

Non-deterministic Action Reversibility: Complexity Results

Jakub Med^{1,2}, Michael Morak³, Lukáš Chrpa¹, Wolfgang Faber³

¹Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague

²Faculty of Electrical Engineering, Czech Technical University in Prague

³University of Klagenfurt

{jakub.med, chrpaluk}@cvut.cz, {michael.morak, wolfgang.faber}@aau.at

Abstract

With the recent interest in the reversibility of action effects, i.e., whether the effects of the action can be undone by applying other actions, the question arose how hard it is to reverse an action in a non-deterministic domain. With the use of φ -reversibility, the paper investigates the computational complexity of weak and strong non-deterministic action reversibility in fully observable non-deterministic domains, showing PSPACE-completeness for all weak variants in question and EXPTIME-hardness and EXPTIME, or NEXPTIME memberships for strong variants.

1 Introduction

In *Automated Planning* the central task concerns finding a plan that transforms an initial state to some goal state (Ghalab, Nau, and Traverso 2016). However, an interesting avenue of research is to reason about properties of actions in order to investigate the dynamics of the system or guide the search for such a solution. One such property is *action reversibility*, that is, whether the effects of an action can be undone by other actions (Daum et al. 2016; Morak et al. 2020; Chrpa, Faber, and Morak 2021; Med et al. 2024).

Action reversibility offers several valuable advantages in automated planning, motivating its study across various applications. Primarily, reversible actions ensure safety by avoiding dead-end states, crucial for online planning (Cserna et al. 2018), and suggesting applications in robotics to verify recovery from failures. Reversibility can aid post-planning plan optimization by identifying redundant action cycles (Med and Chrpa 2022). In fully observable non-deterministic (FOND) planning, it helps assess recoverability from undesirable effects, potentially enhancing strong cyclic plan generation (e.g., by integration with the PRP planner (Muise, McIlraith, and Belle 2014; Camacho, Muise, and McIlraith 2016)) and determining the safety of deterministic replanning like FF-replan (Yoon, Fern, and Givan 2007). Furthermore, it is essential in planning with exogenous events to infer the robustness of the plan (Chrpa, Pilát, and Med 2021; Chrpa and Karpas 2024b; Chrpa and Karpas 2024a). In cybersecurity, action reversibility offers potential as a remedy for attacks (by extending (Boddy et al. 2005)), and allows classification of actions by risk. The utility of reversibility for verification of security protocols has been

demonstrated in cloud management, confirming the undoability of system operations (or informing the user about permanent and irreversible changes in the system) (Weber et al. 2012; Weber et al. 2013). Further research on action reversibility could yield more efficient computations, increasing scalability, and enabling broader practical applications.

Deciding whether a deterministic action is φ -reversible is PSPACE-complete (Morak et al. 2020). That is the same as deciding plan existence for propositional classical planning (Bylander 1994). Morak et al. (2020) proved that the uniform variant is PSPACE-hard, but no upper bound was provided. They also studied classes of action φ -reversibility where the plans have a polynomially bounded length.

The recent work of Med et al. (2024), which introduced weak and strong action reversibility of non-deterministic actions in fully observable settings, left the complexity questions open. Regarding non-deterministic planning in general, Littman, Goldsmith, and Mundhenk (1998) have shown that the probabilistic planning problem with success chance greater than $\theta \in (0, 1)$ is EXPTIME-complete. Rintanen (2004) has given an alternative proof of EXPTIME-completeness of FOND planning. This paper aims to address the knowledge gap by establishing computational complexity results for weak and strong φ -reversibility by deriving their lower and upper bounds.

2 Preliminaries

Formalism. In order to align with the literature, we adopt notation from Med et al. (2024), that is, a combination of SAS⁺ formalism (Bäckström and Nebel 1995) and of FOND planning used by Cimatti et al. (2003), as well as standard computational complexity notation (Papadimitriou 1994).

