

Automatic encoding and repair of reactive high-level tasks with learned abstract representations

The International Journal of
Robotics Research
2023, Vol. 42(4-5) 263–288
© The Author(s) 2023
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649231167207
journals.sagepub.com/home/ijr



Adam Pacheck¹ , Steven James² , George Konidaris³ and Hadas Kress-Gazit¹

Abstract

We present a framework for the automatic encoding and repair of high-level tasks. Given a set of skills a robot can perform, our approach first abstracts sensor data into symbols and then automatically encodes the robot's capabilities in Linear Temporal Logic (LTL). Using this encoding, a user can specify reactive high-level tasks, for which we can automatically synthesize a strategy that executes on the robot, if the task is feasible. If a task is not feasible given the robot's capabilities, we present two methods, one enumeration-based and one synthesis-based, for automatically suggesting additional skills for the robot or modifications to existing skills that would make the task feasible. We demonstrate our framework on a Baxter robot manipulating blocks on a table, a Baxter robot manipulating plates on a table, and a Kinova arm manipulating vials, with multiple sensor modalities, including raw images.

Keywords

Task repair, skill encoding, abstraction generation

1. Introduction

Generally useful robots will be required to generate intelligent behavior from high-level task specifications, especially if they are to be used by non-experts. Robots should have the ability to reason about their actions (or *skills*), a task's goals and its constraints, and generate the behavior necessary to achieve the task, autonomously. One promising formalism for describing tasks and skills is Linear Temporal Logic (LTL) (Pnueli 1977). LTL allows one to encode (i) skills that have nondeterministic outcomes, (ii) safety constraints, (iii) reactive tasks, where the robot responds to the environment, and (iv) tasks with complex goals that go beyond reaching a goal state (as is typical in planning languages such as PDDL (McDermott et al., 1998)). Furthermore, there exist different algorithms that enable a robot to synthesize a controller that is guaranteed to complete a specified task for fragments of LTL, such as generalized reactivity(1) (GR(1)) (Bloem et al., 2012).

However, writing LTL specifications is not trivial. Since it is a discrete logic, it requires an abstraction of the problem. Often, abstractions are handcrafted or constructed from a simplified model of the world which may not fully capture the outcomes of the robot's skills. Recently, work has looked at creating abstractions directly from sensor data (Konidaris et al., 2018; Jetchev et al., 2013; Mugan and Kuipers 2009; Ugur and Piater 2015a), but those approaches have not been extended to seamlessly integrate with mission

specification approaches that employ formal languages such as LTL.

In approaches to synthesizing controllers from LTL specifications, the robot skills and task are encoded as LTL formulas and then the algorithms find a strategy such that the task is guaranteed to be achieved, if feasible (Kress-Gazit et al., 2018). However, if the task is not possible given the current skills of the robot, it is difficult to understand why, much less what needs to be done to make the task possible. Several methods for debugging LTL specifications have been proposed (e.g., Raman and Kress-Gazit (2013); Chatterjee et al. (2008); Könighofer et al. (2009)) along with methods for suggesting modifications to specifications (e.g., Pacheck et al. (2019, 2020); Fainekos (2011); Kim et al. (2015)). When debugging specifications, the user still needs to decide how to repair the specification. In general, methods for finding suggestions that repair a specification often require changing the task, rather than changing the robot's skills to allow the robot to complete the task as specified.

¹Cornell University, Ithaca, NY, USA

²University of the Witwatersrand, Johannesburg, South Africa

³Brown University, Providence, RI, USA

Corresponding author:

Adam Pacheck, Department of Mechanical Engineering, Cornell University, 124 Hoy Rd, Ithaca, NY 14853, USA.
Email: akp84@cornell.edu

Here, we build on the work in [Pacheck et al. \(2019\)](#) and [Pacheck et al. \(2020\)](#). As in [Pacheck et al. \(2019\)](#), we demonstrate automatically encoding robot skills (that are given to the robot by a user) in LTL with symbols generated from sensor data using the methods proposed in [Konidaris et al. \(2018\)](#). In addition, we use the methods in [Konidaris et al. \(2018\)](#) to generate symbols from raw images, which has not been demonstrated before. By creating abstractions, and encoding skills directly from sensor data, we are able to take into account unmodeled nondeterminism in a robot's skills without having to hand design abstractions. A user is then able to use these abstractions to specify a high-level task for the robot. If the robot is always able to accomplish the task, we can use existing methods to generate a strategy that will guarantee that the robot will accomplish the task (e.g., [Bloem et al. \(2012\)](#)). If the task is not possible, we propose two algorithms for repair—*enumeration-based* and *synthesis-based*—that will automatically suggest additional skills or modifications to skills that would allow the robot to successfully accomplish the task as specified by the user. In this work, checking whether the task is possible and, if not, repairing it, is done offline, before task execution. The work in this paper extends [Pacheck et al. \(2019\)](#) and [Pacheck et al. \(2020\)](#); we extend the enumeration-based repair algorithm of [Pacheck et al. \(2019\)](#) to allow for suggestions containing more than one skill. Additionally, we expand the synthesis-based repair algorithm of [Pacheck et al. \(2020\)](#) to allow for the repair of reactive tasks, where the robot behavior depends on the behavior of the uncontrolled environment. Based on the skill suggestions, it is the responsibility of the user to implement the physical skills on the robot. We demonstrate our approach on two new tasks and a new physical robot (the Kinova arm in addition to the Baxter).

Contributions: Given a set of skills a robot is able to perform, we present a framework that uses sensor data to automatically create an abstraction and encode skills in LTL and then, given a task written as an LTL formula over the abstraction, provides skill suggestions to repair infeasible tasks. Specifically, (a) We propose a method to automatically encode the robot capabilities into LTL, using propositions generated directly from sensor data as proposed in [Konidaris et al. \(2018\)](#). The encoded capabilities are then used to automatically synthesize high-level robot behaviors to accomplish a user-specified task. (b) If a user-specified task is not feasible due to a missing skill, we present two approaches—enumeration-based and synthesis-based—for automatically suggesting skills that repair the task (i.e., make it executable by the robot). (c) We demonstrate our approach on two physical systems: a Baxter robot manipulating blocks and pushing plates, and a Kinova arm manipulating vials; furthermore, we demonstrate the symbol generation process of [Konidaris et al. \(2018\)](#) using raw images in addition to previously used sensor modalities.

2. Related work

This work deals with abstractions, planning, synthesis, and repair. Each of these areas is often dealt with individually, with overlap between some, but rarely all, of the areas.

Abstraction Creation: To enable robots to perform high-level tasks, the robot's capabilities, state, and environment are typically abstracted into predicates that include the robot's skills and their effects on the robot's state and environment. These predicates are often abstractions of the state space ([Kress-Gazit et al., 2018](#); [Mazo et al., 2010](#); [Finucane et al., 2010](#)).

There are several existing approaches to generating abstractions directly from low-level observations. These include learning symbols to model an agent's skills ([Konidaris et al., 2018](#)), modeling an agent with parameterized actions ([Ames et al., 2018](#)), and learning agent-centric symbols that can be transferred to new tasks in simple video games ([James et al., 2020](#)). [Ugur and Piater \(2015a, b\)](#) learn object-centric representations for a manipulation task. While their system can be used for symbolic planning on a physical robot, object features are specified prior to learning. This approach is extended to learn representations directly from raw image data using a neural network ([Ahmetoglu et al., 2020](#)). In both cases, however, certain predicates are manually inserted to generate a sound representation.

[Jetchev et al. \(2013\)](#) learn relational symbols and operators directly from geometric data. However, the size of the search space is large, requiring one dimension for every parameter of every symbol, which restricts its ability to scale to large problems. [Mugan and Kuipers \(2009, 2011\)](#) iteratively discretize a continuous state space to construct a model suitable for planning. Skills are then learned to reach these discretized states. This can be seen as a “symbols-first” approach, where skills are learned to achieve an initial discretization, which is then refined as necessary.

[Asai and Fukunaga \(2018\)](#) learn deterministic action operators directly from pixels using an autoencoder, where the bottleneck layer represents the set of propositions set to true and false. However, it is unclear how to extend the approach to the stochastic setting. A similar approach produces deterministic object-centric representations ([Asai 2019](#)), but the symbols are encoded implicitly and cannot be transformed into a language that can be used by existing planners.

While all of these approaches to generate abstraction can learn representations for planning, they lack a mechanism to correct a model that is imperfect or insufficient to solve a given task. Recent works have begun to bridge the gap of modifying models that are insufficient to solve a task ([Pacheck et al., 2019, 2020](#)).

Planning: Planning algorithms and frameworks use abstractions to find a sequence of commands to reach a goal state ([Fikes and Nilsson 1971](#); [Fox and Long 2003](#); [Ghallab et al., 2016](#)). If there is uncertainty in the outcome of skills,

planners exist that will return the sequence of skills that is most likely to accomplish the task, but often require re-planning when an unexpected effect occurs (e.g., Yoon et al. (2007)). When there is uncertainty in sensing, initial state, and actuation, conditional planning can return a plan that will take these into account (Ghallab et al., 2016). If the robot cannot observe all of its environment, conformant planning generates a plan for a robot to accomplish its goal (Ghallab et al., 2016). However, these goals are typically defined as a desired end state, while we consider more complex tasks. Additionally, if a planner fails to find a plan, to the best of our knowledge, planners are unable to suggest new skills (“actions” in the planning domain) that result in a valid plan.

