

REVIEW

Reformulation techniques for automated planning: a systematic review

Diaeddin Alarnaouti, George Baryannis , and Mauro Vallati 

School of Computing and Engineering, University of Huddersfield, Huddersfield HD1 3DH, UK

Corresponding author: Mauro Vallati; Email: m.vallati@hud.ac.uk

Received: 14 June 2022; **Revised:** 22 November 2022; **Accepted:** 11 October 2023

Abstract

Automated planning is a prominent area of Artificial Intelligence and an important component for intelligent autonomous agents. A cornerstone of domain-independent planning is the separation between planning logic, that is the automated reasoning side, and the knowledge model, that encodes a formal representation of domain knowledge needed to reason upon a given problem to synthesize a solution plan. Such a separation enables the use of reformulation techniques, which transform how a model is represented in order to improve the efficiency of plan generation. Over the past decades, significant research effort has been devoted to the design of reformulation techniques. In this paper, we present a systematic review of the large body of work on reformulation techniques for classical planning, aiming to provide a holistic view of the field and to foster future research in the area. As a tangible outcome, we provide a qualitative comparison of the existing classes of techniques, that can help researchers gain an overview of their strengths and weaknesses.

1. Introduction

Automated planning is a research discipline that addresses the problem of generating a totally or partially ordered sequence of actions that transform the environment from some initial state to a desired goal state. Within this discipline, domain-independent planning refers to those approaches that keep the knowledge model, the domain knowledge related to the problem at hand, separate from planning logic, that enables automated reasoning to generate plans. The development of domain-independent planners within the AI Planning community facilitates the use of this ‘off-the-shelf’ technology for a wide range of applications, including UAV maneuvering (Ramírez *et al.*, 2018), space exploration (Ai-Chang *et al.*, 2004), and train dispatching (Cardellini *et al.*, 2021). This is despite the complexity issues inherent in plan generation, which are exacerbated by the separation of planner logic from domain knowledge. On the other hand, this separation has the advantage that planning engines can be interchanged in a modular way, provided that they accept the same language for describing planning problems and deliver the same type of plans.

This modular approach has fostered the development of planning engines, as well as the design and exploitation of reformulation techniques. These refer to the ability to automatically re-formulate, re-represent or tune the domain model and/or a problem description, while keeping to the same input language, in order to increase the efficiency of a planning engine and expand the scope of problems solved (Riddle *et al.*, 2011). The aim is to make these techniques independent of domain and planner to some degree (i.e. applicable to a range of domains and planning engine technology), and use them to form a wrapper around a planner, improving its overall performance for the particular domain to which it is applied.

Cite this article: D. Alarnaouti, G. Baryannis and M. Vallati. Reformulation techniques for automated planning: a systematic review. *The Knowledge Engineering Review* 38(e9): 1–17. <https://doi.org/10.1017/S0269888923000097>

In the past few decades, a number of reformulation techniques have been introduced. This is particularly true for classical planning which, despite its intrinsic simplicity, provides an ideal ground to study and investigate planning techniques that can then be extended beyond a specific framework. For this reason, classical planning has been studied for several decades and is still a focal point of research within the automated planning community, as evidenced by the papers devoted to it in the flagship conference, the International Conference on Automated Planning and Scheduling (ICAPS)¹ and by the number of papers focusing on knowledge engineering for classical planning that are accepted at the Workshop on Knowledge Engineering for Planning and Scheduling (KEPS).

In this paper, we systematically review the state of the art of reformulation techniques for classical planning, with the aim of providing a holistic view of the field and answering the following research questions:

- What reformulation techniques have been proposed in literature?
- How can they be applied on an indicative planning problem?
- What are the particular strengths and weaknesses of each individual reformulation technique?

The main purpose of this review is to provide a qualitative comparison of the existing reformulation approaches, with the aim of helping experts and practitioners in identifying the most promising technique to be used, and to highlight research gaps that can foster further research in the area. Note that empirical analysis of the reviewed reformulation techniques is beyond the scope of this work due to the high variability, among different papers, of considered benchmarks, planning engines, and software/hardware infrastructure (Bocchese *et al.*, 2018).

The remainder of this paper is organized as follows. Section 2 the necessary background on planning and reformulation, and we introduce an example domain model to be used as a running example throughout the paper. Then, Section 3 describes the methodology used for this literature review. Section 4 is the main part of the review, presenting the existing reformulation techniques and showing how they can be implemented on the running example. We then compare the reviewed techniques in Section 5, while in Section 6 we briefly discuss some reformulation techniques that target nonclassical planning problems. Finally, we provide conclusions and suggested directions for future research in the area of reformulation techniques for automated planning.

2. Background

This section is devoted to providing the required background in terms of automated classical planning, the Gripper domain model used as a running example, and the notion of reformulation.

2.1. Classical planning

Classical planning is concerned with finding a (partially or totally ordered) sequence of actions transforming the static, deterministic, and fully observable environment from the given initial state to a desired goal state (Ghallab *et al.*, 2004).

In the classical representation, a *planning task* consists of a *planning domain model* and a *planning problem*, where the planning domain model describes the environment and defines planning operators while the planning problem defines concrete objects, an initial state and a set of goals. The environment is described by *predicates* that are specified via a unique identifier and terms (variable symbols or constants).

Formally, a *planning task* is a pair $\Pi = (Dom_{\Pi}, Prob_{\Pi})$ where a *planning domain model* $Dom_{\Pi} = (P_{\Pi}, Ops_{\Pi})$ is a pair consisting of a finite set of predicates P_{Π} and planning operators Ops_{Π} , and a

¹The interested reader is referred to the ICAPS website: <https://www.icaps-conference.org/>

planning problem $Prob_{\Pi} = (Objs_{\Pi}, I_{\Pi}, G_{\Pi})$ is a triple consisting of a finite set of objects $Objs_{\Pi}$, initial state I_{Π} , and goal G_{Π} .

Let ats_{Π} be the set of all *atoms* that are formed from the predicates P_{Π} by applying all possible substitution mappings from the predicate parameters (variable symbols) to the objects from $Objs_{\Pi}$. In other words, an atom is an *instance* of a predicate (in this article, when we use the term instance, we mean an instance that is fully *ground*). A *state* is a subset of ats_{Π} , and the *initial state* I_{Π} is a distinguished state. The *goal* G_{Π} is a nonempty subset of ats_{Π} , and a *goal state* is any state that contains the goal G_{Π} . Note that the semantics of *state* reflect the full observability of the environment; that is, for a state s , atoms present in s are assumed to be true in s , while atoms not present in s are assumed to be false in s .

Planning operators are ‘modifiers’ of the environment. They consist of *preconditions*, that is, what must hold prior to an operator’s application, and *effects*, that is, what is changed after its application. We distinguish between *negative effects*, that is, what becomes false, and *positive effects*, that is, what becomes true after an operator’s application. *Actions* are instances of planning operators, that is, an operator’s parameters, as well as corresponding variable symbols in its preconditions and effects, are substituted by objects (constants). Planning operators capture general types of activities that can be performed. While predicates can be instantiated to atoms to capture given relations between concrete objects, planning operators can be instantiated to actions to capture given activities between concrete objects.