Let v be a **variable** (or a **state variable**) and $dom(v)$ be its domain. A **fact** (v, x) is a pair of a variable v and its value $x \in dom(v)$. Let $\mathcal{F}(\mathcal{V})$ be the **set of all facts** over the set of variables $\mathcal{V}$. A set $\Sigma \subseteq \mathcal{F}(\mathcal{V})$ is called a **variable assignment** over $\mathcal{V}$ if and only if for all $(v, x) \in \Sigma$ no $y \in dom(v)$ exists such that $y \neq x$ and $(v, y) \in \Sigma$. Let $vars(\Sigma)$ be the **set of all variables** in Σ , i.e., $\{v \mid (v, x) \in \Sigma\}$. A variable assignment Σ over $\mathcal{V}$ is called **complete** if and only if $vars(\Sigma) = \mathcal{V}$. Otherwise, it is **partial**. The **set of all complete variable assignments** over $\mathcal{V}$ is denoted as $\mathcal{S}(\mathcal{V})$.

An **action** a over a set of variables $\mathcal{V}$ is a pair $(pre(a), eff(a))$, where $pre(a)$ is a variable assignment over $\mathcal{V}$ rep-

representing the **precondition** of the action a , and $\text{eff}(a)$ is a non-empty set of variable assignments over $\mathcal{V}$ representing the **set of possible effects** of a . The action a is called **deterministic** if and only if $|\text{eff}(a)| = 1$. A **determinization** of a with respect to the effect $e \in \text{eff}(a)$, denoted as a_e^d , is an action $a_e^d = (\text{pre}(a), \{e\})$. An action a is **applicable** in a variable assignment Σ if and only if $\text{pre}(a) \subseteq \Sigma$. The **application of a deterministic action** $a = (\text{pre}(a), \{e\})$ in a variable assignment Σ is the variable assignment $\gamma(\Sigma, a) = \{(v, x) \in \Sigma \mid v \notin \text{vars}(e)\} \cup e$ if a is applicable in Σ (undefined otherwise). The **application of a (non-deterministic) action** a in a variable assignment Σ is the set $\delta(\Sigma, a) = \{\gamma(\Sigma, a_e^d) \mid e \in \text{eff}(a)\}$.

A **FOND planning domain** $\mathcal{D}$ is a pair $\langle \mathcal{V}, \mathcal{A} \rangle$, where $\mathcal{A}$ is a set of actions over $\mathcal{V}$. A **state** s of the domain $\mathcal{D}$ is a complete variable assignment over the set $\mathcal{V}$. A **FOND planning task** $\mathcal{T}$ is a triple $\langle \mathcal{D}, s_I, G \rangle$, where $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$ is a FOND planning domain, $s_I \in \mathcal{S}(\mathcal{V})$ is an **initial state** and $G \subseteq \mathcal{F}(\mathcal{V})$ is a variable assignment representing a **goal**.

A **policy**¹ Π for the domain $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$ is a binary relation such that $\Pi \subseteq \{(s, a) \mid s \in \mathcal{S}(\mathcal{V}), a \in \mathcal{A}, \text{pre}(a) \subseteq s\}$. Note it may relate only pairs in which the action is applicable. The **set of all states related in Π** is the set $\sigma(\Pi) = \{s \mid (s, a) \in \Pi\}$. The **n -step application** of Π in a state s , is the set $\delta^n(s, \Pi) = \bigcup_{s' \in \delta^{n-1}(s, \Pi)} \bigcup_{(s', a) \in \Pi} \delta(s', a)$ for $n \geq 1$, and $\delta^0(s, \Pi) = \{s\}$. We say that s' is **reachable** from s with a policy Π if and only if $s' \in \delta^i(s, \Pi)$ for some $i \geq 0$. We say that s' is a **terminal state** for Π w.r.t. a state s if and only if s' is reachable from s with Π and $s' \notin \sigma(\Pi)$. The **set of all such terminal states** is denoted $\tau(\Pi, s)$.

A policy Π is called a **weak goal policy** for $\mathcal{T}$ if and only if $\tau(\Pi, s_I) \cap \{s_G \in \mathcal{S}(\mathcal{V}) \mid G \subseteq s_G\} \neq \emptyset$. A policy Π is called a **strong goal policy** for $\mathcal{T}$ if and only if $\tau(\Pi, s_I) \subseteq \{s_G \in \mathcal{S}(\mathcal{V}) \mid G \subseteq s_G\}$ and for each $s \in \sigma(\Pi)$ at least one state that satisfies the goal G is reachable by Π .