Task and Motion Planning: Task and Motion Planning (TAMP) uses actions or skills that have preconditions and postconditions, often defined using a STRIPS-style framework (Fikes and Nilsson 1971) (for a review see Garrett et al. (2021)). In TAMP, tasks are often to achieve a goal state, or set of states, while in this work, we specify objectives as LTL formulas, that naturally capture different desired task properties such as conditional statements, safety constraints, sequencing of goals and reactions to uncontrolled events. Given a task, TAMP finds a solution by typically iterating between symbolic search, and geometric motion planners; TAMP may be performed offline in a simulated environment, or online with a physical robot. In our approach, we check before execution if a given set of skills can be composed to solve the task, and if not, perform repair offline. In general, TAMP works well for tasks that require the robot to ensure the environment reaches a desired state (e.g., rearranging objects on shelves) and where motion and manipulation occurs in highly cluttered environments (e.g., Srivastava et al. (2014); Muhayyuddin et al. (2017); Lee et al. (2021); Dantam et al. (2018)). The approach proposed here is better suited for tasks that can be performed with a small, well designed, set of skills but that capture complex, temporally extended objectives that include reachability, safety, and conditional constraints.

Synthesis: Work in synthesis for robotics from temporal logic specifications (Kress-Gazit et al., 2018) allows one to specify a reactive high-level task for a robot and produce either a strategy guaranteed to succeed, or a proof that the task cannot be accomplished (e.g., Lahijanian et al. (2012); He et al. (2018); Wongpiromsarn et al. (2010); DeCastro and Kress-Gazit (2016); Kress-Gazit et al. (2009); He et al. (2019)).

Specification Debugging: Using synthesis, robots can find a strategy to accomplish a task that accounts for all possible outcomes of their skills and changes in the environment. However, if there does not exist a strategy to accomplish a task, it is difficult even for expert users to determine the cause, much less find a solution. If completion of a task cannot be guaranteed, work has enabled synthesis algorithms to provide explanations as to what caused the problem (e.g., Raman and Kress-Gazit (2013); Chatterjee et al. (2008); Könighofer et al. (2009)). For fragments of

LTL, such as GR(1), synthesis algorithms can also produce counterstrategies that provide details on why the specification cannot be satisfied (Könighofer et al., 2009). Tools such as Slugs (Ehlers and Raman 2016) exist that allow users to step through strategies and counterstrategies for debugging purposes.

Specification Repair: Beyond providing methods to debug specifications, work has proposed methods to provide repair suggestions and automated fixes to be made to specifications (Alur et al., 2013; Li et al., 2011). Some frameworks allow for certain aspects of a specification to be skipped at runtime if they are not possible (Lahijanian et al., 2016). In Lahijanian et al. (2016), the task is split into safety constraints that must always be satisfied and liveness guarantees that should be satisfied if possible. Work by Fainekos (2011) and Kim et al. (2015) also consider revising specifications.

The work mentioned above (Fainekos 2011; Kim et al., 2015; Alur et al., 2013; Li et al., 2011; Lahijanian et al., 2016) focuses on restricting the behavior of the environment or modifying the goals of the robot to make the task possible. In this work, we provide suggestions that extend the capabilities of the robot through additional skills or modifications to skills that are grounded in the sensor-based abstract representation. Instead of changing what we would like the robot to do, we give it additional capabilities that allow it to accomplish the original desired task.

3. Preliminaries

3.1. Skills

We model the abilities of the robot as a set of skills, \mathcal{A} , designed and given to the robot by the user, operating over a world with a continuous state space $(x_1, \dots, x_n) \in X \subseteq \mathbb{R}^n$. This state space may include the location of objects of interest, the robot’s end effector, etc. In the case of a high-dimensional state space, such as an image, we use independent component analysis to reduce the dimensionality of the state space (Hyvärinen and Oja 2000; Konidaris et al., 2018). Each skill $a \in \mathcal{A}$ has a region from which it is applicable, termed the *precondition* of a , $\text{Pre}(a) \subseteq X$. The application of a will result in the state being in the *effect set* $\text{Eff}(a) \subseteq X$. We introduce the example in Figure 1 to illustrate the main ideas of skills and symbol generation. In this two-dimensional space, a robot has two skills a_1 and a_2 that allow it to move between regions, as shown by the arrows.

3.2. Symbol generation

The process of symbol generation (Konidaris et al., 2018) automatically constructs a set of symbols which are used for planning. This process assumes a set of data, $D \subset X \times \mathcal{A} \times X$ has been collected by randomly executing robot skills. Each element of the data consists of $(x^{pre}, a, x^{post}) \in D$ where $x^{pre} \in \text{Pre}(a)$, $a \in \mathcal{A}$, and $x^{post} \in \text{Eff}(a)$. During

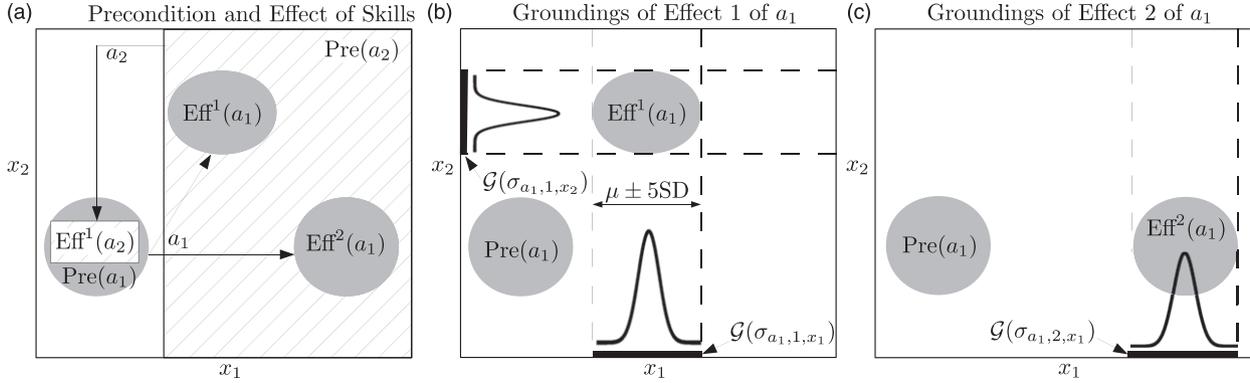


Figure 1. Example demonstrating the symbol generation process. (a) Two skills a_1 and a_2 and their precondition and effect sets. (b, c) The grounding sets of the symbols generated from skill a_1 . The robot needs to consider the value of both x_1 and x_2 when deciding if it can apply a_1 , so $\text{pre-mask}(a_1) = \{\text{True}, \text{True}\}$. The robot only needs to consider the value of x_1 when deciding if it can apply a_2 , so $\text{pre-mask}(a_2) = \{\text{True}, \text{False}\}$. The application of a_1 either changes x_1 and x_2 or only x_1 , so $\text{eff-mask}^1(a_1) = \{\text{True}, \text{True}\}$ and $\text{eff-mask}^2(a_1) = \{\text{True}, \text{False}\}$. In effect 1 of a_1 , $\sigma_{\text{eff}^1(a_1)}^{\text{True}} = \{\sigma_{a_1,1,x_1}, \sigma_{a_1,1,x_2}\}$ become True and $\sigma_{\text{eff}^1(a_1)}^{\text{False}} = \{\sigma_{a_1,2,x_1}, \sigma_{a_2,1,x_1}, \sigma_{a_2,1,x_2}\}$ become False. In effect 2 of a_1 , $\sigma_{\text{eff}^2(a_1)}^{\text{True}} = \{\sigma_{a_1,2,x_1}\}$ becomes True, $\sigma_{\text{eff}^2(a_1)}^{\text{False}} = \{\sigma_{a_1,1,x_1}, \sigma_{a_2,1,x_1}\}$ becomes False, and $\sigma_{\text{eff}^2(a_1)}^{\text{stay}} = \{\sigma_{a_1,1,x_2}, \sigma_{a_2,1,x_2}\}$ do not change. Figure from Pacheck et al. (2019).

symbol generation, we cluster the effect set of a skill into distinct effect sets using DBSCAN (Ester et al., 1996; Konidaris et al., 2018). Each effect set $\text{Eff}(a)$ is split, as needed, into $j \in \{1, \dots, k(a)\}$ possible *effect sets*, denoted by $\text{Eff}^j(a) \subseteq X$. In Figure 1, a_1 can be executed from $\text{Pre}(a_1)$ and has a nondeterministic outcome, resulting in either $\text{Eff}^1(a_1)$ or $\text{Eff}^2(a_1)$.