A *planning operator* $o = (name(o), pre(o), eff(o))$ is specified such that $name(o) = op_name(x_1, \dots, x_k)$, where op_name is a unique identifier and x_1, \dots, x_k are all the variable symbols (parameters) appearing in the operator, $pre(o)$ is a set of predicates representing its precondition, and $eff(o)$ represents its effects, divided into $eff^{-}(o)$ and $eff^{+}(o)$ (i.e. $eff(o) = eff^{-}(o) \cup eff^{+}(o)$) that are sets of predicates representing the operator’s negative and positive effects, respectively. *Actions* are instances of planning operators that are formed by substituting objects, which are defined in a planning problem, for operators’ parameters as well as for the corresponding variable symbols in operators’ preconditions and effects. An action $a = (pre(a), eff^{-}(a) \cup eff^{+}(a))$ is *applicable* in a state s if and only if $pre(a) \subseteq s$. Application of a in s , if possible, results in a state $(s \setminus eff^{-}(a)) \cup eff^{+}(a)$.

A solution of a planning task is a sequence of actions transforming the environment from the given initial state to a goal state. A *plan* is a sequence of actions. A plan is a *solution* of a planning task Π , a *solution plan* of Π in other words, if and only if a consecutive application of the actions from the plan starting in the initial state of Π results in the goal state of Π .

The standardized language for describing classical planning tasks is PDDL (McDermott, 2000), that was introduced in 1998 by the organizers of the first International Planning Competition², building on top of STRIPS (Fikes & Nilsson, 1971) and the Action Description Language (ADL) (Pednault, 1987).

2.2. The Gripper domain

As a running example in this paper, we consider the well-known Gripper domain model, initially introduced in the first International Planning Competition in 1998 by Jana Kohler (McDermott, 2000). This domain was selected in this paper because of its simplicity and due to its suitability for being reformulated. In a nutshell, the Gripper domain consists of a robot that has two grippers, and is tasked to move a number of balls between two rooms. The domain model includes three operators:

Move: to move the robot between rooms.

Pick: to use a gripper to pick up a ball.

Drop: to drop a ball that the robot is holding in one of its grippers.

²<https://www.icaps-conference.org/competitions/>.

```

(:predicates (room ?r) (ball ?b) (gripper ?g) (at-robby ?r) (at ?b ?r) (free ?g) (carry ?o ?g))

(:action move
  :parameters (?from ?to)
  :precondition (and (room ?from) (room ?to) (at-robby ?from))
  :effect (and (at-robby ?to) (not (at-robby ?from))))

(:action pick
  :parameters (?obj ?room ?gripper)
  :precondition (and (ball ?obj) (room ?room) (gripper ?gripper) (at ?obj ?room)
    (at-robby ?room) (free ?gripper))
  :effect (and (carry ?obj ?gripper) (not (at ?obj ?room) (not (free ?gripper))))

(:action drop
  :parameters (?obj ?room ?gripper)
  :precondition (and (ball ?obj) (room ?room) (gripper ?gripper)
    (carry ?obj ?gripper) (at-robby ?room))
  :effect (and (at ?obj ?room) (free ?gripper) (not (carry ?obj ?gripper))))

```

Figure 1. *The Gripper PDDL domain model.*

```

(:objects rooma roomb ball4 ball3 ball2 ball1 left right)
(:init (room rooma) (room roomb)
  (ball ball4) (ball ball3) (ball ball2) (ball ball1)
  (at-robby rooma)
  (free left) (free right)
  (at ball4 rooma) (at ball3 rooma) (at ball2 rooma) (at ball1 rooma)
  (gripper left) (gripper right))
(:goal (and (at ball4 roomb) (at ball3 roomb) (at ball2 roomb) (at ball1 roomb)))

```

Figure 2. *An example Gripper planning problem that consists of 4 balls to be carried from rooma to roomb.*

Figure 1 shows the PDDL code of the domain model, including the relevant predicates and the three mentioned operators. Figure 2 shows an example Gripper planning problem, where the robot has two grippers and has to move four balls between the considered rooms. The initial state defines the initial position of balls and robot, and the initial status of the grippers. The goal section specifies the desired goal position of the 4 balls.

2.3. Reformulation in Domain-Independent planning

Taking a general perspective, reformulation is a term meaning a change to the way in which one thinks about a problem, and it has been demonstrated to be a common practice employed also by people to tackle challenging problems (Riddle *et al.*, 2013). A theoretical framework for describing reformulation schemes in automated planning has been presented by Chrupa *et al.* (2012). Early research in the area of reformulation in AI started in the 1960s (Amarel, 1968), and reformulation has rapidly become a major field of research. In the field of automated planning and, more general, of state-space search, reformulation is intended as a change of representation. Different representations of the same problem can result in different search spaces, and the use of reformulation techniques can allow to make a problem more amenable for a considered planning engine by providing a search space that is easier to be explored to find a goal state.

Focusing on automated planning, the domain-independent paradigm decouples reasoning from knowledge representation. This supports the use of reformulation techniques which can re-formulate, re-represent or tune the domain model and/or problem description, while keeping to the same input language, in order to increase the efficiency of a planning engine and expand the scope of problems solved. The idea is that these techniques are independent from application domain and planning engines

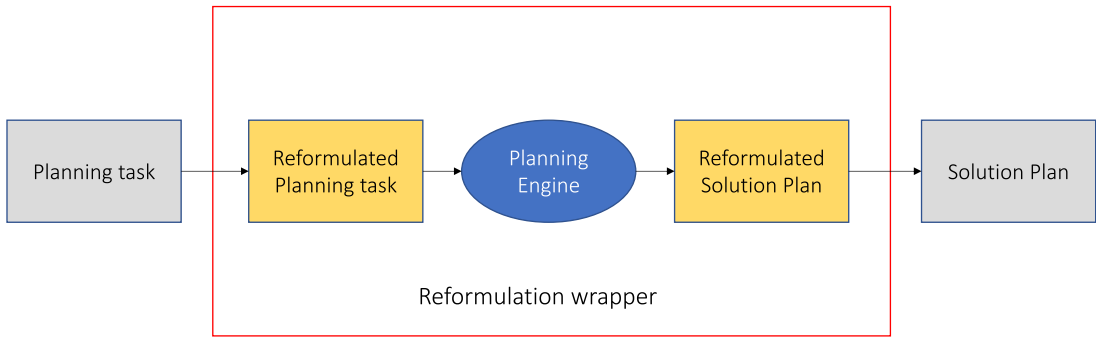


Figure 3. An overview of the use of a domain- and planner-independent reformulation technique.

(i.e. applicable to a range of domains and planning engine technology) and use them to form a wrapper around a planner, improving its overall performance for the domain to which it is applied.

Figure 3 gives an overview of the use of reformulation wrapper to modify the representation of a planning task in a domain- and planner-independent way that allows to reuse the same technique in any domain and with any engine that supports the considered formal language. In this review, we focus on classical planning reformulation approaches that aim at changing the representation of a planning task within the same formal language, with the goal of making the task more amenable for a given (class of) engines.

3. Review methodology

This section is devoted to presenting the methodology used for performing the literature review. To ensure deeper understanding of the reformulation approaches that are applicable in domain-independent classical planning, we applied a variation of the systematic literature review methodology as proposed by Kitchenham and Charters (2017) for software engineering and previously applied for reviews of AI-related literature (Baryannis *et al.*, 2019).

3.1. Search strategy

The collection of literature concerning reformulation techniques in classical planning was performed through automated search using Google Scholar, Scopus, and the University of Huddersfield library search engine (Summon). Two levels of keyword terms were utilized. On the first level, we used a disjunction (OR) of the following three keyword phrases: ‘classical planning reformulation’, ‘PDDL reformulation’, and ‘domain-independent reformulation’. On the basis of the techniques identified using the aforementioned keyword phrases, we provided a second level of keywords, to obtain all the relevant works for each of them. This second-level search was with the following keywords linked to one another by OR: ‘macro-operators’, ‘macro-actions’, ‘entanglements’, ‘actions elimination’, ‘bagged representation’, ‘action schema splitting’, and ‘model configuration’.