Reversibility. Med et al. (2024) defined weak and strong non-deterministic action S -reversibility as follows. The definitions are inspired by the weak and strong solution concepts of FOND planning tasks (Cimatti et al. 2003), i.e., *weak reversibility* refers to the possibility of undoing action effects, whilst *strong reversibility* guarantees undoing effects of the non-deterministic action.

Definition 1. Let $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$ be a planning domain, a be an action over the set of variables $\mathcal{V}$, $S \subseteq \mathcal{S}(\mathcal{V})$ be a set of states of $\mathcal{D}$, Π be the policy for $\mathcal{D}$ and φ be a propositional formula over $\mathcal{F}(\mathcal{V})$. The action a is called:

- **weakly (resp. strongly) S -reversible** in $\mathcal{D}$ if and only if for each state $s \in S$, $\text{pre}(a) \subseteq s$, there exists a policy Π_s (called **weak (resp. strong) $\{s\}$ -reverse policy**) for $\mathcal{D}$ such that for each $s' \in \delta(s, a)$ holds $s \in \tau(\Pi_s, s')$ (resp. $\{s\} = \tau(\Pi_s, s')$) and for each state s'' that is reachable by Π_s from s' it holds that s is reachable by Π_s from s'');
- **weakly (resp. strongly) uniformly S -reversible** in the domain $\mathcal{D}$ by the policy Π if and only if Π is weak (resp.

strong) $\{s\}$ -reverse policy for each $s \in S$, $\text{pre}(a) \subseteq s$;

- **weakly (resp. strongly) (uniformly) universally reversible** if and only if the action a is weakly (resp. strongly) (uniformly) S -reversible and $S = \mathcal{S}(\mathcal{V})$;
- **weakly (resp. strongly) (uniformly) φ -reversible** if and only if the action a is weakly, (resp. strongly) (uniformly) S -reversible such that $S = \{s \mid s \in \mathcal{S}(\mathcal{V}), s \models \varphi\}$.

Informally speaking, S -reversibility is the notion that restricts the investigation of the action reversibility to the set of states S in which the action in question is applied. *Uniform reversibility* requires that there is a common policy that works in each state of the set S in which the action in question can be applied. Intuitively, it is a stronger notion than the non-uniform variant, where there might be a specialized policy for each state from S . *Universal reversibility* is a special case where the set S contains all possible states in the domain. φ -reversibility allows a compact representation of states instead of a possibly exponentially large set S . For further discussion on the differences between the classes of action reversibility, see Med et al. (2024, Examples 1-4).

Decision Problems. Throughout this work, we only consider decision problems in which the set S is represented as a formula φ . Next, we define the two most general of the studied decision problems, weak and strong φ -reversibility.

WEAK- φ - $\mathcal{R}$ (resp. STRONG- φ - $\mathcal{R}$)

Instance: A domain $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$, an action $a \in \mathcal{A}$, a formula φ over $\mathcal{F}(\mathcal{V})$.

Question: Is the action a weakly (resp. strongly) φ -reversible in the domain $\mathcal{D}$?

Decision problems corresponding to universal variants, uniform variants, and their combinations are defined analogously. Notice that universal variants are special cases of the general φ -reversibility problems where $\varphi \equiv \top$.

3 Complexity Results

Weak Reversibility. We start with the hardness of *weak universal uniform reversibility*, as it builds upon the knowledge acquired by a study on deterministic action reversibility. It shows how deterministic universal uniform reversibility generalizes to weak universal uniform reversibility, thus proving hardness.

Proposition 2. WEAK-UNIVERSAL-UNIFORM- $\mathcal{R}$ is PSPACE-hard, which implies the same for WEAK-UNIFORM- φ - $\mathcal{R}$, WEAK-UNIVERSAL- $\mathcal{R}$, and WEAK- φ - $\mathcal{R}$.

Proof. Consider a planning domain that only consists of deterministic actions (i.e. where all actions have only one set of effects). In this case, the universal uniform reversibility for STRIPS studied by Morak et al. (2020) is a special case of our WEAK-UNIVERSAL-UNIFORM- $\mathcal{R}$ problem. Thus, we inherit its PSPACE-hardness. Moreover, since universal reversibility is a special case of φ -reversibility, hardness is immediately inherited by WEAK-UNIFORM- φ - $\mathcal{R}$. As hardness for deterministic reversibility already holds for an action applicable in a single state only, and since then, the notions of uniform and non-uniform reversibility coincide, the

¹A policy specifies which actions are applied in which states. Intuitively, it represents a “non-deterministic plan”.




