified by overriding more general rules. The set R_∞ collects rules that remain exceptional throughout the entire stratification process; these rules cannot be falsified by a transition in any model of R.

From the Z-partition of R we can now construct a ranked transition schema modelling R that will be later used for inference. The central idea is to assign each transition a rank depending on the latest part of the Z-partition it falsifies: a transition that satisfies every rule receives rank 0, while a transition that violates rules from lower layers but respects those from rank i onwards receives rank i. Transitions falsifying rules in R_∞ are excluded entirely.

Definition 7. *Let R be a set of defeasible action rules with the Z-partition $(R_1, \dots, R_n, R_\infty)$. The Z-ranked transition schema $\mathcal{T}_R^z = (V, T, \kappa^z)$ is constructed as follows. Let $V = \Omega_\Sigma$. Let $T \subseteq V \times V \times A$ be set of all transitions that do not falsify R_∞. For every such $t \in T$, let $\kappa^z(t) = i$ where i is the lowest number such that t does not falsify any rule in $E_i = R_i \cup \dots \cup R_n$.*

From now on we will assume that for every action $\alpha \in A$ there is a transition $(\omega, \omega', \alpha)$ that does not falsify any rule in R. This corresponds to the assumption of *weak consistency* [16] made commonly when reasoning with propositional conditionals. With this assumption, we can verify that $\mathcal{T}_R^z$ satisfies the normalization constraint and is indeed a ranked transition model.

Additionally, $\mathcal{T}_R^z$ is faithful to the original set of rules R in the sense that it is a model R.

Proposition 1. *Let R be a set of defeasible action rules. Then $\mathcal{T}_R^z$ is a model of R.*

Proof. Let $\mathcal{T}_R^z = (V, T, \kappa^z)$ be the Z-ranked transition schema of R and let $(R_1, \dots, R_n, R_\infty)$ be the Z-partition of R. Let $r \in R$ be arbitrary.

If $r \in R_\infty$, then by construction no transition in $\mathcal{T}_R^z$ violates r and therefore $\mathcal{T}_R^z$ is a model of r.

For the remainder of the proof assume that $r \notin R_\infty$ implying that there is an $i \in \{1, \dots, n\}$ such that $r \in R_i$. By definition of exceptionality, there must be a transition t_v verifying r without falsifying any rule in $E_i = R_i \cup \dots \cup R_n \cup R_\infty$. Therefore, $\kappa^z(t) \leq i$. Every transition t_f falsifying r will have a rank $\kappa^z(t') > i$ by construction of T_R^z. Therefore, $\min_{\substack{t_v \in E, \\ t_v \text{ verifies } r}} \kappa^z(t_v) \leq i < \min_{\substack{t_f \in E, \\ t_f \text{ falsifies } r}} \kappa^z(t_f)$, and thus $\mathcal{T}_R^z$ is a model of r. □

Example 3. *With $\Sigma = \{l, f, h, b\}$ and $A = \{n, p\}$ consider the set of defeasible action rules $R = \{r_1, r_2, r_3\}$ with $r_1 = \neg l, p \rightsquigarrow \{l\}$, $r_2 = \neg l \land \neg f, p \rightsquigarrow \{\neg l\}$, and $r_3 = h \land \neg b, p \rightsquigarrow \{b\}$. This can be seen as a version of Example 1 where atoms and actions are abbreviated with single-letter symbols. The Z-partition of R is (R_1, R_2) with $R_1 = \{r_1, r_3\}$ and $R_2 = \{r_2\}$. The resulting Z-ranked transition schema $\mathcal{T}_R^z$ is illustrated in Figure 2; to keep the figure readable, we only include the transitions starting in world $\bar{l} f h \bar{b}$ with action p and selected transitions starting in $\bar{l}\,\bar{f} h \bar{b}$. Observe how all transitions with rank 0 do not falsify any rule, transitions with rank 1 falsify one of the rules in R_1, and, e.g., $(\bar{l}\,\bar{f} h \bar{b}, \bar{l} \bar{f} h b, p)$ with rank 2 falsifies the rule in R_2.*

With the Z-ranked transition schema in hand, we can now define rational closure for defeasible action rules.