The previously mentioned libraries were used equally as results vary between them. The covered time period is set from 1980 to January 2022. Notably, thanks to the introduction of the standardized PDDL language and the first edition of the International Planning Competition in 1998, most of the relevant work has been published after 2000.

To maximize coverage and completeness of the conducted search, two ancillary procedures were included: (1) checking reference lists of select primary studies (often referred to as backward snowballing); and (2) identifying existing literature reviews on automated planning that may include reformulation (more details in Section 3.3).

3.2. Search scope

A series of inclusion and exclusion criteria frames the scope of this review, to ensure quality and relevance of the selected works. First, studies must be peer-reviewed and written in English. Then, they should contain a formal description of at least one reformulation technique for automated planning models, accompanied with an empirical analysis of the usefulness of the proposed technique. Third, the proposed reformulation technique must be domain-independent and must focus on classical planning. Fourth, the reformulation must accept planning models as input and provide as output models that can be processed by domain-independent classical planning engines. It should be noted that the last point excludes some well-known approaches such as DISCOPLAN (Gerevini & Schubert, 1998) and TIM (Fox & Long, 1998), that do not allow to include the extracted knowledge in planning models. While the main focus of the review is on techniques for reformulating classical PDDL lifted models, here we also include works proposed before the introduction of the language, and that propose techniques for reformulating STRIPS (Fikes & Nilsson, 1971) models; for the purposes of this review these can be considered to be similar in nature to PDDL ones.

A search based on the keywords and the engines described in Section 3.1 yields several thousands of results. By carefully and thoroughly applying all aforementioned criteria and ancillary search strategies, and by considering also abstracts of the identified studies, 54 studies remain and are analyzed in the rest of this paper. For the sake of readability, the selected studies are divided into different classes, presented in Section 4, according to the reformulation approach that they present.

3.3. Related surveys

The main motivation behind this systematic review is the limited amount of works covering the field of domain-independent reformulation for automated planning. The most relevant published work is a survey of machine learning methods for automated planning (Jiménez *et al.*, 2012) where a section is dedicated to macro-actions. This survey covers only 4 studies that are also covered in this paper, published between 1977 and 2007. The work of Long *et al.* (2002) provides a definition of reformulation for automated planning and briefly discusses some ways in which it has been applied in the field, though it also includes cases not within the scope of this review. To the best of our knowledge, this review is the first attempt at systematically analyzing reformulation approaches for automated planning.

4. Reformulation techniques for classical planning

In this section we present the reviewed studies, classified according to the implemented reformulation approach. For each approach, we also provide an example of applying the reformulation approach to the Gripper domain.

4.1. Entanglements

The entanglements approach was firstly introduced in Chrupa and Barták (2009) and aims at removing useless actions from a ground planning task, to reduce the branching factor of a task at hand. Such useless actions are identified by analyzing the relationship between predicates of the problem, and operators. In a nutshell, the idea is to change the representation of a planning task so that only potentially useful actions can be grounded and considered during search.

In literature, two main classes of entanglements have been proposed: outer and inner. Outer entanglements (Chrupa & McCluskey, 2012; Chrupa *et al.*, 2018) identify useless actions on the basis of their relationships with predicates from the initial state or the goal state. Inner entanglements (Chrupa *et al.*, 2019) focus instead on the relation between pairs of actions. There are two types of inner entanglements: *entanglements by preceding*, which defines the case where a certain predicate is required by an operator

```

(:predicates ... (at-ent ?obj ?room))
...
(:action pick
  :parameters (?obj ?room ?gripper)
  :precondition (and (ball ?obj) (room ?room) (gripper ?gripper) (at ?obj ?room)
    (at-ent ?obj ?room) (at-roby ?room) (free ?gripper))
  :effect (and (carry ?obj ?gripper) (not (at ?obj ?room) (not (free ?gripper))))
...))
....
(:init ... (at ball4 rooma) (at ball3 rooma) (at ball2 rooma) (at ball1 rooma) ...
(at-ent ball4 rooma) (at-ent ball3 rooma) (at-ent ball2 rooma) (at-ent ball1 rooma))
...

```

Figure 4. An excerpt of the modified Gripper planning task when the *pick* operator is entangled by *init* on the basis of the *at* predicate.

as a precondition, and *entanglements by succeeding*, which denotes the case where a particular operator makes a predicate available as one of its effects.

According to the results presented in Chrupa *et al.* (2012), exactly finding entanglements is as hard as solving the planning task itself. For this reason, most of the work in literature focuses on approximate approaches for identifying entanglement relations. To exemplify how entanglements can reformulate a given planning task, here we show how the Gripper domain can be modified by entanglements by *init* for the operator *pick*. The basis of entanglement in this case for the particular encoding is that it is useful to perform a *pick* action only on balls that are in the initial room. In any other occasion, the *pick* action should not be considered. The relationship is therefore captured by the $(at \ ?b \ ?r)$ predicate, that indicates the position of a ball and is a precondition of the considered operator. To implement the entanglement, a new predicate p' $(at-ent \ ?b \ ?r)$ should be created and added to the domain model, having the same arguments of $(at \ ?b \ ?r)$. The *pick* operator should be modified by adding p' as a precondition. In the initial state description of the problem, an additional p' predicate should be created for each existing $(at \ ?b \ ?r)$. This is similar in nature to the auxiliary predicate added to preconditions to account for unforeseen circumstances to address the qualification problem in service specifications (Baryannis & Plexousakis, 2013; Baryannis *et al.*, 2017) An excerpt of the modified domain and problem models is presented in Figure 4.

4.2. Macro-operators

Macro-operators (macros) represent a well-known and well-studied technique for enhancing performance of planning engines. In a nutshell, macros encapsulate sequences of (primitive) planning operators. Macros are encoded as ordinary planning operators and, hence, they can be added into planning domain models. Macros, informally speaking, provide shortcuts in the state space and, consequently, planning engines can generate plans in a smaller number of steps. This comes at the cost of an increased branching factor, since macros often have many more instances than primitive operators and thus their use might introduce additional overheads as well as larger memory requirements.

The notion of macros can be traced back to 1970s and 1980s. REFLECT (Dawson & Siklossy, 1977) builds macros from pairs of primitive operators that can be applied in sequence and share at least one argument. Macro Problem Solver (Korf, 1985) is capable of learning macros for achieving particular non-serializable subgoals (e.g. in Rubik's cube). MORRIS (Minton, 1985) learns macro-operators from parts of plans appearing frequently (S-macros) or being potentially useful despite having low priority (T-macros). FM (McCluskey, 1987) learns sequences of operators frequently used together and combines them in potentially long sequences called chunks. MACLEARN (Iba, 1985, 1989) is a heuristic approach for learning re-usable macros for solving puzzle problems; the system includes three subcomponents that are in charge of proposing promising macros and testing their usefulness for solving problems.

The first International Planning Competition, held in 1998, introduces PDDL that became the ‘de facto’ standard language for planning models. The introduction of a language supported the design and testing of approaches for the generation and identification of macros, that started to thrive. MacroFF (Botea *et al.*, 2005), based on the well-known FF planning engine (Hoffmann & Nebel, 2001), generates macros according to a number of predefined rules (e.g. the ‘locality rule’) that apply on adjacent actions in training plans. MacroFF is capable of generating planner-independent macros, that can be added to domain models as standard operators, or planner-specific macros for FF, that can be provided as additional input to the planning engine. DHG (Armano *et al.*, 2004) is able to learn macro-operators from static domain analysis by exploring a graph of dependencies between operators. WIZARD (Newton & Levine, 2007) is a framework that exploits genetic programming to create macros; starting from the primitive operators of the domain model, WIZARD leverages genetic algorithms to combine them into useful macros for a given planning engine. Chrpa (2010b) propose an approach for identifying suitable macros by looking at action dependencies in generated plans. Dulac *et al.* (2013) propose to exploit n-gram algorithm to analyze training plans to automatically learn macros.