Figure 1: Sketch of construction of $\mathcal{T}_s$ in proof of Theorem 6.

same hardness also follows for WEAK-UNIVERSAL- $\mathcal{R}$ and WEAK- φ - $\mathcal{R}$. $\square$

Hence, hardness for all considered variants of weak reversibility follows easily from reversibility in STRIPS. However, notice that while weak reversibility generalizes reversibility in the deterministic case, there is no one-to-one correspondence between reverse plans (as defined for the deterministic setting) and reverse policies as defined in this paper. For example, a uniform reverse plan allows us to immediately construct an equivalent uniform reverse policy. However, the converse is in general not true, since a policy may apply a different set of actions in each state.

The lemma below is related to this particular observation. It can be viewed as a blueprint for combining several policies together to achieve uniformity. On the other hand, it also claims that once there is some weak uniform reverse policy for at least two states, it can be decomposed into two separate policies for each one. The lemma is derivable by using Theorem 7 of Med et al. (2024).

Lemma 3. *Let $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$ be a planning domain, $S_1, S_2 \subseteq \mathcal{S}(\mathcal{V})$ be a set of states, $a \in \mathcal{A}$ be an action, Π_1 be a weak S_1 -reverse policy for a and Π_2 be a weak S_2 -reverse policy for a . The policy $\Pi_1 \cup \Pi_2$ is a weak $(S_1 \cup S_2)$ -reverse policy for the action a if and only if $\{s \mid s \in S_1, \text{pre}(a) \subseteq s\} \cap \sigma(\Pi_2) = \emptyset$ and $\{s \mid s \in S_2, \text{pre}(a) \subseteq s\} \cap \sigma(\Pi_1) = \emptyset$.*

Lemma 3 will allow us to prove PSPACE-membership of weak uniform φ -reversibility by constructing individual reverse policies for each relevant starting state and merging them to form a uniform reverse policy.

Theorem 4. WEAK-UNIFORM- φ - $\mathcal{R}$ is in PSPACE, which implies the same for WEAK-UNIVERSAL-UNIFORM- $\mathcal{R}$, WEAK- φ - $\mathcal{R}$, and WEAK-UNIVERSAL- $\mathcal{R}$.

Proof. We want to decide uniform φ -reversibility for action a in planning domain $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$. Let e be an effect of a . Existence of an $\{s\}$ -reverse policy for a_e^d that does not visit any states satisfying φ and $\text{pre}(a)$ can be decided using the following NPSpace procedure: Let s' be the current state. Guess an action a' and an effect e' of a' . Verify that applying a_e^d does not lead to a state s'' with $s'' \models \varphi$ and $\text{pre}(a) \subseteq s''$. If so, answer no. If $s'' = s$, answer yes. Otherwise, repeat with $s' = s''$. The procedure starts with $s' = s$, and the first chosen action and effect is forced to be a and e . This decides $\{s\}$ -reversibility of a_e^d with a policy as in Lemma 3. But then there is an NPSpace procedure that establishes φ -reversibility of a : Simply decide the above for all effects e of a and states s where $s \models \varphi$. If and only if all answers are yes, then, per Lemma 3, there must be a φ -reverse policy for




Figure 2: Sketch of construction of $\mathcal{D}'$ in the proof of Theorem 5.

a combining all these individual $\{s\}$ -reverse policies. Via Savitch's theorem (which implies PSPACE = NPSpace) (1970) the desired result follows. Since universal reversibility is a special case of φ -reversibility, and since uniform reversibility implies (non-uniform) reversibility, the upper bound also holds for these variants of reversibility. $\square$

With all upper bounds addressed, we have proved PSPACE-completeness of all variants of weak reversibility. Similarly to weak planning, weak reversibility revolves around whether the state in which the action has been applied can possibly be reached from the resulting states. However, this reachability is checked for many different states at once, either via separate policies for each state (non-uniformity) or by a shared multifunctional policy (uniformity). Yet, our results show that this does not cause a difference from a computational complexity perspective.