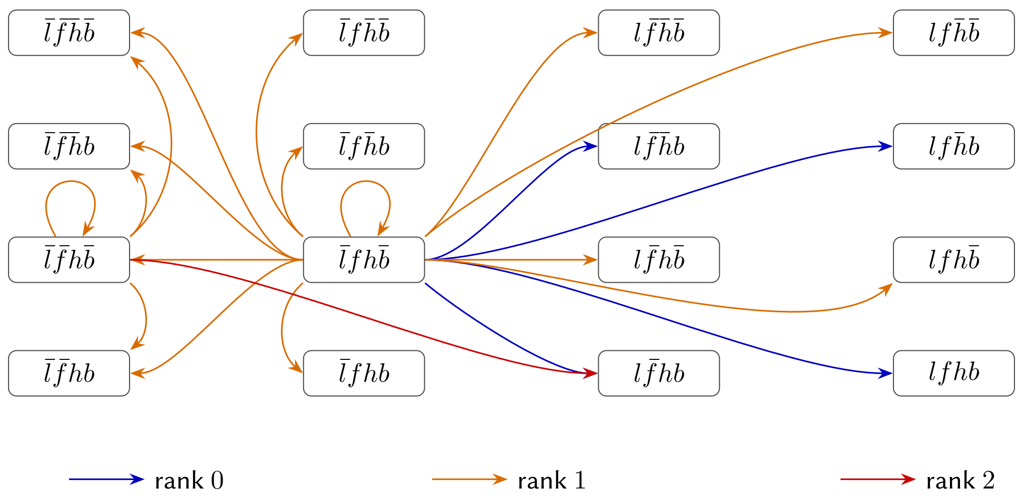

Figure 2: Illustration of the z-ranked transition schema $\mathcal{T}_R^z$ from Example 3. Only selected transitions are shown, and ranks are indicated by colour.

Definition 8. *Given a set of rules R,* rational closure for actions *answers the query $\omega, \alpha \rightsquigarrow e$ with true if it holds in $\mathcal{T}_R^z$.*

This operator inherits the central intuition of rational closure: conclusions are drawn from the most typical transitions consistent with the available rules, with more specific rules taking precedence over more general ones. When the world and action in a query uniquely determine the most typical outcome, the inference collapses to a deterministic prediction; when multiple equally typical outcomes exist, the query is answered positively only if all of them satisfy the queried formula.

To conclude this section, let us illustrate rational closure by an example.

Example 4. *Consider the set R of defeasible action rules from Example 3 again. In the world $\omega_1 = \bar{l}fh\bar{b}$ after performing action p can we expect $l \wedge b$ to hold? With rational closure the answer is* yes*, the query $\omega_1, p \rightsquigarrow l \wedge b$ is answered positively. To verify, we can check the transitions in $T_{\omega_1,p}$ with minimal rank (here rank 0) and see that they all lead to worlds in which $l \wedge b$ holds.*

What about the query $\omega_2, p \rightsquigarrow b$ for $\omega_2 = \bar{l}\bar{f}h\bar{b}$? It is answered negatively by rational closure, because some of the transitions in $T_{\omega_2,p}$ with minimal rank (rank 1 for this query) lead to worlds in which b does not hold.

In the next section we will consider an alternative to ranked transition schemas and rational closure.

5. Transition Schemas with Orderings

Ranking edges requires every two edges to be comparable. However, for some reasoning approaches that is too limiting. For propositional conditionals, some of the considered semantics like preferential models use strict partial orderings (SPO) instead of ranks, allowing for greater expressivity. SPOs are ordering relations that are transitive and irreflexive.

In the following we define an alternative to ranked transition schemas that use a strict partial order instead of a ranking function to represent (a)typicality of transitions.

Definition 9. *An* ordered transition schema *is a tuple $\mathcal{O} = (V, T, <)$ consisting of a set of worlds $V = \Omega_\Sigma$, a set of transitions $T \subseteq V \times V \times A$, and a strict partial ordering $<$ over T.*

Again, the set V captures the set of different worlds, and T the set of possible transitions. The SPO $<$ orders transitions according to their typicality: $t < t'$ indicates that t is more plausible/typical than t'. Completely impossible transitions are again completely omitted from T. Ordered transition schemas are an alternative semantic for defeasible action rules.