DBMP/S (Hofmann *et al.*, 2017) applies Map Reduce for learning macros from a large training plan databases. More recently, the same authors propose an approach for generating macros for ADL domain models, that includes PDDL features that are rarely supported by more traditional methods (Hofmann *et al.*, 2020). CAP (Asai & Fukunaga, 2015) exploits component abstraction, that allows to cluster together similar objects (introduced by MacroFF), for generating sub-goal specific macros. In other words, CAP divides complex planning problems into independent subproblems by abstracting the components of the original problem. Then it finds subplans for each subproblem and connects the actions of every subplan into a solo macro operator. BloMa (Chrpa & Siddiqui, 2015) leverages block deordering, which rearranges plans into ‘blocks’ that can no longer be deordered (Siddiqui & Haslum, 2012), for generating longer macros. In particular, BloMa generates a large pool of macros from ‘macroblocks’, which are derived from ‘blocks’ by applying a set of rules. Chrpa & Vallati (2019) introduce the idea of critical section macros, that are inspired by parallel computing critical sections; such macros aim at capturing a whole activity that deals with a limited resource (or more limited resources). Finally, Castellanos-Paez *et al.* (2021a,b) introduce ERA, an approach for extracting macros from plans that is based on pattern mining; an important feature of ERA is its ability to identify macros even if the included operators are not always adjacent in the considered plans.

A different line of work looks into exploiting the notion of entanglements to identify promising macros. Chrpa (2010a) focuses on combining both macro-operators and entanglements in order to get the benefit behind each of them, as macros can reduce the size of the search space whereas the usage of entanglements is capable of reducing the branching factor which may occur because of the generated instances of macros. Chrpa *et al.* (2013) propose an automated approach that combine two primitive operators that are linked by inner entanglement relationships, and leverages on such relationships to eliminate one or both of the primitive operators from the domain model. MUM (Chrpa *et al.*, 2014) is a learning system that exploits outer entanglements as heuristics in the process of generating macros. Macros generated by MUM have a limited number of instances, specifically, the number of macro instances has to be in the same order of magnitude as the number of primitive operator instances. OMA (Chrpa *et al.*, 2015b) is capable of generating macros online, that is without the need for offline training, by considering entanglement relations between operators of the domain model. Notably, there is also a line of work that considers the problem of identifying the best (set of) macros to be used, given an initial pool of candidates. Alhossaini & Beck (2013) select problem-specific macros from a given pool of macros (hand-coded or generated by another technique) using a specifically trained predictor. ASAP (Vallati *et al.*, 2013) uses a set of provided training plans to identify the best combination of planning engine and macro set (also considering entanglements) to be used on a given domain. PbP (Gerevini *et al.*, 2014) uses statistical tests to identify the most promising portfolio of planning engines and macro-actions to be used for solving challenging planning instance. Finally, MeVo (Vallati *et al.*, 2020), given a large pool of macros, can evolve over time the best set of macros to be used by a planning engine for solving a continuous stream of problems from a considered domain. This features


```

(:action pick-move-drop
 :parameters (?from ?to ?obj ?gripper)
 :precondition (and (ball ?obj) (gripper ?gripper)
                   (at ?obj ?from) (at-robby ?from) (free ?gripper)
                   (room ?from) (room ?to))
 :effect (and
          (not (at ?obj ?from))
          (not (at-robby ?from))
          (at ?obj ?to)
          (at-robby ?to)))

```

Figure 5. An example macro-operator that encapsulates the sequence of operators *pick(?obj ?from ?gripper)*, *move(?from ?to)*, *drop(?obj ?to ?gripper)*. For the sake of clarity, we use the same name for matched variables between operators.

allows MeVo to overcome the issue of having training instances that are not representative of the testing ones.

Technically speaking, the way in which a macro-operator is generated by assembling two operators is straightforward and quite similar to the way composition works in the case of services (Baryannis & Plexousakis, 2014). Considering two operators a_i and a_j , the resulting macro encapsulating their execution in sequence can be generated as follows:

$$\begin{aligned}
 pre(a_{i,j}) &= Pre(a_i) \cup (Pre(a_j) \setminus eff^+(a_i)) \\
 eff^-(a_{i,j}) &= (eff^-(a_i) \cup eff^-(a_j)) \setminus eff^+(a_j) \\
 eff^+(a_{i,j}) &= (eff^+(a_i) \cup eff^+(a_j)) \setminus eff^-(a_j)
 \end{aligned}$$

More than two operators can be encapsulated in a single macro by iteratively repeating the described process. Figure 5 shows the macro-operator *pick-move-drop* for the Gripper domain model. It has been generated by composing the primitive operators *pick*, *move*, and *drop*, with the idea of providing a single operator that can represent a whole movement of a ball from its initial position to its goal position. It can be added to the original domain model, and planning engines will take it into account when solving a given planning task. Notably, the resulting plans may include the macro operator and therefore, to be valid with regard to the original domain model, will need to be parsed.

4.3. Operators elimination

Operators elimination aims at reducing the branching factor by removing from the domain model those operators whose effects can be achieved by executing a sequence of other operators. In a nutshell, the main point is to minimize the size of the model by identifying operators that can be considered redundant. It should of course be noted that it may not always be possible to remove operators from a domain model. Haslum and Jonsson (2000) introduce a technique for performing operators elimination, by formally defining the notions of a redundant operator and of a minimal set of operators, also proposing a greedy algorithm to identify the minimal set of operators for a given domain model.

Operators elimination may be of limited impact when used on its own, particularly in the case of highly engineered domain models. However, it can be particularly helpful when performed after another reformulation approach. For instance, Chrpa and Barták (2009) discuss operators elimination after the exploitation of entanglements, and there has been a large body of work investigating the elimination of primitive operators that are encapsulated into macros (e.g. Chrpa, 2010b; Alhossaini & Beck, 2013; Chrpa & Vallati, 2019).

Considering the Gripper domain, in the original model shown in Figure 1 there are no operators that are redundant and can be removed. However, if the domain model is extended by adding the macro shown in Figure 5, then both *pick* and *drop* operators can be removed without compromising solvability of

```

(:predicates ... (count ?b ?r ?n) (more ?n1 ?n2))

(:action pick
 :parameters (?n1 ?n0 ?obj ?room ?gripper)
 :precondition (and (ball ?obj)(room ?room)
                   (gripper ?gripper)(at-robby ?room)(free ?gripper)
                   (more ?n1 ?n0)(count ?obj ?room ?n0))
 :effect (and (carry ?obj ?gripper)
              (not (count ?obj ?room ?n0))
              (count ?obj ?room ?n1)(not (free ?gripper))))
...)
....

(:objects n4 n3 n2 n1 n0 rooma roomb ballX left right)
(:init .... (ball ballX) (more n0 n1) (more n1 n2)(count ballX rooma n4)
(count ballX roomb n0) ...)
(:goal (count ballX roomb n4))
)

```

Figure 6. An excerpt of the Gripper domain and problem models reformulated using bagged representation for objects of the type *ball*.

problems sharing the structure of the one shown in Figure 2—that is where the goal is not for the robot to hold a ball in its gripper.