Strong Reversibility. Also for *strong reversibility*, at the heart of the problem lie several planning tasks, one for each potential starting state. However, these tasks now need to check for strong solutions, that is, whether the goal state is reached no matter which action effects occur.

We again start with hardness for the universal uniform case. The idea of the proof is an adaptation of the proof provided by Morak et al. (2020) in the deterministic setting.

Theorem 5. STRONG-UNIVERSAL-UNIFORM- $\mathcal{R}$ is EXPTIME-hard. This implies the same for STRONG-UNIFORM- φ - $\mathcal{R}$, STRONG-UNIVERSAL- $\mathcal{R}$, and STRONG- φ - $\mathcal{R}$.

Proof Sketch. The proof is via reduction from the existence of a strong cyclic solution. The idea is sketched in Figure 2. Given a planning task $\mathcal{T} = \langle \mathcal{D}, s_I, \mathcal{G} \rangle$ we need to check whether there is a strong goal policy for $\mathcal{T}$. In order to convert this into a reversibility task, we introduce a fresh, special value “undefined” for each variable. An action a_{start} is added, which is only applicable wWhen all variables are undefined and it sets their values to the starting state s_I . An action a_{end} is added, which is only applicable in states where $\mathcal{G}$ holds, and which again sets all variables to undefined. The original actions remain the same. Clearly, the procedure runs in polynomial time.

In order to decide the existence of a strong goal policy for $\mathcal{T}$ via strong universal uniform reversibility, we simply ask whether the action a_{start} is strongly reversible. Since a_{start} is only applicable in single state θ , where all variables are undefined, it is strongly universally uniformly reversible if and

only if there is a strong goal policy for $\mathcal{T}$, which leads from s_I the goal $\mathcal{G}$. Hence, EXPTIME-hardness is inherited from the existence of a strong cyclic solution (Rintanen 2004).

For φ -reversibility and the non-uniform versions, recall that the above hardness result already holds for a single starting state θ . But in this case, the hardness is immediately inherited by these variants, proving the desired result. $\square$

Since solving multiple planning tasks lies at the heart of reversibility, it is clear that the relevant lower bounds will be inherited from the corresponding planning problems. However, up to exponentially many planning tasks need to be solved within a single reversibility task. For non-uniform reversibility, we can solve these tasks separately, resulting in exponential time overall.

Theorem 6. STRONG- φ - $\mathcal{R}$ is in EXPTIME.

Proof Sketch. Given a planning domain $\mathcal{D} = \langle \mathcal{V}, \mathcal{A} \rangle$ where strong φ -reversibility of action a should be tested, we construct, for each starting state $s \in \mathcal{S}(\mathcal{V})$ such that $\text{pre}(a) \subseteq s$ and $s \models \varphi$, a planning task $\mathcal{T}_s$ in the following way: (1) in $\mathcal{T}_s$ we require that the first action that is applied is a ; and (2) where the goal is s . Condition (1) is achieved by adding a fresh variable v_a with possible values $\top$ and $\perp$ and augmenting action a by adding as a precondition that v_a must be $\perp$, modifying all its effects so that v_a is set to $\top$, and augmenting all other actions to require v_a to be $\top$. We establish (2) by adding that v_a must be $\top$ to our goal; see Figure 1.

In this way, strong $\{s\}$ -reversibility can be checked by solving the non-deterministic planning task $\mathcal{T}_s$, where strong $\{s\}$ -reversibility holds if and only if a strong goal policy exists for $\mathcal{T}_s$, proving EXPTIME-membership for $\{s\}$ -reversibility by inheriting the upper bound from strong planning (Rintanen 2004). Since $\mathcal{S}(\mathcal{V})$ is of exponential size, we can solve strong $\{s\}$ -reversibility for each state $s \in \mathcal{S}(\mathcal{V})$ in turn, taking exponential time overall. $\square$

When $\varphi \equiv \top$, we immediately obtain the following:

Corollary 7. STRONG-UNIVERSAL- $\mathcal{R}$ is in EXPTIME.

The above establishes EXPTIME-completeness for the non-uniform versions of strong reversibility. We will now turn our attention to the uniform variants. The following upper bound can be obtained straightforwardly.

Theorem 8. STRONG-UNIFORM- φ - $\mathcal{R}$ and STRONG-UNIVERSAL-UNIFORM- $\mathcal{R}$ are in NEXPTIME.