Definition 10. *An ordered transition schema $\mathcal{O} = (V, T, <)$, models a defeasible action rule $r = P, \alpha \rightsquigarrow E$ if for every $t_f \in T$ that falsifies r there is a $t_v \in T$ that verifies r.*

For a set R of defeasible action rules, $\mathcal{O}$ models R if $\mathcal{O}$ models every rule in R. Just as ranked transition schemas, ordered transition schemas can be used to answer defeasible queries.

Definition 11. *In the context of an ordered transition schema $\mathcal{O} = (V, T, <)$, a query $\omega, \alpha \rightsquigarrow e$ is answered positive if $\{\omega' \mid (\omega, \omega', \alpha) \in \min_< T_{\omega,\alpha}\} \subseteq Mod(e)$, i.e., if the minimal transitions starting in ω with action α lead to worlds in which e holds.*

Example 5. *Consider the signature $\Sigma = \{l, f\}$ and the set of actions $A = \{n, p\}$ (cf. Example 2). The tuple $\mathcal{T} = (V, T, <)$ with*

- $V = \{lf, l\overline{f}, \overline{l}f, \overline{l}\,\overline{f}\}$,
- $T = \{t_1, t_2, t_3, t_4, t_5, t_6, t_7, t_8, t_9\}$ *with*

$$\begin{array}{lllll} t_1 = (\overline{l}f, lf, p) & t_2 = (\overline{l}f, \overline{l}\,\overline{f}, p) & t_3 = (lf, \overline{l}f, p) & t_4 = (\overline{l}\,\overline{f}, \overline{l}\,\overline{f}, p) & t_5 = (l\overline{f}), l\overline{f}, p) \\ t_6 = (lf, lf, n) & t_7 = (\overline{l}f, \overline{l}f, n) & t_8 = (l\overline{f}, l\overline{f}, n) & t_4 = (\overline{l}\,\overline{f}, \overline{l}\,\overline{f}, n) & \end{array}$$

- $< \subseteq T \times T$ *given by*

$$t_1,\ t_3,\ t_6,\ t_7,\ t_8 \quad < \quad t_4,\ t_9 \quad < \quad t_2 \quad < \quad t_5$$

is an example for an ordered transition schema. For world $\overline{l}f$ and action p for example, the transition t_1 to lf is more plausible than the transition t_2 to $\overline{l}\,\overline{f}$ as $t_1 < t_2$. For an illustration of the transitions (without their ordering) cf. Figure 1.

Observe that every combination of query answers that are induced by a ranked transition schema can also be induced by an ordered transition schema. The reverse does not hold in general.

Proposition 2. *Let $\mathcal{T} = (V, T, \kappa)$ be a ranked transition schema. Construct the ordered transition schema $\mathcal{O} = (V, T, <)$ by letting $t < t'$ iff $\kappa(t) < \kappa(t')$. Then $\mathcal{T}$ answers a query $\omega, \alpha \rightsquigarrow e$ positively iff $\mathcal{O}$ answers it positively.*

Using the additional expressivity of ordered transition schemas we can define the following inference operator called *w-closure for actions*, that is inspired by the inference operator system W for reasoning from propositional conditionals [10] and is more fine-grained than rational closure.

At the core of w-closure is the following ordered transition schema.

Definition 12. *Let R be a set of defeasible action rules with the Z-partition $(R_1, \ldots, R_n, R_\infty)$. Let $\xi^i_R(t) = \{r \in R_i \mid r \text{ falsifies } t\}$ be the set of rules in R_i falsified by r.*

The W-ordered transition schema $\mathcal{O}^w_R = (V, E, <^w)$ is constructed as follows. Let $V = \Omega_\Sigma$ and let $T \subseteq V \times V \times A$ be the set of all transitions that do not falsify R_∞. For every $t_1 = (\omega_1, \omega'_1, \alpha_1), t_2 = (\omega_2, \omega'_2, \alpha_2)$ in T let $t_1 <^w t_2$ if there is an $i \in \{0, \ldots, n\}$ such that $\xi^i_R(t_1) \subsetneqq \xi^i_R(t_2)$ and $\xi^j_R(t_1) = \xi^j_R(t_2)$ for every $j > i$.