4.4. Bagged representation

In domain models encoded in PDDL, it is usually the case that each object is uniquely identified, even if it is not important to distinguish between objects. In the presence of large sets of objects, this can lead to an explosion of the combinatorial problem, that needs to take into account the specific information of each individual object. The bagged representation reformulation addresses the above criticism: in cases where only the number of objects is relevant, and it is not important to have the ability to distinguish between objects, they can be represented as ‘bags’ of identical objects. The main advantage is to reduce the branching factor, by basically pruning states that are identical but for the specific object.

Bagged representation was first introduced in 2013 (Riddle *et al.*, 2013) and then further extended by providing an in-depth analysis of its impact on well-known benchmarks and automated techniques to perform the reformulation (Riddle *et al.*, 2015b,c,a, 2016).

When it comes to reformulating the Gripper domain model, objects of type *ball* are an excellent candidate to be represented using bagged representation. It is important to know their number in every room, but it is not important to know which specific ball is where. Figure 6 presents an excerpt of the reformulated domain and problem models. Objects of type *ball* are removed and are substituted by counters. The predicate *count* is used to record the number of balls in a room. Predicate *more* is used to link together the different possible values of a counter. This is required because in classical planning there is no notion of numeric elements, so this allows to use Boolean predicates such as *more* as counters. The reformulated operator *pick* is also shown in Figure 6, with the main change being that it now updates the counter of the room where it is applied. The *drop* operator (not shown in the figure) is modified in a similar way. Finally, the problem model is reformulated by removing the *ball* objects, appropriately setting the initial values of the counters for the considered rooms, and expressing the goal in terms of the number of balls that need to be in the final room.

4.5. Action schema splitting

The presence of operators with a large number of parameters can be problematic for the grounding step of domain-independent planning engines, as they usually have to instantiate all possible actions before

```

(:action drop
  :parameters (?obj ?room ?gripper)
  :precondition (and (ball ?obj) (room ?room) (gripper ?gripper)
    (carry ?obj ?gripper) (at-roby ?room))
  :effect (and (at ?obj ?room) (free ?gripper) (not (carry ?obj ?gripper))))

(:action pick-move
  :parameters (?obj ?from ?to ?gripper)
  :precondition (and (ball ?obj) (room ?from) (room ?to) (gripper ?gripper)
    (at ?obj ?from) (at-roby ?from) (free ?gripper))
  :effect (and (carry ?obj ?gripper) (at-roby ?to) (not (at ?obj ?from))
    (not (at-roby ?from)) (not (free ?gripper))))

```

Figure 7. Action Schema Splitting applied to the macro operator shown in Figure 5, and resulting in two operators.

filtering out those that are irrelevant. The number of ground actions grows exponentially with the number of parameters and the number of objects of the problem to solve. In order to address this issue, the idea of action schema splitting (Areces *et al.*, 2014) is to split large operators into smaller ones (with regard to the number of parameters). The main aim is to break the exponential growth by dividing parameters between different operators. While the idea is intuitively easy, its exploitation has a number of hidden issues to be taken into account. When an operator is broken into two smaller operators, the order in which these two operators are executed becomes important, as is to consider whether their execution can be interleaved with different operators—that can change the state of the world in some unexpected ways.

Considering the example Gripper domain model, in the original version of the model there is no operator that is suitable to be split using the described methodology. However, action schema splitting can be understood also as the ‘opposite’ of macro-operators, and in fact, it can also be used to find reformulations of models by first encapsulating primitive operators into macros and then splitting them in different ways. For this reason, to exemplify the use of this technique, we consider the macro encapsulating pick, move, drop operators shown in Figure 5 and we split it into two operators, pick-move and drop. Figure 7 shows the resulting operators. It is worth highlighting that the example also shows a drawback of the action schema splitting technique, that is the fact that the resulting operators may have the same number of parameters as the initial large operator.

4.6. Domain model configuration

It is well known that the way in which elements of planning models are ordered can have an impact on the performance of domain-independent planning engines (Howe & Dahlgren, 2002). In this context, the term elements can refer to operators, pre- and post-conditions, predicate definitions, etc. for the domain models; objects, initial, and goal state predicates listing for problem definitions.

The idea behind domain model configuration is to identify an ordering of elements that can improve the performance of a considered domain-independent planning engine. Vallati *et al.* (2015) introduce an approach that leverages on algorithm configuration techniques to identify a suitable configuration of a domain model to improve the performance of a planning engine. The work has been subsequently extended (Vallati *et al.*, 2021) to consider also cases where macro-operators have to be added to the domain model. Vallati *et al.* (2017) describe a method for the online reordering of domain models by means of dedicated heuristics, based on aspects such as the number of preconditions, number of effects, etc. In a different line of work, Vallati and Serina (2018) explore the configuration of planning problem models, by considering a structure called Planning Encoding Graph (PEG) (Serina, 2010) to produce information that helps in creating the basis on which the reordering of elements should be done. In a nutshell, the PEG can provide information about how important some objects of the

Table 1. Qualitative comparison of the main strengths and weaknesses of the reviewed reformulation techniques.

Reformulation approach	Benefits	Drawbacks
Macro-Operators	- Reduce depth (reduce the transitions needed to reach goal states)	- Increase branching factor - Increase of the ground size
Entanglements	- Reduce the branching factor	- Potentially incomplete
Actions Elimination	- Reduce the branching factor	- Rarely applicable - Potentially incomplete
Bagged Representation	- Reduce the branching factor	- Only applicable when having indistinguishable numerable objects
Action Schema Splitting	- Reduce the ground size of the problem	- Potential increase of the branching factor
Domain Model Configuration	- Ease the exploration of the search space	- Can have limited impact if planners re-order elements internally - Potential reduction of performance if wrong ordering is used

problem are based on their involvement in predicates of both initial and goal descriptions. This knowledge can then be exploited in the ordering of the predicates, according to the objects that they deal with.

Considering the guidelines in Vallati *et al.* (2015), there are no unpromising operators that we may prefer to put last in the Gripper domain model. However, following the introduced notion of *directionality*, we may change the ordering of operators to pick, move, drop. In this way, the ordering of operators follow the expected typical ordering of corresponding actions in solution plans that has been demonstrated to be useful for a range of planning engines.

5. Qualitative comparison

Having completed an overview of the existing literature on domain- and planner-independent reformulation for classical planning, we are now able to qualitatively compare the considered techniques. In fact, we will now focus on the advantages and disadvantages of the reviewed techniques, with the aim of providing useful guidelines for planning experts and practitioners in the process of selecting a promising technique to be used to improve planning performance on a domain of interest. Table 1 gives an overview of the reviewed reformulation techniques in terms of their main advantages and the major potential drawbacks.

As indicated in Table 1, most reformulation techniques aim at making the exploration of the search space easier by reducing the branching factor, reducing the number of steps needed to reach a goal state, or reducing the ground size of the problem to be solved. However, every technique has potential drawbacks to be weighed in when deciding whether to exploit it. Drawbacks range from the serious potential loss of completeness to a more shallow limited impact on performance in the worst case. The use of macros can boost the performance of a planning engine by reducing the number of steps needed to reach a goal state, at the cost of a potentially significant increase in terms of the branching factor and a larger ground problem. Entanglements aim at reducing the branching factor and the ground size of a problem by eliminating unpromising actions, but in doing so can remove some or all of the paths to goal states. In a similar fashion, action elimination directly reduces the ground size by eliminating operators

deemed to be useless, but on its own it is rarely applicable to well-formed domain models. Further, if applied in an ‘aggressive’ way, this may undermine completeness. Bagged representation provides an elegant way to reduce branching factor by removing the differences between objects, but it is again a technique that is rarely applicable on its own. Action schema splitting aims at reducing the ground size of a problem by splitting complex operators, but this comes at the cost of increasing the branching factor. Finally, configuring a domain model can help improve the performance of planning engines by listing the most important planning elements first, but if used inappropriately it can also lead to adverse effects.