The values of some x_i may matter in determining whether a skill can be applied, while the values of others may not. We denote this in the *precondition mask* of a , $\text{pre-mask}(a)[a] \in \mathbb{B}^n$, where $\text{pre-mask}(a)(i) = \text{True}$ whether the value of x_i influences if a can be applied, and False otherwise. We create a classifier to test inclusion in $\text{Pre}(a)$, which is defined for x_i for which $\text{pre-mask}(a)(i) = \text{True}$. Similarly, when a skill is applied, it may change some or all of the state variables. We denote this in the *effect mask*, $\text{eff-mask}^j(a) \in \mathbb{B}^n$, where $\text{eff-mask}^j(a)(i) = \text{True}$ if the value of x_i is modified by the application of a in the j^{th} outcome and False otherwise. In Figure 1, to apply a_1 , the values of both x_1 and x_2 matter, so $\text{pre-mask}(a_1) = \{\text{True}, \text{True}\}$. However, we need only consider the value of x_1 to determine if a_2 can be applied, so $\text{pre-mask}(a_2) = \{\text{True}, \text{False}\}$. Effect 1 of a_1 changes the value of x_1 and x_2 so $\text{eff-mask}^1(a_1) = \{\text{True}, \text{True}\}$, while effect 2 only changes the value of x_1 , so $\text{eff-mask}^2(a_1) = \{\text{True}, \text{False}\}$.

We define factors $f \in F \subset 2^X$ that denote which state variables x_i always change together. Here, F is the set of all factors. A separate symbol $\sigma \in \Sigma$ is created¹ for each a, j , and f when $\text{eff-mask}^j(a)(i) = \text{True} \forall x_i \in f$. We add subscripts to σ and say each $\sigma_{a,j,f}$ grounds to a set over the state variables $x_i \in f$. We fit either a Gaussian or use Kernel Density Estimation with a Gaussian kernel with mean μ and standard deviation sd . We consider $\mathcal{G}(\sigma_{a,j,f}) = \{x \in X \mid |x - \mu| < 5sd\} \subseteq X$ to be the set of states spanned by five standard deviations from the mean of the fit. If the raw data is found to be the same by a two-sample Kolmogorov-Smirnov test

(Hollander et al., 2014) or in the case of higher dimensions if the mean and variance are similar (Konidaris et al., 2018), the two symbols are merged into one symbol. The set of symbols referring to a single factor f is $\Sigma_f = \{\sigma_{a,j,f} \mid a \in \mathcal{A}, j \in \{1, \dots, k(a)\}\}$. The set of all symbols is $\Sigma = \cup_{f \in F} \Sigma_f$. For readability, when $f = \{x_i\}$, we denote $\sigma_{a,j,f}$ as σ_{a,j,x_i} .

In Figure 1(b), $\text{Eff}^1(a_1)$ results in two symbols, $\sigma_{a_1,1,x_1}$ and $\sigma_{a_1,1,x_2}$, because $\text{eff-mask}^1(a_1) = \{\text{True}, \text{True}\}$. Only one symbol, $\sigma_{a_1,2,x_1}$ is generated from $\text{Eff}^2(a_1)$ as $\text{eff-mask}^2(a_1) = \{\text{True}, \text{False}\}$. In this example, all factors are singletons.

During the symbol generation process, skills that have different effects from different preconditions are partitioned into multiple skills (Konidaris et al., 2018).

3.3. Linear Temporal Logic

Let \mathcal{AP} be a set of atomic propositions and $\pi \in \mathcal{AP}$ be a Boolean variable. The syntax of a formula in LTL (Pnueli 1977) obeys the following grammar:

$$\varphi ::= \pi \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

where negation (\neg , “not”) and disjunction (\vee , “or”) are Boolean operators and \bigcirc (“next”) and \mathcal{U} (“until”) are temporal operators. We define $\text{True} = \varphi \vee \neg\varphi$ and $\text{False} = \neg\text{True}$. Given these operators, one can derive conjunction ($\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2)$), implication ($\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$), equivalence ($\varphi_1 \leftrightarrow \varphi_2 \equiv (\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_2 \rightarrow \varphi_1)$), eventually ($\diamond\varphi \equiv \text{True} \mathcal{U} \varphi$), and always ($\square\varphi \equiv \neg\diamond\neg\varphi$). We define a *symbolic state* as the set of all propositions that are currently True and denote all possible symbolic states by $\mathcal{V}_{\mathcal{AP}} = 2^{\mathcal{AP}}$. We use $\mathcal{AP}' = \{\pi' \mid \pi \in \mathcal{AP}\}$ as the set of primed versions of the variables in \mathcal{AP} to denote variables at the next time step. The set of all possible symbolic states at the next time step is $\mathcal{V}_{\mathcal{AP}'} = 2^{\mathcal{AP}'}$.

The semantics of an LTL formula φ are defined over an infinite sequence $w = w_1w_2\dots$ (Pnueli 1977). Each w_i corresponds to the set of π that are True at step i . We denote that a sequence w satisfies an LTL formula at instance i by $w, i \models \varphi$. Intuitively, $w, i \models \bigcirc\varphi$ if φ is True at step $i+1$, $w, i \models \square\varphi$ if φ holds at every step after and including i in w , and $w, i \models \diamond\varphi$ if φ holds at some step on or after i in w .

We consider the GR(1) fragment of LTL (Bloem et al., 2012). Let $\mathcal{AP} = \mathcal{E} \cup \mathcal{S} = \Sigma \cup \mathcal{R} \cup \mathcal{S}$ be the set of atomic propositions, where $\mathcal{E} = \Sigma \cup \mathcal{R}$ is the learned symbols Σ and additional user-defined symbols \mathcal{R} (which correspond to uncontrolled signals that a user requires the robot to respond to), and \mathcal{S} refers to the activation of robot skills \mathcal{A} (i.e., when $a \in \mathcal{S}$ is True, $a \in \mathcal{A}$ is executed). The state of the world is represented by the truth values of the propositions in \mathcal{E} . In GR(1), formulas are of the form:

$$\begin{aligned} \varphi &= \varphi_e \rightarrow \varphi_s \\ \varphi_e &= \varphi_i^e \wedge \varphi_t^e \wedge \varphi_g^e \\ \varphi_s &= \varphi_i^s \wedge \varphi_t^s \wedge \varphi_g^s \end{aligned} \quad (1)$$

where φ_e are assumptions about the environment's behavior and φ_s are guarantees for the robot, also referred to as the system, and:

- φ_i^e and φ_i^s are predicates over \mathcal{E} and $\mathcal{E} \cup \mathcal{S}$, respectively, characterizing the initial states.
- φ_t^e and φ_t^s are safety constraints of the form $\bigwedge_i \square \psi_i$ where ψ_i are over v and $\bigcirc u$ where $v \in \mathcal{E} \cup \mathcal{S}$ for φ_t^e and φ_t^s , and $u \in \mathcal{E}$ for φ_t^e and $u \in \mathcal{E} \cup \mathcal{S}$ for φ_t^s .
- φ_g^e and φ_g^s are the liveness requirements and characterize events that should occur infinitely often. Here, $\varphi_g^e = \bigwedge_{i=1}^m \square \diamond J_i^e$ and $\varphi_g^s = \bigwedge_{j=1}^n \square \diamond J_j^s$ where J_i^e and J_j^s are predicates over $\mathcal{E} \cup \mathcal{S}$.

An implementation of the specification is guaranteed to satisfy φ_s , provided that the environment satisfies φ_e .

3.4. Synthesis

We use GR(1) synthesis (Bloem et al., 2012) to find a strategy to accomplish a task. In this work, a task consists of a set of system liveness guarantees (φ_g^s), initial conditions ($\varphi_i^e \wedge \varphi_i^s$), and a set of ‘‘hard’’ system safety guarantees (φ_t^s). When synthesizing, we solve a two-player game played between a system and its environment where the system reacts to the environment. The environment is considered to be adversarial and attempts to keep the system from accomplishing its task. This ensures that the system is able to accomplish its task regardless of what happens in the environment. We define a *game structure* $G = (\mathcal{AP}, \mathcal{E}, \mathcal{S}, \theta^{init}, \tau_e, \tau_s, \tau_s^{\text{hard}}, \Phi)$ where \mathcal{AP} , \mathcal{E} , and \mathcal{S} are as defined in Section 3.3. We define θ^{init} as the set of states that satisfy $\varphi_i^e \wedge \varphi_i^s$. We define $\tau_e \subseteq \mathcal{V}_{\mathcal{AP}} \times \mathcal{V}_{\mathcal{E}'}$ as the set of current and next states satisfying φ_t^e , $\tau_s \subseteq \mathcal{V}_{\mathcal{AP}} \times \mathcal{V}_{\mathcal{AP}'}$ as the set of current and next states satisfying φ_t^s , and $\tau_s^{\text{hard}} \subseteq \mathcal{V}_{\mathcal{AP}} \times \mathcal{V}_{\mathcal{AP}'}$ as the set of current

and next states satisfying the hard system constraints in φ_t^s . Hard system constraints in τ_s^{hard} cannot be modified by the synthesis-based repair process in Section 6.2. Note that in Bloem et al. (2012), τ_e and τ_s are defined as logical formulas; here, we define them as sets of states. The winning condition is given by $\Phi = \varphi_g^e \rightarrow \varphi_g^s$.