While $\mathcal{T}^z_R$ ranks worlds only based on what parts of the Z-partition contain a falsified rule, $\mathcal{O}^w_R$ also takes into what rules are falsified. Thus, $t <^w t'$ if and only if t falsifies strictly fewer defeasible action rules than t' in the latest part of the Z-partition where t and t' falsify different rules.

Example 6. *Consider again the set of rules R from Example 3. From the Z-partition of R we can get the w-ordered transition model $\mathcal{O}^w_R = (V, T, <^w)$. For the transitions $t_1, \ldots, t_5$ as illustrated in Figure 3 we obtain the following relations in $<^w$:*

$$\begin{tikzcd} t_2 \arrow[r] \arrow[dr] & t_1 \arrow[r] & t_5 \\ t_4 \arrow[r] \arrow[ur] & t_3 \arrow[ur] & \end{tikzcd}$$

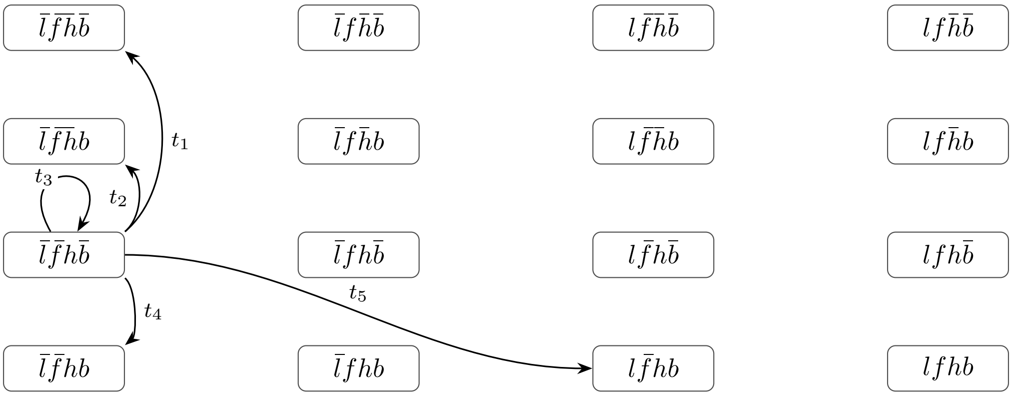

Figure 3: Illustration of selected transitions $t_1, \dots, t_5$ in the ordered transition schema $\mathcal{O}_R^w$ from Example 6.

We can verify that t_2 and t_4 each falsify r_1, while transitions t_1 and t_3 falsify both r_1 and r_3 making the latter two less plausible according to $<^w$. Finally, t_5 falsifies r_2 from R_2 making it the least plausible among the considered transitions.

$\mathcal{O}_R^w$ is faithful to the original set of rules R in that it is a model of R.

Proposition 3. *Let R be a set of defeasible action rules. Then $\mathcal{O}_R^w$ is a model of R.*

Proof. Let $\mathcal{O}_R^w = (V, T, <^w)$ be the w-ordered transition schema of R and let $(R_1, \dots, R_n, R_\infty)$ be the Z-partition of R. Let $r \in R$ be arbitrary.

If $r \in R_\infty$, then by construction no transition in $\mathcal{O}_R^w$ falsifies r and therefore $\mathcal{O}_R^w$ is a model of r.

For the remainder of the proof assume that $r \notin R_\infty$ implying that there is an $i \in \{1, \dots, n\}$ such that $r \in R_i$. By definition of exceptionality, there must be a transition t_v verifying r without falsifying any rule in $E_i = R_i \cup \dots \cup R_n \cup R_\infty$. For every transition t_f falsifying r it will be $\xi_R^i(t_v) = \emptyset \subsetneq \{r\} \subseteq \xi_R^i(t_f)$ and $\xi_R^j(t_v) = \emptyset \subseteq \xi_R^j(t_f)$ for all $j > i$; and therefore $t_v <^w t_f$. Thus, $\mathcal{O}_R^w$ is a model of r. □

With w-ordered transition schemas we can now define w-closure for defeasible action rules.