It should be noted, however, that in many cases it is possible to combine different techniques to maximize performance boost and mitigate drawbacks of individual techniques (Vallati & Kitchin, 2020). As mentioned in previous sections, macros are frequently combined with other reformulation techniques. For instance, they have been used with entanglements (Chrupa *et al.*, 2014), with action elimination (Chrupa & Vallati, 2022), or with domain model configuration (Vallati *et al.*, 2021). While domain model configuration can be straightforwardly combined with any other reformulation technique, the combination of macros and entanglements can lead to significant performance improvements, as entanglements are tackling the main drawback that is associated with the use of macros, that is the branching factor.

The main take-home messages that can be diluted by the performed systematic review of the literature, can be summarized as follows.

- The majority of existing literature is dedicated to techniques for generating macros, while comparatively limited effort has been devoted to investigating alternative or different reformulation techniques. Admittedly, the planning community should aim at expanding the spectrum of reformulation techniques as much as possible, to foster their combination and the possibility to use them fruitfully in large and challenging domain models—where they are needed the most.
- A substantial number of existing reformulation techniques aims at addressing issues that are typical of planning engines that reason on the basis of a ground representation. This focus is historically motivated, as domain-independent planning engines traditionally ground the lifted PDDL representation. However, as engines capable of reasoning with lifted or partially ground representation are gaining momentum (see for instance Corrêa *et al.*, 2020; Horčík & Fišer, 2021), reformulations that are effective also on lifted models can and should be investigated more convincingly.
- On a similar note as the previous point, there is a lack of reformulations that look into exploiting the potential synergies with planning approaches that are based on compiling the planning problem into an equivalent SAT or ASP problem to be solved. Given the fact that reformulations for such languages have been proposed (see for instance Dodaro *et al.*, 2022), there may be changes in PDDL models that can result in performance boost in the resulting compiled instances.
- A potentially interesting area to explore is the reformulation of problem models. Existing techniques tend to have a strong focus on the domain model, to be widely applicable on problems from the same domain. However, as planning is more and more used in real-world applications, the complexity of problems is increasing as well. Further, in some domains like urban traffic control (McCluskey & Vallati, 2017) or energy network balancing (Piacentini *et al.*, 2013), large parts of the problem models remain the same between problems (i.e. the network description); so it would be worth exploring if problem models can benefit from a targeted reformulation.
- We recognize that further work is needed to improve the usability of planning techniques in real-world applications, as the main focus of existing reformulation techniques is still on classical planning. However, as demonstrated by a number of recent works, in many cases classical planning reformulation approaches can be extended to, or provide inspiration for, nonclassical systems.

6. Beyond classical planning

In this review, we focused on reformulation techniques for classical planning. However, it is worth mentioning that there is also a (limited) body of work looking at reformulation for problems beyond classical planning. In this section we provide an overview, by no means complete, of some of the reformulation approaches introduced for planning problems beyond classical. On the one hand, nonclassical planning models include additional language features that can complicate the reformulation process. Examples of languages that have the expressive power to represent nonclassical problems include PDDL 2 (Fox & Long, 2003) and PDDL 3 (Gerevini *et al.*, 2009), for numeric, temporal planning and for encoding preferences; PDDL + (Fox & Long, 2006) for mixed discrete-continuous problems, PPDDL (Younes & Littman, 2004) and RDDDL (Sanner, 2010) for probabilistic planning, and MA-PDDL (Kovács, 2012) for multi-agent problems. On the other hand, it can be argued that nonclassical planning is needed in most real-world applications, and it is therefore imperative to investigate techniques and approaches to boost planning performance in such circumstances.

A line of research that looks into extending reformulation techniques for classical planning to more expressive cases is that of Chrpá *et al.* (2015a), which investigates the use of entanglements in numeric planning problems. Specifically, the authors extend the notion of outer entanglements to allow them to handle numeric variables. Similarly, Scala (2014) extended macros to be used in the presence of numeric fluents. Finally, Franco *et al.* (2019) present a technique to reduce the ground size of PDDL+ planning problems by reducing the arity of sparse predicates, drawing a parallel to bagged representation for classical planning.

A different line of research focuses instead on translating a planning model from an original input language, to a different less-expressive one. The main advantage of this approach is increasing the number of planning engines that are able to reason upon the planning problem and leverage existing robust technologies devised for solving more restricted cases. Examples of this class of reformulation approaches include (Percassi *et al.*, 2021), that translates PDDL+ problems into PDDL2.1 ones, (Taig & Brafman, 2013; Grastien & Scala, 2020) that provides approaches for translating of conformant planning problems into classical problems, the re-representation of uncertainty in conformant planning problems (Palacios & Geffner, 2009), the translation of complex temporal aspects in PDDL2.1 (Cooper *et al.*, 2010), and the removal from PDDL3 of soft trajectory constraints (Percassi & Gerevini, 2019).

7. Conclusion

Reformulation represents a well-known class of approaches for improving the performance of domain-independent planning systems. The main idea is to re-represent a given planning problem in a way that allows to increase the efficiency of a selected planning engine to be used to solve the considered problem or class of problems. In this paper, we reviewed the state of the art of reformulation approaches for classical planning. We first presented in detail the different techniques and then took the opportunity to provide a qualitative comparison of their expected benefits and drawbacks, with the aim of providing some useful guidelines to select the most appropriate reformulation to be used for a considered problem. Notably, the provided comparison helps also in highlighting potentially fruitful combinations of reformulation techniques. Finally, to emphasize the importance of reformulation for planning models beyond classical, we briefly presented well-known work in the area.

Acknowledgement. Mauro Vallati is supported by the UKRI Future Leaders Fellowship [grant number MR/T041196/1].

References

- Ai-Chang, M., Bresina, J., Charest, L., Chase, A., Hsu, J.-J., Jonsson, A., Kanefsky, B., Morris, P., Rajan, K., Yglesias, J., *et al.* 2004. Mapgen: mixed-initiative planning and scheduling for the mars exploration rover mission. *IEEE Intelligent Systems* **19**(1), 8–12.