Given a game structure G , the *realizability* problem is to decide if the game is winning for the system; either (a) for every environment action, the system is able to achieve φ_s or (b) the system is able to falsify φ_e . To determine if a specification is realizable, we find via a fixed point computation all the states Z from which the system is able to win (Bloem et al., 2012). We iterate through every system liveness guarantee, J_j^s , and determine the set of states the system can always either transition to the next liveness goal from or falsify φ_e . The *synthesis* problem is to compute a strategy for the system to make the specification realizable (Bloem et al., 2012).

We define a strategy computed using the synthesis process as $C = (\mathcal{E}, \mathcal{S}, Q, Q_0, \delta, L)$, where:

- \mathcal{E} and \mathcal{S} are the environment and system propositions, respectively, defined above
- Q is a set of states
- $Q_0 \subseteq Q$ is the set of initial states
- $\delta: Q \times 2^{\mathcal{E}} \rightarrow Q$ is the transition function
- $L: Q \rightarrow 2^{\mathcal{E}} \times 2^{\mathcal{S}}$ is a labeling function that returns the propositions in $\mathcal{E} \cup \mathcal{S}$ that are True in state $q \in Q$

Here, δ depends on \mathcal{E} as the system reacts to the environment state.

If equation (1) is unrealizable, meaning that there does not exist a strategy C that will satisfy the task, the synthesis algorithm can provide a *counter-strategy* that represents the behavior of the environment that will cause the system to fail to accomplish its task (Könighofer et al., 2009; Chatterjee et al., 2008). We define a counter-strategy as $C_{c.s.} = (\mathcal{E}, \mathcal{S}, Q, Q_0, Q_{n.o.t.}, \delta_{c.s.}, L_t, L_{n.o.t.})$, where $\mathcal{E}, \mathcal{S}, Q, Q_0$ are the same as in C and:

- $Q_{n.o.t.} \subseteq Q$ is the set of states from which there are no outgoing transitions (n.o.t.)
- $\delta_{c.s.}: Q \setminus Q_{n.o.t.} \times 2^{\mathcal{E}} \rightarrow Q$ is the transition function
- $L_t: Q \setminus Q_{n.o.t.} \rightarrow 2^{\mathcal{E}} \times 2^{\mathcal{S}}$ is the labeling function for states with outgoing transitions
- $L_{n.o.t.}: Q_{n.o.t.} \rightarrow 2^{\mathcal{E}}$ is the labeling function for states with no outgoing transitions. The system has no valid transitions from $Q_{n.o.t.}$, so only \mathcal{E} is needed to label $Q_{n.o.t.}$.

In Section 6.1, we use the states with no outgoing transitions, $Q_{n.o.t.}$, to narrow the search for skills to repair unrealizable specifications in the enumeration-based repair approach.

4. Problem formulation

Our goal is to automatically encode the capabilities of a robot in an LTL formula and find a strategy for a reactive

high-level task. If no strategy can be found, we find additional skills or modifications to skills that would allow the robot to complete the given task.

Problem 1. Given a set of skills \mathcal{A} , automatically abstract and encode the capabilities of the robot in an LTL formula, φ_{skills} . Allow a user to specify a reactive high-level task and find a strategy to fulfill it.

Problem 2. Given an unrealizable specification φ_{unreal} , find skill suggestions \mathcal{A}_{new} , in the form of additional skills or modifications to current skills, such that constructing φ_{skills} with $\mathcal{A} \cup \mathcal{A}_{\text{new}}$ makes the specification φ_{unreal} realizable.

5. Specification encoding

To address Problem 1, we automatically encode the robot's capabilities in φ_{skills} using the symbols in Σ , which are learned from low-level sensor information (Konidaris et al., 2018), and the skills \mathcal{A} of the robot. The skills-based specification, φ_{skills} , can be reused for different tasks performed by the same robot. The user then writes the task-specific specification, φ_{task} , over \mathcal{AP} , which is combined through conjunction with φ_{skills} to create φ_{full} . We use a synthesis tool, such as Slugs (Ehlers and Raman 2016), to either find a strategy, C , for accomplishing φ_{full} if the specification is realizable or a counter-strategy, $C_{\text{c.s.}}$, if the specification is unrealizable. An overview of the framework is depicted in Figure 2.

5.1. Skills-based specification (φ_{skills})

The skills-based specification encodes the preconditions and postconditions of skills, along with mutual exclusion constraints on the skills and symbols.

Given a set of skills \mathcal{A} , we first create symbols $\sigma \in \Sigma$, representing the effects of $a \in \mathcal{A}$ (Konidaris et al., 2018). We slightly abuse notation and use a as a proposition that is True when the skill a is active, and False otherwise.

The skills-based specification (φ_{skills}) is composed of the system safety ($\varphi_{\text{t,skills}}^{\text{s}} = \varphi_{\text{t,pre}}^{\text{s}} \wedge \varphi_{\text{t,mx_skills}}^{\text{s}}$) and environment safety ($\varphi_{\text{t,skills}}^{\text{e}} = \varphi_{\text{t,eff}}^{\text{e}} \wedge \varphi_{\text{t,no_act}}^{\text{e}} \wedge \varphi_{\text{t,mx_syms}}^{\text{e}}$) specifications. The system safety specification includes constraints on when the system is allowed to perform skills ($\varphi_{\text{t,pre}}^{\text{s}}$) and optionally the mutual exclusion of skills ($\varphi_{\text{t,mx_skills}}^{\text{s}}$). The environment safety specification includes how each σ is allowed to change with the application of a skill ($\varphi_{\text{t,eff}}^{\text{e}}$), the effect of no skill being performed ($\varphi_{\text{t,no_act}}^{\text{e}}$), and the mutual exclusion of symbols over the same factor ($\varphi_{\text{t,mx_syms}}^{\text{e}}$).

5.1.1. System safety ($\varphi_{\text{t,skills}}^{\text{s}}$). We encode constraints on when skills can be performed in $\varphi_{\text{t,pre}}^{\text{s}}$ based on the preconditions of the skills. For each action, we find all possible combinations of symbols that overlap with the precondition mask and determine which combinations fall within the precondition set (Konidaris

et al., 2018). Note that there may be multiple combinations of symbols from which the robot is allowed to execute a skill. We define $\sigma_{\text{pre}(a)} = \{\sigma_p \in \Sigma_{\text{pre-mask}(a)} \mid \mathcal{G}(\sigma_p) \subseteq \text{Pre}(a)\}$, where $\Sigma_{\text{pre-mask}(a)} = \prod_{f \in \mathcal{F}, \text{s.t. } \forall x_i \in f, \text{pre-mask}(a)[a(i)=\text{True}] \Sigma_f$. The set $\sigma_{\text{pre}(a)}$ contains all the combinations of σ that satisfy the precondition of a . We encode in $\varphi_{\text{t,pre}}^{\text{s}}$ that when none of the preconditions in the set $\sigma_{\text{pre}(a)}$ are satisfied, the robot is not allowed to perform a as shown in equation (2). Equation (2) states that skill a cannot be executed at the next step when no combinations of symbols $\sigma_p \in \sigma_{\text{pre}(a)}$ are True at the next step. This allows the robot to choose to execute a skill only when the preconditions of a skill are satisfied, essentially encoding “do not use the skill unless a precondition is met.” This is in contrast to encoding “if the precondition is true then execute the skill” which requires the robot to use a skill when it is available, even if the skill is not needed or it may lead to undesired behavior. We write $\varphi_{\text{t,pre}}^{\text{s}}$ over $\mathcal{V}_{\mathcal{AP}'}$ instead of over $\mathcal{V}_{\mathcal{AP}}$ as in Pacheck et al. (2019) to generate additional types of skill suggestions in Section 6.2 and match assumptions made in Pacheck et al. (2020). In Figure 1, $\sigma_{\text{pre}(a_2)}[a_2] = \{\{\sigma_{a_1,1,x_1}\}, \{\sigma_{a_1,2,x_1}\}\}$. Following $\varphi_{\text{t,pre}}^{\text{s}}$, this means that when both $\{\sigma_{a_1,1,x_1}\}$ and $\{\sigma_{a_1,2,x_1}\}$ are False, the robot cannot choose to execute skill a_2 . Or, equivalently, when either $\{\sigma_{a_1,1,x_1}\}$ or $\{\sigma_{a_1,2,x_1}\}$ are True, a_2 can be executed.

We can encode mutual exclusion of skills in $\varphi_{\text{t,mx_skills}}^{\text{s}}$ at both the current and next step. In the examples presented, skills are mutually exclusive, although in general they need not be.

Note that $\varphi_{\text{t,mx_skills}}^{\text{s}}$ is considered a “hard” constraint ($\varphi_{\text{t,hard}}^{\text{s}}$) for synthesis-based repair (Section 6.2) and so is not allowed to be changed. On the other hand, $\varphi_{\text{t,pre}}^{\text{s}}$ is not a “hard” constraint and can be modified during synthesis-based repair, meaning that we can modify the preconditions of the skills.

5.1.2. Environment safety ($\varphi_{\text{t,skills}}^{\text{e}}$). To encode a skill's (possibly nondeterministic) effects, we consider the skill outcome to be determined by the environment.