Definition 13. *Given a set of rules R,* w-closure for actions *answers the query $\omega, \alpha \rightsquigarrow e$ with true if it holds in $\mathcal{O}_R^w$.*

Example 7. *Consider the set R of defeasible action rules from Example 3, inducing the w-ordered transition schema $\mathcal{O}_R^z$ (cf. 6). How does w-closure answer the query $\omega_2, p \rightsquigarrow b$ for $\omega_2 = \bar{l}\bar{f}h\bar{b}$? Because all $<^w$-minimal transitions in $T_{\omega_2,p}$ lead to worlds in which h holds, this query is answered positively. Note that this is different from the answer rational closure gave for the same query (cf. Example 4).*

Example 7 illustrates that w-closure can sometimes derive more than rational closure. This is true in general, and we can show that w-closure is an extension of rational closure in the sense that it answers strictly more queries positively. The reverse is not true, as shown by the examples.

Proposition 4. *Let R be a set of defeasible action rules. If rational closure answers the query $\omega, \alpha \rightsquigarrow e$ with true, then W-closure also answers it with true.*

Proof. Let $\mathcal{T}_R^z = (V, T, \kappa^z)$ be the z-ranked transition schema of R, let $\mathcal{O}_R^w = (V, T, <^w)$ be the w-ordered transition schema of R, and let $(R_1, \dots, R_n, R_\infty)$ be the Z-partition of R. Let $q : \omega, \alpha \rightsquigarrow e$ be an arbitrary query that is answered positively by rational closure.

Let $t = (\omega, \omega', \alpha)$ be any transition in $\min_{<^w} T_{\omega,\alpha}$. Towards a contradiction assume that $\omega' \not\models e$. Because $\mathcal{T}_R^z$ answers the query positively, it must be that $\kappa_{\omega,\alpha}^z(\omega') > 0$. This implies that there is a transition $t^* \in T_{\omega,\alpha}$ with $\kappa_{\omega,\alpha}^z(t^*) = 0$ and thus $\kappa^z(t^*) < \kappa^z(t)$. Let $i = \kappa^z(t^*)$ and $j = \kappa^z(t)$. Because of $j > i \geq 0$ there must be an $r \in R_j$ that is falsified by t while $\xi_R^j(t^*) = \emptyset$ for every $k > i$. Thus, $\xi_R^j(t^*) = \emptyset \subsetneq \{r\} \subseteq \xi_R^j(t)$ and $\xi_R^j(t^*) = \emptyset \subseteq \xi_R^j(t_f)$ for all $j > i$; and therefore $t^* <^w t$. Contradiction, because t is minimal in $T_{\omega,\alpha}$ with respect to $<^w$. □

In summary, ordered transition schemas are a more versatile and more expressive alternative to ranked transition schemas. However, this additional expressivity comes at a cost: ordered transition schemas do not support the arithmetic operations offered by ranks. For example, the intuitive combination of ranks for reasoning about traces is not possible with rankings. Ordering traces will therefore require the development of more involved techniques.

6. Keeping Things as They Are

The defeasible action rules introduced so far describe what typically holds after an action is executed and thus license changes in the world. But we have not discussed what should stay unchanged. In classical action formalisms this is known as the *frame problem*: the need to specify which properties of the world persist across an action [2]. In this section we discuss ways of incorporating frame rules into our defeasible action rule-framework.

One approach to solve this would be to modify the inference operators to prefer transitions that introduce fewer changes to the world, thus realizing a kind of minimal change principle. While this can be surely realized using ranked or ordered transition schemas, this is fiddly to integrate into defeasible reasoning operators. Furthermore, it is unflexible as it fixes the minimal change in the inference operator, not giving the designer of a knowledge base control over what should remain unchanged.

A second approach would be to express frame conditions as defeasible action rules of the form $f, \alpha \rightsquigarrow \{f\}$ asserting that f typically holds after α whenever it held before. This, however, runs into a problem when we try to use it in context of other defeasible action rules. The rule would have to be repeated for every f we do not want to change. At the same time such a rule cannot be overridden by another rule with precondition f and action α that does change the valuation of f, since none of the rules is more specific. Therefore, general rules expressing "in all situations, don't change a unless necessary" interact poorly with other defeasible action rules.