- Alhossaini, M. & Beck, J. C. 2013. Instance-specific remodelling of planning domains by adding macros and removing operators. In *Tenth Symposium of Abstraction, Reformulation, and Approximation*.
- Amarel, S. 1968. On representations of problems of reasoning about actions. Machine Intelligence.
- Areces, C. E., Bustos, F., Dominguez, M. & Hoffmann, J. 2014. Optimizing planning domains by automatic action schema splitting. In *Twenty-Fourth International Conference on Automated Planning and Scheduling*.
- Armano, G., Cherchi, G. & Vargiu, E. 2004. Automatic generation of macro-operators from static domain analysis. In *Proceedings of the 16th European Conference on Artificial Intelligence*, 955–956.
- Asai, M. & Fukunaga, A. 2015. Solving large-scale planning problems by decomposition and macro generation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 25, 16–24.
- Baryannis, G., Kritikos, K. & Plexousakis, D. 2017. A specification-based QoS-aware design framework for service-based applications. *Service Oriented Computing and Applications* 11(3), 301–314. ISSN 1863-2394. doi: [10.1007/s11761-017-0210-4](https://doi.org/10.1007/s11761-017-0210-4).
- Baryannis, G. & Plexousakis, D. 2013. WSSL: a fluent calculus-based language for web service specifications. In *25th International Conference on Advanced Information Systems Engineering (CAiSE 2013)*, Salinesi, C., Norrie, M. C. & Pastor, Ó. (eds), Lecture Notes in Computer Science 7908, 256–271. Springer Berlin Heidelberg. ISBN 978-3-642-38708-1. doi: [10.1007/978-3-642-38709-8_17](https://doi.org/10.1007/978-3-642-38709-8_17).
- Baryannis, G. & Plexousakis, D. 2014. Fluent calculus-based semantic web service composition and verification using WSSL. In *9th International Workshop on Semantic Web Enabled Software Engineering (SWESE2013), co-located with ICSSOC 2013*, Lomuscio, A., et al. (eds), Lecture Notes in Computer Science 8377, 256–270. Springer International Publishing Switzerland. doi: [10.1007/978-3-319-06859-6_23](https://doi.org/10.1007/978-3-319-06859-6_23).
- Baryannis, G., Validi, S., Dani, S. & Antoniou, G. 2019. Supply chain risk management and artificial intelligence: state of the art and future research directions. *International Journal of Production Research* 57(7), 2179–2202. doi: [10.1080/00207543.2018.1530476](https://doi.org/10.1080/00207543.2018.1530476).
- Bocchese, A. F., Fawcett, C., Vallati, M., Gerevini, A. E. & Hoos, H. H. 2018. Performance robustness of AI planners in the 2014 international planning competition. *AI Community* 31(6), 445–463.
- Botea, A., Enzenberger, M., Müller, M. & Schaeffer, J. 2005. Macro-ff: improving ai planning with automatically learned macro-operators. *Journal of Artificial Intelligence Research* 24, 581–621.
- Cardellini, M., Maratea, M., Vallati, M., Boletto, G. & Oneto, L. 2021. In-station train dispatching: a PDDL+ planning approach. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 450–458.
- Castellanos-Paez, S., Rombourg, R. & Lalanda, P. 2021a. ERA: Extracting planning macro-operators from adjacent and non-adjacent sequences. In *Pacific Rim Knowledge Acquisition Workshop*, 30–45. Springer.
- Castellanos-Paez, S., Rombourg, R. & Lalanda, P. 2021b. On the relevance of extracting macro-operators with non-adjacent actions: does it matter? In *13th International Conference on Agents and Artificial Intelligence*.
- Chrpa, L. 2010a. Combining learning techniques for classical planning: Macro-operators and entanglements. In *2010 22nd IEEE International Conference on Tools with Artificial Intelligence*, 2, 79–86. IEEE.
- Chrpa, L. 2010b. Generation of macro-operators via investigation of action dependencies in plans. *The Knowledge Engineering Review* 25(3), 281–297.
- Chrpa, L. and Barták, R. 2009. Reformulating planning problems by eliminating unpromising actions. In *Eighth Symposium on Abstraction, Reformulation, and Approximation*.
- Chrpa, L. and McCluskey, T. L. 2012. On exploiting structures of classical planning problems: Generalizing entanglements. In *ECAI*, 240–245.
- Chrpa, L., McCluskey, T. L. & Osborne, H. 2012. Reformulating planning problems: a theoretical point of view. In *25th International Florida Artificial Intelligence Research Society Conference*, 14–19.
- Chrpa, L., Scala, E. & Vallati, M. 2015a. Towards a reformulation based approach for efficient numeric planning: numeric outer entanglements. In *Eighth Annual Symposium on Combinatorial Search*.
- Chrpa, L. and Siddiqui, F. H. 2015. Exploiting block reordering for improving planners efficiency. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Chrpa, L. and Vallati, M. 2019. Improving domain-independent planning via critical section macro-operators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7546–7553.
- Chrpa, L. and Vallati, M. 2022. Planning with critical section macros: theory and practice. *Journal of Artificial Intelligence Research*.
- Chrpa, L., Vallati, M., McCluskey, T. L. and Kitchin, D. 2013. Generating macro-operators by exploiting inner entanglements. In *Tenth Symposium of Abstraction, Reformulation, and Approximation*.
- Chrpa, L., Vallati, M. & McCluskey, T. L. 2014. Mum: a technique for maximising the utility of macro-operators by constrained generation and use. In *Twenty-Fourth International Conference on Automated Planning and Scheduling*.
- Chrpa, L., Vallati, M. & McCluskey, T. L. 2015b. On the online generation of effective macro-operators. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Chrpa, L., Vallati, M. & McCluskey, T. L. 2018. Outer entanglements: a general heuristic technique for improving the efficiency of planning algorithms. *Journal of Experimental & Theoretical Artificial Intelligence* 30(6), 831–856.
- Chrpa, L., Vallati, M. & McCluskey, T. L. 2019. Inner entanglements: narrowing the search in classical planning by problem reformulation. *Computational Intelligence* 35(2), 395–429.
- Cooper, M. C., Maris, F. & Régnier, P. 2010. Compilation of a high-level temporal planning language into PDDL 2.1. In *2010 22nd IEEE International Conference on Tools with Artificial Intelligence*, 181–188.
- Corrêa, A. B., Pommerening, F., Helmert, M. & Frances, G. 2020. Lifted successor generation using query optimization techniques. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 30, 80–89.