Proof. A non-deterministic exponential time procedure for deciding the above problems works as follows: First, guess a strong reverse policy Π for action a (which is of exponential size w.r.t. the input planning domain). For each state s , check in exponential time whether Π reverts action a and whether $s \models \varphi$ in case of φ -reversibility. If all checks succeed, answer *yes*, otherwise *no*. Since there are exponentially many states s , overall, these checks take exponential time. This establishes NEXPTIME-membership. $\square$

Unfortunately, this leaves a gap in our complexity investigation. However, from our investigation, it seems non-trivial to close, as both a membership proof for EXPTIME and a

Problem	Complexity
Weak reversibility	PSPACE-complete
Weak uniform reversibility	PSPACE-complete
Strong reversibility	EXPTIME-complete
Strong uniform reversibility	EXPTIME-hard, in NEXPTIME

Table 1: Summary of computational complexity results. All bounds hold for all types of reversibility: S -reversibility, φ -reversibility, and universal reversibility, respectively.

hardness proof for NEXPTIME seem to run into structural issues w.r.t. how strong uniform reversibility works.

Notice that in the strong reversibility setting, it might be the case that a (uniform) reverse policy for two starting states, say s_1 and s_2 , exists, but then, after applying some actions, it may end up in a single state s for both, from which there are then several paths leading back to s_1 and s_2 . While this is not a problem for weak reversibility, it renders the policy invalid for the purposes of strong reversibility. In the latter case, the policy must guarantee that state s_1 is always reverted back to s_1 and the same for s_2 . But a shared state s renders this impossible: the policy cannot distinguish whether it should now return to s_1 or s_2 from s (cf. Med et al. (2024, Theorem 8)). This is because the policy decides the next action based on the current state only. Hence, strong uniform reversibility offers not only the guarantee that the initial state is restored in all cases, but also that individual parts of the policy do not interfere with each other.

Hence, we can verify in EXPTIME that a given policy works, but it seems more difficult to establish the existence of such a policy, since there are superexponentially many candidate policies for a given reversibility task. Attempting to construct non-interfering partial policies for each state that can later be combined to form a single strong reverse policy seems to require more effort than can be spent in EXPTIME. On the other hand, this clear separation and a complete independence of partial policies for each state also seem to prevent NEXPTIME-hardness from being established. This mirrors a complexity-theoretic gap which exists already for deterministic STRIPS planning tasks: namely, that uniform reversibility of STRIPS actions with polynomial-length plans is harder than NP (in fact it is hard for the complexity class of UNiquesat), but seemingly not hard for Σ_2^P (Morak et al. 2020). From our investigation, it seems difficult to establish completeness for one of the standard complexity classes for strong uniform reversibility.

4 Conclusions

We established PSPACE-completeness of weak reversibility for FOND planning, as well as the EXPTIME-completeness for strong reversibility in general (i.e. its non-uniform variant). In the uniform case, strong reversibility seems to fall between EXPTIME and NEXPTIME, where we established a lower and an upper bound, respectively (see Table 1).

In the future, we plan to investigate refined notions of reversibility, e.g. polynomial-size or acyclic policies. If possible, we plan to further narrow down the complexity of strong uniform reversibility.

Acknowledgments

This research was partially supported by the Czech Science Foundation (project no. 24-13337L) and Austrian Science Fund (FWF) under grant numbers 10.55776/PIN8782623 and 10.55776/COE12, under the joint bilateral project, and by the Grant Agency of Czech Technical University (project no. SGS24/140/OHK3/3T/13).