We therefore introduce a third approach, that extends the language of defeasible action rules with a dedicated notation for a formula that is not changed by an action. For this, we add a new type of effect expression $nc(f)$ expressing that f is the same before and after the execution of an action.

Definition 14. *A* frame rule *is a rule of the form*

$$p, \alpha \rightsquigarrow F$$

where p is a formula, $\alpha \in A$, and F is a set that may contain formulas as well as expressions of the form $nc(f)$ for formulas f.

Example 8. *Consider again the agent from Example 1. To represent that the button usually stays clean (unless another rule tells us otherwise) we can use the following frame rule:*

$$\top, push\_button \rightsquigarrow \{nc(button\_dirty)\}. \tag{r_4}$$

Frame rules generalize defeasible action rules: every defeasible action rule is trivially a frame rule with no $nc(\cdot)$ expressions in its effects. For a frame rule $p, \alpha \rightsquigarrow F$ let F^{clas} denote all formulas in F and F^{nc} denote all statements of the form $nc(f)$ in F. The verification and falsification conditions for transitions are then extended to account for this richer effect language. For a given transition $t = (\omega, \omega', \alpha)$ and a frame rule $r = p, \alpha \rightsquigarrow F$ we say

- t *verifies* r if $\omega \models p$ and $\omega' \models F^{clas}$ and for every $nc(f) \in F^{nc}$ it holds that $\omega \models f$ iff $\omega' \models f$.
- t *falsifies* r if $\omega \models p$ but t does not verify the rule, and
- r is *not applicable* to t if $\omega \not\models p$.

This three-valued semantics mirrors that of ordinary defeasible action rules, with the additional condition that $nc(f)$-statements require a matching evaluation of f in the worlds before and after the action.

The modelling condition for ranked transition schemes carries over without modification: A ranked transition scheme $\mathcal{T} = (V, T, \kappa)$, models a frame rule $r = p, \alpha \rightsquigarrow F$ if there is no transition $t_f \in T$ falsifying r, or if there is at least one transition $t_v \in T$ verifying r and $\min_{\substack{t_v \in T, \\ t_v \text{ verifies } r}} \kappa(t_v) < \min_{\substack{t_f \in T, \\ t_f \text{ falsifies } r}} \kappa(t_f)$. Queries can similarly be extended to ask whether a property is typically preserved by an action: in the context of a ranked transition scheme $\mathcal{T} = (V, T, \kappa)$, a query $\omega, \alpha \rightsquigarrow nc(f)$ is answered positive if for every $t = (\omega, \omega', \alpha)$ in $T_{\omega,\alpha}$ with $\kappa_{\omega,\alpha}(t) = 0$ we have $\omega \models f$ iff $\omega' \models f$.

This holds analogous for ordered transition schemes: An ordered transition scheme $\mathcal{O} = (V, T, <)$, models a frame rule $r = P, \alpha \rightsquigarrow F$ if for every $t_f \in T$ that falsifies r there is a $t_v \in T$ that verifies r with $t_v < t_f$. A query $\omega, \alpha \rightsquigarrow nc(f)$ is answered positive in the context of an ordered transition scheme $\mathcal{O} = (V, T, <)$ if for every transition $(\omega, \omega', \alpha) \in \min_< T_{\omega,\alpha}$ we have $\omega \models f$ iff $\omega' \models f$.

Using these new definitions, we can adapt the inference operators considered in Sections 4 and 5 to frame rules. The most notable change is adapting the notion of exceptionality that underlies both rational closure and w-closure. The key to this is in evaluating the $nc(\cdot)$ effects depending on the considered world.

Definition 15. *Let R be a set of frame rules. For an effect e let*

$$eval(e, \omega) = \begin{cases} e & \text{if } e \text{ is a formula} \\ f & \text{if } e = nc(f) \text{ and } \omega \models f \\ \neg f & \text{if } e = nc(f) \text{ and } \omega \models \neg f. \end{cases}$$

For a set F of formulas and $nc(\cdot)$ statements define $eval(\omega, F) = \{eval(\omega, e) \mid e \in F\}$.