We denote the symbols which become True with the application of a skill a as $\sigma_{\text{eff}(a)}^{\text{True}} = \bigcup_{f \in \mathcal{F}, \text{s.t. } \forall x_i \in f, \text{eff-mask}(a)(i)=\text{True}} \sigma_{a,j,f}$ (Konidaris et al., 2018). In Figure 1, $\sigma_{\text{eff}^2(a_1)}^{\text{True}} = \{\sigma_{a_1,2,x_1}\}$.

When a is applied, symbols belonging to the same factor f whose grounding sets do not overlap with those in $\sigma_{\text{eff}(a)}^{\text{True}}$ become False due to mutual exclusion. We denote this set of symbols $\sigma_{\text{eff}(a)}^{\text{False}}$ in equation (3). In Figure 1, $\sigma_{\text{eff}^2(a_1)}^{\text{False}} = \{\sigma_{a_1,1,x_1}, \sigma_{a_2,1,x_1}\}$. The physical interpretation of the set of symbols that become False is that the robot is not in the union of those symbols. In effect 2 of a_2 , the robot is not in the union of $\mathcal{G}(\sigma_{a_1,1,x_1})$ and $\mathcal{G}(\sigma_{a_2,1,x_1})$.

When performing synthesis (Kress-Gazit et al., 2018), if a symbol is not constrained, it can be set to any value. We must therefore consider the “frame problem” (Ghallab et al., 2004) and constrain symbols that are not modified by the current skill to stay the same. The set $\sigma_{\text{eff}(a)}^{\text{stay}} = \bigcup_{f \in \mathcal{F}, \text{s.t. } \forall x_i \in f, \text{eff-mask}(a)(i)=\text{False}} \Sigma_f$ contains the σ not modified

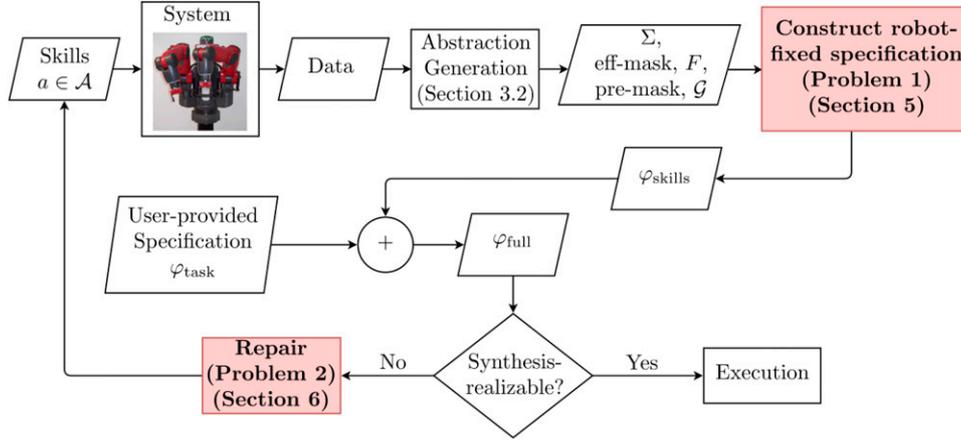


Figure 2. Framework for automatically encoding robot capabilities, executing tasks, and repairing unrealizable tasks. Novel contributions are in bold and red. Figure adapted from Pacheck et al. (2019).

by skill a in the j^{th} outcome. In Figure 1, because x_2 is not modified in effect 2 of a_1 , $\sigma_{\text{eff}^2(a_1)}^{\text{stay}} = \{\sigma_{a_1,1,x_2}, \sigma_{a_2,1,x_2}\}$.

For example, after applying skill a_1 when $\{\sigma_{a_2,1,x_1}, \sigma_{a_2,1,x_2}\}$ are True, one possible outcome is that $\sigma_{\text{eff}^2(a_1)}^{\text{True}} = \{\sigma_{a_1,2,x_1}\}$ become True, $\sigma_{\text{eff}^2(a_1)}^{\text{False}} = \{\sigma_{a_1,1,x_1}, \sigma_{a_2,1,x_1}\}$ become False, and the robot is now in the state where $\{\sigma_{a_1,2,x_1}, \sigma_{a_2,1,x_2}\}$ are True and the rest of the symbols False. The physical interpretation is that the robot in the intersection of $\mathcal{G}(\sigma_{a_1,2,x_1})$ and $\mathcal{G}(\sigma_{a_2,1,x_2})$.

$$\varphi_{t,\text{pre}}^s = \bigwedge_{a \in \mathcal{A}} \square \left[\neg \left(\bigvee_{\sigma_p \in \sigma_{\text{pre}}(a)} \left(\bigwedge_{\sigma \in \sigma_p} \sigma \right) \right) \rightarrow \neg \bigcirc a \right] \quad (2)$$

$$\begin{aligned} \sigma_{\text{eff}^j(a)}^{\text{False}} \\ = \bigcup_{f.s.t. \forall x_i \in f, \text{eff-mask}^j(a)(i) = \text{True}} \{ \sigma \in \Sigma_f \mid \mathcal{G}(\sigma) \cap \mathcal{G}(\sigma_{a,j,f}) = \emptyset \} \end{aligned} \quad (3)$$

$$\begin{aligned} \varphi_{t,\text{eff}}^c = \bigwedge_{a \in \mathcal{A}} \square \left[a \rightarrow \bigvee_{j \in \{1, \dots, k(a)\}} \left(\left(\bigwedge_{\sigma \in \sigma_{\text{eff}^j(a)}^{\text{True}}} \sigma \right) \right. \right. \\ \left. \left. \wedge \left(\bigwedge_{\sigma \in \sigma_{\text{eff}^j(a)}^{\text{False}}} \neg \sigma \right) \wedge \left(\bigwedge_{\sigma \in \sigma_{\text{eff}^j(a)}^{\text{stay}}} (\sigma \leftrightarrow \bigcirc \sigma) \right) \right) \right] \end{aligned} \quad (4)$$

We encode how the truth values for σ can change when a skill is applied in $\varphi_{t,\text{eff}}^c$ in equation (4). Equation (4) states that when skill a is performed, it leads to one of j nondeterministic outcomes with $\sigma \in \sigma_{\text{eff}^j(a)}^{\text{True}}$ becoming True, $\sigma \in \sigma_{\text{eff}^j(a)}^{\text{False}}$ becoming False, and the truth value of $\sigma \in \sigma_{\text{eff}^j(a)}^{\text{stay}}$ remaining the same. Symbols whose grounding sets overlap with those in $\sigma_{\text{eff}^j(a)}^{\text{True}}$ and are therefore not in $\sigma_{\text{eff}^j(a)}^{\text{True}}$, $\sigma_{\text{eff}^j(a)}^{\text{False}}$, or $\sigma_{\text{eff}^j(a)}^{\text{stay}}$ are not constrained. In the examples presented in this work, there are no symbols whose grounding sets overlap that have not been merged into one symbol. During the synthesis process, the adversarial environment chooses which nondeterministic outcome j would result in the worst case scenario for the system. This enables us to guarantee

that no matter what the nondeterministic effect of an action is, the system is still able to complete its task.

When no skill is performed, we encode in $\varphi_{t,\text{no-act}}^c$ that the truth values of σ remain the same.

$$\varphi_{t,\text{no-act}}^c = \square \left[\left(\bigwedge_{a \in \mathcal{A}} \neg a \right) \rightarrow \left(\bigwedge_{\sigma \in \Sigma} (\sigma \leftrightarrow \bigcirc \sigma) \right) \right] \quad (5)$$

We encode the mutual exclusion of non-overlapping symbols over the same factor in $\varphi_{t,\text{mx-syms}}^c$ at both the current and next step. We enforce that only one of the symbols in a factor is True at a time. In Figure 1, $\sigma_{a_1,1,x_1}$, $\sigma_{a_1,2,x_1}$, and $\sigma_{a_2,1,x_1}$ are all grounded over x_1 and do not overlap, so only one of them can be True at any time.

5.2. Task specification, synthesis, and execution

The user writes the task-specific specification, φ_{task} , which may include additional environment propositions $v_u \in \mathcal{R}$. The task-specific specification can include constraints on the initial state(s) of the system and environment, system liveness, and environment liveness in $\varphi_{i,\text{task}}^s$, $\varphi_{e,\text{task}}^c$, $\varphi_{g,\text{task}}^s$, and $\varphi_{g,\text{task}}^c$, respectively. Additional system safety constraints are added in $\varphi_{t,\text{task}}^s$, which we consider to be a “hard” constraint and which is not allowed to be changed during the synthesis-based repair. Tasks can encode objectives such as repeatedly accomplishing a goal or goals, always avoiding some states, always making sure a constraint holds, or reacting to environment events. We give examples of tasks and $\varphi_{t,\text{task}}^s$, $\varphi_{g,\text{task}}^c$, and $\varphi_{t,\text{task}}^s$ in Section 7.