References

- Boddy, M. S.; Gohde, J.; Haigh, T.; and Harp, S. A. 2005. Course of Action Generation for Cyber Security Using Classical Planning. In *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling, June 5-10 2005, Monterey, California, USA*, 12–21. AAAI.
- Bylander, T. 1994. The Computational Complexity of Propositional STRIPS Planning. *Artificial Intelligence* 69(1-2):165–204. Number: 1-2.
- Bäckström, C., and Nebel, B. 1995. Complexity Results for SAS+ Planning. *Computational Intelligence* 11(4):625–655. Number: 4.
- Camacho, A.; Muise, C. J.; and McIlraith, S. A. 2016. From FOND to Robust Probabilistic Planning: Computing Compact Policies that Bypass Avoidable Deadends. In *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling, ICAPS 2016, London, UK, June 12-17, 2016*, 65–69. AAAI Press.
- Chrapa, L., and Karpas, E. 2024a. On Verifying and Generating Robust Plans for Planning Tasks with Exogenous Events. In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning, KR 2024, Hanoi, Vietnam. November 2-8, 2024*.
- Chrapa, L., and Karpas, E. 2024b. On Verifying Linear Execution Strategies in Planning Against Nature. In *Proceedings of the Thirty-Fourth International Conference on Automated Planning and Scheduling, ICAPS 2024, Banff, Alberta, Canada, June 1-6, 2024*, 86–94. AAAI Press.
- Chrapa, L.; Faber, W.; and Morak, M. 2021. Universal and Uniform Action Reversibility. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, 651–654.
- Chrapa, L.; Pilát, M.; and Med, J. 2021. On Eventual Applicability of Plans in Dynamic Environments with Cyclic Phenomena. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, 184–193.
- Cimatti, A.; Pistore, M.; Roveri, M.; and Traverso, P. 2003. Weak, Strong, and Strong Cyclic Planning via Symbolic Model Checking. *Artificial Intelligence* 147(1-2):35–84.
- Cserna, B.; Doyle, W. J.; Ramsdell, J. S.; and Ruml, W. 2018. Avoiding Dead Ends in Real-Time Heuristic Search. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 1306–1313. AAAI Press.
- Daum, J.; Torralba, Á.; Hoffmann, J.; Haslum, P.; and Weber, I. 2016. Practical Undoability Checking via Contingent Planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 26, 106–114. AAAI Press.
- Ghallab, M.; Nau, D.; and Traverso, P. 2016. *Automated Planning and Acting*. New York, NY: Cambridge University Press.
- Littman, M. L.; Goldsmith, J.; and Mundhenk, M. 1998. The Computational Complexity of Probabilistic Planning. *Journal of Artificial Intelligence Research* 9:1–36.
- Med, J., and Chrapa, L. 2022. On Speeding Up Methods for Identifying Redundant Actions in Plans. In *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling, ICAPS 2022, Singapore (virtual), June 13-24, 2022*, 252–260. AAAI Press.
- Med, J.; Chrapa, L.; Morak, M.; and Faber, W. 2024. Weak and Strong Reversibility of Non-deterministic Actions: Universality and Uniformity. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, 369–377. AAAI Press.
- Morak, M.; Chrapa, L.; Faber, W.; and Fiser, D. 2020. On the Reversibility of Actions in Planning. In *Proceedings of the Seventeenth International Conference on Principles of Knowledge Representation and Reasoning*, 652–661. Rhodes, Greece: International Joint Conferences on Artificial Intelligence Organization.
- Muise, C. J.; McIlraith, S. A.; and Belle, V. 2014. Non-Deterministic Planning with Conditional Effects. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, Portsmouth, New Hampshire, USA, June 21-26, 2014*. AAAI.
- Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley.
- Rintanen, J. 2004. Complexity of Planning with Partial Observability. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004), June 3-7 2004, Whistler, British Columbia, Canada*, 345–354. AAAI.
- Savitch, W. J. 1970. Relationships between Nondeterministic and Deterministic Tape Complexities. *Journal of Computer and System Sciences* 4(2):177–192.
- Weber, I.; Wada, H.; Fekete, A. D.; Liu, A.; and Bass, L. 2012. Automatic Undo for Cloud Management via AI Planning. In *Proceedings of the Eighth USENIX Conference on Hot Topics in System Dependability, HotDep’12*, 10. USA: USENIX Association. Hollywood, CA.
- Weber, I.; Wada, H.; Fekete, A. D.; Liu, A.; and Bass, L. 2013. Supporting Undoability in Systems Operations. In *27th Large Installation System Administration Conference (LISA 13)*, 75–88. Washington, D.C.: USENIX Association.
- Yoon, S. W.; Fern, A.; and Givan, R. 2007. FF-Replan: A Baseline for Probabilistic Planning. In *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling, ICAPS 2007, Providence, Rhode Island, USA, September 22-26, 2007*, 352. AAAI.