A frame rule $(p, \alpha \rightsquigarrow F) \in R$ is exceptional *in R if there is no world $\omega \in \Omega_\Sigma$ and action $\alpha \in A$ such that $\omega \models p$ and $F \cup \bigcup_{(P', \alpha \rightsquigarrow F') \in R,\ \omega \models P} eval(F')$ is consistent. $\mathcal{F}(R)$ denotes the set of rules in R that are exceptional in R.*

Using the updated definitions of exceptionality and query answering, we can use the definitions for the Z-partition, rational closure and w-closure without change.

The following proposition confirms that the adapted definitions are a conservative extension of the earlier framework.

Proposition 5. *For sets of frame rules that do not contain $nc(\cdot)$ expressions,*

- *rational closure for frame rules coincides with rational closure for defeasible action rules, and*
- *w-closure for frame rules coincides with w-closure for defeasible action rules.*

This follows immediately from the observation that for classical effect formulas e we have $eval(e, \omega) = e$ for all ω, so the consistency condition in Definition 15 reduces to that of Definition 5. The following example indicates that w-closure may be more suitable for reasoning with frame rules than rational closure.

Example 9. *Consider the set $R = \{r_1, r_2, r_3, r_4\}$ of rules from Examples 1 and 8 with abbreviated symbols for atoms and actions as in Example 3. The Z-partition of R is R_1, R_2 with $R_1 = \{r_1, r_4\}$ and $R_2 = \{r_2, r_3\}$. For the world $\omega = \overline{l}\,\overline{f}\,\overline{h}\,\overline{b}$ the query $\omega, p \rightsquigarrow \neg b$ is answered negatively. Because any transition complying with r_2 must violate r_1, the lowest rank a transition in $T_{\omega,p}$ can have is 1, sidelining the rule r_4. This is an instance of the* drowning problem *[17] that also affects rational closure in the context of propositional conditionals. w-Closure however does answer the query $\omega, p \rightsquigarrow \neg b$ positively, as it prefers transitions complying with r_4 even if r_1 is falsified.*

7. Summary and Future Work

In this paper we have presented a framework for defeasible reasoning about actions and their effects. Motivated by the observation that real-world actions often fail to produce their typical outcome in exceptional circumstances, we introduced defeasible action rules as a formalism for capturing the expected effects of actions while still allowing for other outcomes in exceptional cases. The semantics of these rules is given by ranked transition schemas, which assign ranks to transitions between worlds to represent their typicality. Ordered transition schemas are introduced as an alternative semantic based on strict partial orderings. Building on this semantic foundation, we defined two inference operators: rational closure for action rules and w-closure. Both based on counterparts for conditional reasoning, these inference operators support defeasible reasoning and exceptional rules overriding more general rules. Finally, we extended our framwork with frame rules, which add the notation $nc(f)$ to the set of effects, allowing to express that certain properties are preserved across an action.

This paper opens several directions for future work. The most immediate is a thorough investigation of the formal properties of the two inference operators with respect to postulates and complexity. In particular developing postulates describing desirable properties and seeing if they are satisfied would interesting. A second direction concerns the connection to classical planning: defeasible action rules with literal effects have a natural operational reading reminiscent of STRIPS operators, and it would be valuable to investigate this connection formally. The ranked and ordered transition schemas also open up richer notions of planning, such as safe plans that minimise the risk of undesirable actions. Finally, the relationship between the present framework and probabilistic approaches to uncertain action deserves further investigation, in particular to understand whether the qualitative rankings can be seen as an abstraction of an underlying probability distribution.

Acknowledgments

This work was supported in part by the National Research Foundation of South Africa (REFERENCE NO: SAI240823262612) and the Austrian Science Fund (FWF) projects 10.55776/COE12 and 10.55776/PIN8884924.

Declaration on Generative AI

During the preparation of this work, the authors used Claude Sonnet 4.6 in order to draft content as well as paraphrase and reword text. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] R. Fikes, N. J. Nilsson, STRIPS: A new approach to the application of theorem proving to problem solving, Artif. Intell. 2 (1971) 189–208. URL: https://doi.org/10.1016/0004-3702(71)90010-5. doi:10.1016/0004-3702(71)90010-5.