The full specification φ_{full} is shown in equation (6). We generate a strategy for satisfying φ_{full} using a synthesis tool, such as Slugs (Ehlers and Raman 2016). If φ_{full} is realizable, the resulting strategy $C = (\mathcal{E}, \mathcal{S}, \mathcal{Q}, \mathcal{Q}_0, \delta, L)$, where $\mathcal{E} = \Sigma \cup \mathcal{R}$ and $\mathcal{S} = \mathcal{A}$, is used to control the robot. If φ_{full} is not realizable, we repair the specification using either an enumeration-based or synthesis-based repair approach (Section 6).

$$\varphi_{\text{full}} = \varphi_{i, \text{task}}^c \wedge \overbrace{\varphi_{t, \text{eff}}^c \wedge \varphi_{t, \text{no_act}}^c \wedge \varphi_{t, \text{mx_syms}}^c}^{\varphi_{t, \text{skills}}^c} \wedge \varphi_{g, \text{task}}^c \rightarrow \varphi_{i, \text{task}}^s \wedge \overbrace{\varphi_{t, \text{pre}}^s \wedge \varphi_{t, \text{mx_skills}}^s \wedge \varphi_{t, \text{task}}^s}^{\varphi_{t, \text{skills}}^s} \wedge \varphi_{g, \text{task}}^s \quad (6)$$

6. Specification repair

We address Problem 2 of making an unrealizable specification realizable by searching for additional skills or modifications to existing skills. We present and compare both an enumeration-based and synthesis-based approach, based on methods first proposed in [Pacheck et al. \(2019\)](#) and [Pacheck et al. \(2020\)](#), respectively.

6.1. Enumeration-based repair

In the enumeration-based repair approach, we search for one or more skills, $a_{\text{new}} \in \mathcal{A}_{\text{new}}$, that would make an unrealizable task realizable when φ_{skills} is constructed with $\mathcal{A} \cup \mathcal{A}_{\text{new}}$. We build on the enumeration-based repair process presented in [Pacheck et al. \(2019\)](#). There, we assume that only one skill, a_{new} is required to repair the specification. In this work, we relax that assumption and repair specifications that may need more than one additional skill. We assume the robot has all the symbols it needs to define the task. We also assume that our new skills will consist of a precondition set and effect mask we have already seen, restricting the search space for the new skills. By assuming our new skills will consist of a precondition set and effect mask we have already seen, we will not find all possible skills to repair the task and may even be unable to repair the specification. It is possible to relax these assumptions to consider all possible preconditions and postconditions; however, without these assumptions, the number of possible skills is too large to reasonably consider. In this work, we are able to find skill suggestions for all examples in Section 7 with the enumeration-based approach while making the above assumptions.

We leverage the structure of $C_{c.s.}$ to focus the repair process. The counter-strategy, $C_{c.s.}$, contains the environment behaviors that make a specification unrealizable. In general, a GR(1) specification is unrealizable either because (i) the robot violates safety constraints, (ii) gets stuck in a loop when trying to satisfy its liveness goals, or (iii) is unable to reach the liveness goals from its initial conditions. When the robot can only satisfy at least one of its liveness goals by using skills that leave the environment unable to act, the counter-strategy contains states with no successors (i.e., $Q_{n.o.t.} \neq \emptyset$). We find the skills that lead to these states, and use their precondition sets to narrow the search space for \mathcal{A}_{new} . Then, we generate new effect sets, based on existing effect masks, and combine them with existing precondition sets to create new skills.

Algorithm 1 shows our enumeration-based procedure for repairing unrealizable specifications. On Line 1, we create

new effect sets, Σ^+ , based on existing effect masks, based on the assumption that new skills will change similar states as current skills. For each existing effect mask, we find all the state variables that are in the mask. We then compute all possible combinations of $\sigma_{a,j,f}$ that ground to those state variables, regardless of which skill they were originally generated from.

On Line 2, we find $\mathcal{A}_{n.o.t.}$: the set of skills whose preconditions were satisfied that lead to states with no outgoing transitions. Based on $\mathcal{A}_{n.o.t.}$, we then construct a set of candidate skills, $\mathcal{A}_{\text{candidate-skills-n.o.t}}$ on Line 3. Each one consists of the precondition set of a skill in $\mathcal{A}_{n.o.t.}$ and a new effect set found in Σ^+ . We then construct a second set of new skills, $\mathcal{A}_{\text{candidate-skills-all}}$, in Line 4 based on the precondition sets of all skills \mathcal{A} . Slightly abusing notation, we denote candidate skills in $\mathcal{A}_{\text{candidate-skills-n.o.t}}$ and $\mathcal{A}_{\text{candidate-skills-all}}$ as pairs containing the preconditions of a skill and which symbols become True, while using the name of a skill to denote skills in $\mathcal{A}_{n.o.t.}$.

We then consider combinations of a skill in $\mathcal{A}_{\text{candidate-skills-n.o.t}}$ and $n_{\text{new-skills-desired}} - 1$ skills in $\mathcal{A}_{\text{candidate-skills-all}}$. We assume that one new skill needs to include a precondition from the skills $\mathcal{A}_{n.o.t.}$, but do not assume any other new skills need to start from one such precondition. In Lines 5–13, we write the unrealizable specification with the new skills and attempt to synthesize a strategy. If the specification is realizable, we store the skill combination. All \mathcal{A}_{new} which make φ_{full} realizable are returned to the user, enabling them to select the skill they deem easiest to physically implement.

6.2. Synthesis-based repair

Our second approach to repair unrealizable specification is synthesis-based repair where, as opposed to the enumeration-based approach, we take advantage of the synthesis process to guide the repair. We extend the synthesis-based repair introduced in [Pacheck et al. \(2020\)](#) to find suggestions for skills that can repair reactive tasks. Here, we give an overview of the process (see [Pacheck et al. \(2020\)](#) for a full description), and describe modifications we have made to the repair process that allow us to find suggestions for a larger class of specifications. With these modifications we are able to find repair suggestions for specifications with reactive liveness guarantees, which is not possible with the repair process in [Pacheck et al. \(2020\)](#).

The repair process takes an unrealizable specification and finds suggestions of new skills or modifications to existing skills. Modifications to existing skills are in the form of additional preconditions that should be added to

Algorithm 1: Enumeration-Based Repair

Input: $C_{c.s.}(\mathcal{E}, \mathcal{S}, Q, Q_0, Q_{n.o.t.}, \delta_{c.s.}, L_t, L_{n.o.t.}), n_{\text{new-skills}}, \mathcal{A}, \varphi_{\text{task}}$

Output: $\alpha_{\text{suggestions}}$

- 1 $\Sigma^+ := \bigcup_{a \in \mathcal{A}, j \in \{1, \dots, k(a)\}} \prod_{f \in F_{s.t.} v_{x_i} \in f, \text{eff-mask}^j(a)(i) = \text{True}} \Sigma_f$ // Create new postconditions
- 2 $\mathcal{A}_{n.o.t.} := \{a \in \mathcal{A} \mid \exists \sigma_p \in \sigma_{\text{pre}(a)}, q \in Q, e \in 2^{\mathcal{E}} \text{ s.t. } \delta(q, e) \in Q_{n.o.t.}, \sigma_p \in L_t(q)\}$ // Skills leading to states with no outgoing transitions
- 3 $\mathcal{A}_{\text{candidate-skills-n.o.t.}} := \{(\sigma_{\text{pre}(a_{n.o.t.})}, \sigma) \mid a_{n.o.t.} \in \mathcal{A}_{n.o.t.}, \sigma \in \Sigma^+ \text{ s.t. } \nexists a_{\text{orig}} \in \mathcal{A} \text{ where } \sigma_{\text{pre}(a_{n.o.t.})} = \sigma_{\text{pre}(a_{\text{orig}})} \text{ and } \sigma_{\text{eff}(a_{n.o.t.})}^{\text{True}} = \sigma_{\text{eff}(a_{\text{orig}})}^{\text{True}}\}$ // New skills based on $\mathcal{A}_{n.o.t.}$
- 4 $\mathcal{A}_{\text{candidate-skills-all}} := \{(\sigma_{\text{pre}(a)}, \sigma) \mid a \in \mathcal{A}, \sigma \in \Sigma^+ \text{ s.t. } \nexists a_{\text{orig}} \in \mathcal{A} \text{ where } \sigma_{\text{pre}(a)} = \sigma_{\text{pre}(a_{\text{orig}})} \text{ and } \sigma_{\text{eff}(a)}^{\text{True}} = \sigma_{\text{eff}(a_{\text{orig}})}^{\text{True}}\}$ // New skills based on \mathcal{A}
- 5 $\alpha_{\text{suggestions}} := \emptyset$
- 6 $\alpha_{\text{combinations}} = \text{combinations of } \mathcal{A}_{\text{candidate-skills-n.o.t.}} \text{ and } (n_{\text{new-skills}} - 1) \text{ skills from } \mathcal{A}_{\text{candidate-skills-all}}$
- 7 **for** $\mathcal{A}_{\text{new}} \in \alpha_{\text{combinations}}$ **do**
- 8 Write φ_{skills} with $\mathcal{A} \cup \mathcal{A}_{\text{new}}$
- 9 Synthesize φ_{full} with φ_{skills} and φ_{task}
- 10 **if** *Realizable* **then**
- 11 $\alpha_{\text{suggestions}} := \alpha_{\text{suggestions}} \cup \mathcal{A}_{\text{new}}$
- 12 **end**
- 13 **end**
- 14 **return** $\alpha_{\text{suggestions}}$