- Dawson, C. & Siklossy, L. 1977. The role of preprocessing in problem solving systems: “an ounce of reflection is worth a pound of backtracking”. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence-Volume 1*, 465–471.
- Dodaro, C., Maratea, M. & Vallati, M. 2022. On the configuration of more and less expressive logic programs. In *Theory and Practice of Logic Programming*, 1–29. doi: [10.1017/S1471068422000096](https://doi.org/10.1017/S1471068422000096).
- Dulac, A., Pellier, D., Fiorino, H. & Janiszek, D. 2013. Learning useful macro-actions for planning with n-grams. In *2013 IEEE 25th International Conference on Tools with Artificial Intelligence*, 803–810.
- Fikes, R. E. & Nilsson, N. J. 1971. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* **2**(3–4), 189–208.
- Fox, M. & Long, D. 1998. The automatic inference of state invariants in TIM. *Journal of Artificial Intelligence Research* **9**, 367–421.
- Fox, M. & Long, D. 2003. PDDL2.1: an extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research* **20**, 61–124.
- Fox, M. & Long, D. 2006. Modelling mixed discrete-continuous domains for planning. *Journal of Artificial Intelligence Research* **27**, 235–297.
- Franco, S., Vallati, M., Lindsay, A. & McCluskey, T. L. 2019. Improving planning performance in PDDL+ domains via automated predicate reformulation. In *Computational Science - ICCS* 491–498.
- Gerevini, A. E., Haslum, P., Long, D., Saetti, A. & Dimopoulos, Y. 2009. Deterministic planning in the fifth international planning competition: PDDL3 and experimental evaluation of the planners. *Artificial Intelligence* **173**(5–6), 619–668.
- Gerevini, A., Saetti, A. & Vallati, M. 2014. Planning through automatic portfolio configuration: the pbp approach. *Journal of Artificial Intelligence Research* **50**, 639–696.
- Gerevini, A. & Schubert, L. 1998. Inferring state constraints for domain-independent planning. In *Proceedings of the American Artificial Intelligence Conference, AAAI/IAAI*, 905–912.
- Ghallab, M., Nau, D. & Traverso, P. 2004. *Automated Planning: Theory and Practice*. Elsevier.
- Grastien, A. & Scala, E. 2020. CPCES: a planning framework to solve conformant planning problems through a counterexample guided refinement. *Artificial Intelligence* **284**, 103271.
- Haslum, P. & Jonsson, P. 2000. Planning with reduced operator sets. In *AIPS*, 150–158.
- Hoffmann, J. & Nebel, B. 2001. The ff planning system: fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* **14**, 253–302.
- Hofmann, T., Niemueller, T. & Lakemeyer, G. 2017. Initial results on generating macro actions from a plan database for planning on autonomous mobile robots. In *Twenty-Seventh International Conference on Automated Planning and Scheduling*.
- Hofmann, T., Niemueller, T. & Lakemeyer, G. 2020. Macro operator synthesis for adl domains. In *ECAI 2020*, 761–768. IOS Press.
- Horčák, R. and Fišer, D. 2021. Endomorphisms of lifted planning problems. In *Proceedings of the International Conference on Automated Planning and Scheduling*, **31**, 174–183.
- Howe, A. E. and Dahlman, E. 2002. A critical assessment of benchmark comparison in planning. *Journal of Artificial Intelligence Research* **17**, 1–33.
- Iba, G. A. 1985. Learning by discovering macros in puzzle solving. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence - Volume 1*, 640–642.
- Iba, G. A. 1989. A heuristic approach to the discovery of macro-operators. *Machine Learning* **3**(4), 285–317.
- Jiménez, S., De La Rosa, T., Fernández, S., Fernández, F. & Borrajo, D. 2012. A review of machine learning for automated planning. *The Knowledge Engineering Review* **27**(4), 433–467.
- Kitchenham, B. & Charters, S. 2017. Guidelines for performing systematic literature reviews in software engineering. Technical Report EBSE-2007-01, EBSE Technical Report.
- Korf, R. E. 1985. Macro-operators: a weak method for learning. *Artificial Intelligence* **26**(1), 35–77.
- Kovács, D. L. 2012. A multi-agent extension of PDDL3.1. In *Proceedings of the 3rd Workshop on the International Planning Competition (IPC), ICAPS-2012*, Atibaia, Sao Paulo, Brazil, 19–27.
- Long, D., Fox, M. & Hamdi, M. 2002. Reformulation in planning. In *International Symposium on Abstraction, Reformulation, and Approximation*, 18–32.
- McCluskey, T. L. 1987. Combining weak learning heuristics in general problem solvers. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI’87*, San Francisco, CA, USA, 331–333. Morgan Kaufmann Publishers Inc.
- McCluskey, T. L. & Vallati, M. 2017. Embedding automated planning within urban traffic management operations. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS*, 391–399. AAAI Press.
- McDermott, D. M. 2000. The 1998 AI planning systems competition. *AI Magazine* **21**(2), 35–35.
- Minton, S. 1985. Selectively generalizing plans for problem-solving. In *IJCAI*, 596–599. Citeseer.
- Newton, M. H. & Levine, J. 2007. Wizard: suggesting macro-actions comprehensively. In *Proceedings of the Doctoral Consortium held at ICAPS*, 7.
- Palacios, H. & Geffner, H. 2009. Compiling uncertainty away in conformant planning problems with bounded width. *Journal of Artificial Intelligence Research* **35**, 623–675.
- Pednault, E. P. 1987. Formulating multiagent, dynamic-world problems in the classical planning framework. In *Reasoning About Actions & Plans*, 47–82. Elsevier.
- Percassi, F. & Gerevini, A. E. 2019. On compiling away PDDL3 soft trajectory constraints without using automata. In *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS*, 320–328.

- Percassi, F., Scala, E. & Vallati, M. 2021. Translations from discretised PDDL+ to numeric planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, **31**, 252–261.
- Piacentini, C., Alimisis, V., Fox, M. & Long, D. 2013. Combining a temporal planner with an external solver for the power balancing problem in an electricity network. In *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling, ICAPS, AAAI*.
- Ramrez, M., Papisimeon, M., Lipovetzky, N., Benke, L., Miller, T., Pearce, A. R., Scala, E. & Zamani, M. 2018. Integrated hybrid planning and programmed control for real time UAV maneuvering. In *Proceedings of AAMAS*, 1318–1326.
- Riddle, P., Barley, M. & Franco, S. 2013. Problem reformulation as meta-level search. In *Proceedings of the Conference on Advances in Cognitive Systems*.
- Riddle, P., Barley, M., Franco, S. & Douglas, J. 2015a. Automated transformation of PDDL representations. In *International Symposium on Combinatorial Search*.
- Riddle, P., Barley, M., Franco, S. & Douglas, J. 2015b. Bagged representations in PDDL. In *4th Workshop on the International Planning Competition*, 17–23.
- Riddle, P., Barley, M., Franco, S. & Douglas, J. 2015c. Analysis of bagged representations in PDDL. In *Workshop on Heuristics and Search for Domain-independent Planning*, 17–23.
- Riddle, P., Douglas, J., Barley, M. & Franco, S. 2016. Improving performance by reformulating PDDL into a bagged representation. In *Proceedings of the 8th Workshop on Heuristic Search for Domain-independent Planning (HSDIP@ ICAPS)*, 28–36.
- Riddle, P. J., Holte, R. C. & Barley, M. W. 2011. Does representation matter in the planning competition? In *Ninth Symposium of Abstraction, Reformulation, and Approximation*.
- Sanner, S. 2010. Relational dynamic influence diagram language (rddl): language description. *Unpublished ms. Australian National University* **32**, 27.
- Scala, E. 2014. Plan repair for resource constrained tasks via numeric macro actions. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS'14*, 280–288. AAAI Press. ISBN 9781577356608.
- Serina, I. 2010. Kernel functions for case-based planning. *Artificial Intelligence* **174**(16–17), 1369–1406.
- Siddiqui, F. H. & Haslum, P. 2012. Block-structured plan deordering. In *Australasian Joint Conference on Artificial Intelligence*, 803–814.
- Taig, R. & Brafman, R. I. 2013. Compiling conformant probabilistic planning problems into classical planning. In *Twenty-Third International Conference on Automated Planning and Scheduling*.
- Vallati, M., Chrupa, L. & Kitchin, D. 2013. An automatic algorithm selection approach for planning. In *2013 IEEE 25th International Conference on Tools with Artificial Intelligence*, 1–8. IEEE.
- Vallati, M., Chrupa, L. & McCluskey, T. 2017. Improving a planner's performance through online heuristic configuration of domain models. In *Proceedings of the 10th International Symposium on Combinatorial Search (SoCS 2017)*, 171–172.
- Vallati, M., Chrupa, L., McCluskey, T. L. & Hutter, F. 2021. On the importance of domain model configuration for automated planning engines. *Journal of Automated Reasoning* **65**(6), 727–773.
- Vallati, M., Chrupa, L. & Serina, I. 2020. Mevo: a framework for effective macro sets evolution. *Journal of Experimental & Theoretical Artificial Intelligence* **32**(4), 685–703.
- Vallati, M., Hutter, F., Chrupa, L. & McCluskey, T. L. 2015. On the effective configuration of planning domain models. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Vallati, M. & Kitchin, D. 2020. *Knowledge Engineering Tools and Techniques for AI Planning*. Springer.
- Vallati, M. & Serina, I. 2018. A general approach for configuring PDDL problem models. In *Twenty-Eighth International Conference on Automated Planning and Scheduling*.
- Younes, H. L. & Littman, M. L. 2004. PPDDL1.0: an extension to PDDL for expressing planning domains with probabilistic effects. Technical Report CMU-CS-04-162, **2**, 99.