[2] J. McCarthy, P. Hayes, Some philosophical problems from the standpoint of artificial intelligence, in: B. L. Webber, N. J. Nilsson (Eds.), Readings in Artificial Intelligence, Morgan Kaufmann, 1981, pp. 431–450. URL: https://www.sciencedirect.com/science/article/pii/B9780934613033500337. doi:https://doi.org/10.1016/B978-0-934613-03-3.50033-7.

[3] M. Thielscher, The qualification problem: A solution to the problem of anomalous models, Artificial Intelligence 131 (2001) 1–37. URL: https://www.sciencedirect.com/science/article/pii/S000437020100131X. doi:https://doi.org/10.1016/S0004-3702(01)00131-X.

[4] J. McCarthy, Epistemological problems of artificial intelligence, in: Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'77, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1977, p. 1038–1044.

[5] R. Bellman, Dynamic programming and stochastic control processes, Information and Control 1 (1958) 228–239. URL: https://www.sciencedirect.com/science/article/pii/S0019995858800030. doi:https://doi.org/10.1016/S0019-9958(58)80003-0.

[6] L. P. Kaelbling, M. L. Littman, A. R. Cassandra, Planning and acting in partially observable stochastic domains, Artificial Intelligence 101 (1998) 99–134. URL: https://www.sciencedirect.com/science/article/pii/S000437029800023X. doi:https://doi.org/10.1016/S0004-3702(98)00023-X.
[7] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artif. Intell. 44 (1990) 167–207.
[8] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artif. Intell. 55 (1992) 1–60. URL: https://doi.org/10.1016/0004-3702(92)90041-U. doi:10.1016/0004-3702(92)90041-U.
[9] J. Pearl, System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: R. Parikh (Ed.), Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1990, Morgan Kaufmann, 1990, pp. 121–135.
[10] C. Komo, C. Beierle, Nonmonotonic reasoning from conditional knowledge bases with system W, Ann. Math. Artif. Intell. 90 (2022) 107–144. URL: https://doi.org/10.1007/s10472-021-09777-9. doi:10.1007/s10472-021-09777-9.
[11] W. Spohn, Ordinal conditional functions: a dynamic theory of epistemic states, in: W. L. Harper, B. Skyrms (Eds.), Causation in decision, belief change, and statistics: Proceedings of the Irvine Conference on Probability and Causation, 2, Kluwer, Dordrecht, 1988, pp. 105–134. URL: https://doi.org/10.1007/978-94-009-2865-7_6. doi:10.1007/978-94-009-2865-7_6.
[12] V. Lifschitz, Formal theories of action, in: F. M. BROWN (Ed.), The Frame Problem in Artificial Intelligence, Morgan Kaufmann, 1987, pp. 35–57. URL: https://www.sciencedirect.com/science/article/pii/B9780934613323500095. doi:https://doi.org/10.1016/B978-0-934613-32-3.50009-5.
[13] L. Giordano, A. Martelli, C. Schwind, Ramification and causality in a modal action logic, Journal of Logic and Computation 10 (2000) 625–662. doi:10.1093/logcom/10.5.625.
[14] G. Brewka, T. Eiter, M. Truszczynski, Answer set programming at a glance, Commun. ACM 54 (2011) 92–103. URL: https://doi.org/10.1145/2043174.2043195. doi:10.1145/2043174.2043195.
[15] B. de Finetti, La prévision, ses lois logiques et ses sources subjectives, Ann. Inst. H. Poincaré 7 (1937) 1–68. Engl. transl. *Theory of Probability*, J. Wiley & Sons, 1974.
[16] J. Haldimann, T. Meyer, G. Kern-Isberner, C. Beierle, Rational closure extension in spo-representable inductive inference operators, in: S. A. Gaggl, M. V. Martinez, M. Ortiz (Eds.), Logics in Artificial Intelligence - 18th European Conference, JELIA 2023, Dresden, Germany, September 20-22, 2023, Proceedings, volume 14281 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 561–576. URL: https://doi.org/10.1007/978-3-031-43619-2_38. doi:10.1007/978-3-031-43619-2_38.
[17] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: R. Bajcsy (Ed.), Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993, Morgan Kaufmann, 1993, pp. 640–647. URL: http://ijcai.org/Proceedings/93-1/Papers/090.pdf.