Algorithm 2: restrictPostconditions (additions to Pacheck et al. (2020) are in blue)

Input: Game structure G , Winning states Z , User-defined variables \mathcal{R} , **Extra skills** $\mathcal{A}_{\text{extra-skills}}$, **Extra skill that can be changed** $a_{\text{extra-skill-modify}}$

Output: Updated τ_e

- 1 $R_1 := \{(v_{\mathcal{AP}}, v'_e) \mid v_{\mathcal{AP}} \in \mathcal{V}_{\mathcal{AP}}, v_e \in \mathcal{V}_{\mathcal{E}}, \exists v_s \in \mathcal{V}_{\mathcal{S}} \text{ s.t. } (v_{\mathcal{AP}}, v'_e, v'_s) \in \tau_e \text{ and } (v_e, v_s) \in Z\}$;
- 2 $R_2 := \{(v_{\mathcal{AP}}, v'_e) \mid v_{\mathcal{AP}} \in \mathcal{V}_{\mathcal{AP}}, v_e \in \mathcal{V}_{\mathcal{E}} \text{ s.t. } (v_{\mathcal{AP}}, v'_e) \in \tau_e \text{ and } (v_{\mathcal{AP}}, v'_e) \in R_1\}$;
- 3 $R_2 := R_2 \setminus Z$;
- 4 $R_2 := \{(v_e, v_s) \in R_2 \mid \text{if } v_s \in \mathcal{A}_{\text{extra}}, v_s = a_{\text{extra-skill-modify}}\}$;
- 5 $\tau_e^{\text{new}} := \{(v_{\mathcal{AP}}, v'_e) \mid v_{\mathcal{AP}} \in \mathcal{V}_{\mathcal{AP}}, v_e \in \mathcal{V}_{\mathcal{E}} \text{ s.t. } v_{\mathcal{AP}} \in R_2 \text{ and } (v_{\mathcal{AP}}, v'_e) \in R_1 \cap \tau_e\}$;
- 6 $\tau_e^{\text{new-expanded}} = \{(v_{\mathcal{AP}\text{-expanded}}, v'_{e\text{-expanded}}) \mid \exists (v_{\mathcal{AP}}, v'_e) \in \tau_e^{\text{new}}, r_1, r_2 \in 2^{\mathcal{R}} \text{ s.t. } (v_{\mathcal{AP}\text{-expanded}} = v_{\mathcal{AP}} \setminus r_1 \text{ or } v_{\mathcal{AP}\text{-expanded}} = v_{\mathcal{AP}} \cup r_1) \text{ and } (v'_{e\text{-expanded}} = v'_e \setminus r_2 \text{ or } v'_{e\text{-expanded}} = v'_e \cup r_2)\}$;
- 7 $\tau_e^{\text{old}} := \{(v_{\mathcal{AP}}, v'_e) \mid v_{\mathcal{AP}} \in \mathcal{V}_{\mathcal{AP}}, v_e \in \mathcal{V}_{\mathcal{E}} \text{ s.t. } v_{\mathcal{AP}} \notin R_2 \text{ and } (v_{\mathcal{AP}}, v'_e) \in \tau_e\}$;
- 8 **return** $\tau_e^{\text{new-expanded}} \cup \tau_e^{\text{old}}$;

skills, thereby allowing the robot to use them in additional situations, or postconditions that should be removed from skills, essentially reducing nondeterminism. Additional preconditions are a set of symbol combinations $\sigma_{\text{pre}(a)}^{\text{added}}$ that should be added to the existing preconditions, that is, $\sigma_{\text{pre}(a)}^{\text{new}} = \sigma_{\text{pre}(a)} \cup \sigma_{\text{pre}(a)}^{\text{added}}$. The repair process can also remove one or more postconditions $j \in \{1, \dots, k(a)\}$.

To find new skills, we add a set of additional skills, $\mathcal{A}_{\text{extra-skills}}$, to the specification that are unrestricted—they can be executed from any combination of symbols and can result in any combination of symbols. The preconditions of $a_{\text{extra-skill}} \in \mathcal{A}_{\text{extra-skills}}$ are $\sigma_{\text{pre}(a)}[a_{\text{extra-skill}}] = 2^{\Sigma}$. In practice, we simply do not include constraints on the preconditions of $\mathcal{A}_{\text{extra-skills}}$ in $\varphi_{t, \text{skills}}^s$. Similarly, there are no constraints on the postconditions of $a_{\text{extra-skill}} \in \mathcal{A}_{\text{extra-skills}}$. The extra skills still need to satisfy constraints on the mutual exclusion of symbols in $\varphi_{t, \text{mx_syms}}^e$. New skills are of the form of a set of preconditions, $\sigma_{\text{pre}(a)}^{\text{new}}$ and collection of postconditions for that skill $\sigma_{\text{eff}(a)}^{\text{True-new}}$.

To find skill suggestions, the synthesis-based repair process performs synthesis until it determines the specification is unrealizable. The repair process then iteratively modifies τ_e and τ_s , which correspond to $\varphi_{t, \text{skills}}^e$ and $\varphi_{t, \text{skills}}^s$, respectively, and attempts to perform synthesis until the specification is realizable. The repair process returns τ_e^{new} and τ_s^{new} , from which we extract the skill suggestions.

To modify τ_e , the process of restricting postconditions takes the current game structure G , a set of winning states Z ,

the user-defined variables \mathcal{R} , the set of extra skills $\mathcal{A}_{\text{extra-skills}}$, and which extra skill can be changed on this iteration $a_{\text{extra-skill-modify}}$. The process of restricting the postconditions removes transitions from τ_e as shown in Algorithm 2. We start with the current set of winning states Z and attempt to expand it. In Line 1, we find R_1 , the set of states from which the system has the ability to reach Z . Then, in Line 1, we find R_2 , the set of states from which at least one next state will be in R_1 . This set includes states where one nondeterministic action outcome may reach R_1 (and can therefore reach Z) but another may not. We then remove already winning states from R_2 in Line 1. In Pacheck et al. (2020), we find τ_e^{new} by removing all the postconditions associated with R_2 that do not result in R_1 in Line 1 (skipping Line 1). We then return τ_e^{new} and τ_e^{old} (the transitions that were not modified).

In Pacheck et al. (2020), the repair process can remove any transition, while running restrictPostconditions, in order to attempt to repair the specification. In the case of reactive specifications, this can result in unwanted behavior. The process of restricting postconditions is able to remove transitions that correspond to changes in the values of $v_u \in \mathcal{R}$, as symbols in both \mathcal{R} and Σ are treated the same. As a result, the synthesis-based repair process proposed in (Pacheck et al., 2020) may offer suggestions for skills that can change (or keep the same) the value of symbols $v_u \in \mathcal{R}$. We show an example of this in the Baxter Plates example in Section 7.

In this work, we modify Algorithm 2 to repair reactive specifications. Modifications to the algorithm are shown in blue. We add Lines 1 and 1 to Algorithm 2. We do not allow the repair process to change the value of symbols $v_u \in \mathcal{R}$. The modifications in τ_e^{new} found in Line 1 may include restrictions on the truth value of $v_u \in \mathcal{R}$. Since we do not want to allow such modifications, for each possible modification in τ_e^{new} , we add transitions corresponding to all possible changes in the truth value of $v_u \in \mathcal{R}$ in Line 1. Now, the environment is able to make transitions with no restriction on $v_u \in \mathcal{R}$, while still having restrictions on $\sigma \in \Sigma$. The addition of Line 1 does not allow us to find any suggestions for some specifications. As we show in Section 7, for some specifications, the only way for the synthesis-based repair process to find suggestions is by changing the values of variables $v_u \in \mathcal{R}$, essentially trying to enforce the behavior of external events, which is not desired.

For some specifications, multiple skills are required to find suggestions. To account for this, we only allow the repair process to change one additional skill $a_{\text{extra-skill-modify}} \in \mathcal{A}_{\text{extra-skills}}$ per iteration in (Line 1). Without Line 1, the repair process attempts to change all of the extra skills to the same postconditions at once, which can cause the repair process to fail.

Additionally, we only perform one iteration of the while loop in Algorithm 2: Repair in Pacheck et al. (2020). In the previous work, we modified τ_s and τ_e , then applied the controllable predecessor operator to Z with the new τ_s and τ_e until the current liveness guarantee overlapped with Z . However, because we are making modifications but not allowing the user-defined variables to be restricted, the process of restricting the postconditions may expand Z in ways not captured by only the application of the controllable predecessor. The new skills may allow for the system to have more control over the outcome of skills at other states not currently in Z . By performing synthesis after only one iteration of the repair process, we are able to find suggestions for more specifications.

We find τ_s^{new} by relaxing the preconditions as in Pacheck et al. (2020). The process is similar to restricting the postconditions. We find all the states from which the system can reach Z without violating τ_s^{hard} . We then find the states that will always lead to these states from which the system can reach Z . We add these states τ_s and thereby expand the preconditions.

After finding τ_e^{new} and τ_s^{new} that allow the specification to be synthesized, the repair process finds a strategy for the system to achieve the liveness guarantees from the initial conditions. We then compare the preconditions and postconditions of the skills performed during the strategy to the preconditions and postconditions of the skills initially given to the robot. The new postconditions for the extra skills are then those seen in the strategy. The new preconditions for the extra skills are those seen in the strategy. Similarly, the additional preconditions for extra skills are those seen in the strategy (Pacheck et al., 2020).

To find multiple suggestions, once one suggestion has been found by the synthesis-based repair process, the new skills making up the suggestion are automatically disallowed along with any additional preconditions without any user input. A new strategy is found if possible and another suggestion extracted. We continue finding additional suggestions and disallowing previous suggestions until there are no more suggestions (Pacheck et al., 2020). All suggestions found by the synthesis-based repair process are then returned to the user who can choose which suggestion to implement.

7. Robot demonstrations

We demonstrate automatically creating φ_{skills} , writing and executing task specifications, and the repair process with examples involving a Kinova robot manipulating vials and a Baxter robot manipulating blocks and pushing plates. The Baxter Blocks example shows skills with nondeterminism and several unrealizable specifications. The Baxter Plates example shows the use of raw camera images to create symbols and the benefit of enumeration-based repair over synthesis-based repair. The Kinova Vials example shows the benefits of the synthesis-based repair over the enumeration-based repair.

7.1. Environment setup

We have three different demonstration environments.

Baxter Blocks: In the Baxter Blocks example, a Baxter robot is manipulating blocks on a table as shown in Figure 3(a) and (b). There are three blocks: red, blue and green (also labeled 1, 2, and 3, respectively). The red and blue blocks can be placed at locations A, D and E, and can also be stacked at AT. The green block can be placed at locations B, C, F, and G. The location of the blocks are determined by AprilTags (Wang and Olson 2016) detected by the Baxter’s wrist cameras. The state space is the x , y , and z position of each of the blocks.

Baxter Plates: The Baxter Plates example contains a Baxter robot manipulating plates on a table as shown in Figure 3(c) and (d). There are two plates (a blue square plate and a green oval plate) that can be moved between the clean, set, and dirty regions. The state in this example is the image recorded by the USB camera shown in the right of Figure 3(d). Symbols are generated using a lower dimensional state space found using independent component analysis (Hyvärinen and Oja 2000) as detailed in Section 7.3.

Kinova Vials: The Kinova Vials example deals with a Kinova arm moving three colored vials (green, red and yellow) as shown in Figure 3(e) and (f). The vials can be in six regions: top-left, top-right, rack-left, rack-right, right-top, and right-bottom. There can only be one vial in each region at a time. The position of the vials is determined by a motion capture system using structures placed on top of the

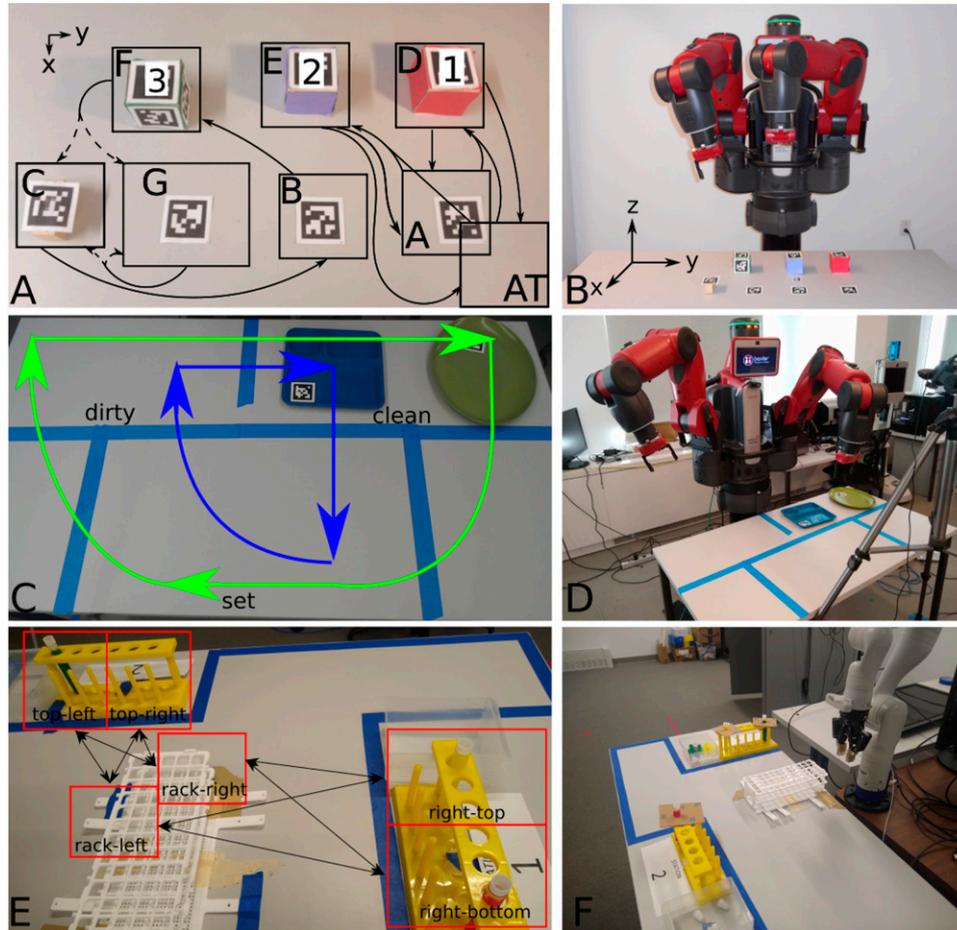


Figure 3. Baxter Blocks: (a) The arrows show the skills given to the Baxter robot to move blocks between lettered locations. Dashed arrows represent skills with nondeterministic outcomes. Note that one skill takes blocks from A and AT to D and another from A and AT to E while all the other skills have preconditions that are only one location. (b) Initial setup of the Baxter Blocks example. Baxter Plates: (c) The colored arrows represent the skills given to the Baxter robot to move the two plates between the clean, set, and dirty locations. The skill moving the green (oval) plate from clean to set is not possible when the blue (square) plate is set and the skill moving the green plate from dirty to clean is not possible when the blue plate is dirty or clean. The skill moving the blue plate from set to dirty is not possible when the green plate is set. (d) Initial setup of the Baxter Plates example. Kinova Vials: (e) The arrows represent the skills given to the Kinova robot to move the vials. The skills move the green, red, and yellow vials between the top-left, top-right, rack-left, rack-right, right-top, and right-bottom locations. The skills move the vials between the yellow (top and right) and white (rack) holders (and vice-versa) but not between the two yellow holders. (f) Initial setup of the Kinova Vials example. (a) and (b) are from Pacheck et al. (2019).

vials as shown in Figure 3(f). The state is the x and y location of all of the vials.

7.2. Skills

The robots are given a set of executable skills.

Baxter Blocks: The skills for the Baxter Blocks example are implemented as controllers that move the arm of the Baxter over the position of the block as determined by the AprilTags (Wang and Olson 2016) attached to each block. The skill then lifts up the block and moves it to over the destination location. Finally, the skill lowers the block and releases it.

The left arm of the Baxter moves the red and blue blocks while the right arm moves the green block. The skills allow the robot to move the red and blue blocks

from D and E to A and AT (and vice-versa). Another skill attempts to move the green block from F to C, but due to C being elevated, the skill sometimes results in the block ending in G. Similarly, the skill from G to C sometimes results in the green block ending in C and sometimes in G. Finally, there are skills that move the green block reliably from C to B and B to F. Figure 3(a) shows the skills available to the Baxter where skills with dashed lines have nondeterministic outcomes.

Baxter Plates: In the Baxter plates example, there are six skills that involve reproducing a trajectory demonstrated to the robot to move the green and blue plates from clean to set, set to dirty, and dirty to clean. Skills require the plate to be in the named initial location in order to be executed (e.g., the blue plate must be clean in order for $a_{blue-clean-to-set}$ to be executed). Some skills have additional restrictions on when

they can be executed—the skill $a_{green-clean-to-set}$ cannot be executed when the blue plate is in set, $a_{green-dirty-to-clean}$ cannot be executed when the blue plate is dirty or clean, and $a_{blue-set-to-dirty}$ cannot be executed when the green plate is in set. Figure 3(c) shows the skills available to the Baxter and the approximate paths the plates follow between locations.

Kinova Vials: The skills in the Kinova Vials example allow the robot to move the vials between the yellow outer racks and the white rack (and vice-versa), but not between the two yellow racks. For each skill, the arm moves between predetermined waypoints to move above the initial location, move down, grasp the vial, move up, move to above the destination location, move down, release the vial into the rack, move up, and then move back to a home position shown in Figure 3(f). Figure 3(e) shows the skills available to the Kinova arm; note that all the arrows are bi-directional.

7.3. Symbol generation

We collected data to automatically generate symbols and encode the skills into an LTL formula for each example.

During the data collection process, the robot randomly executes a valid skill. A user-defined oracle tells the robot which skills can be executed based on the locations of the blocks, plates, or vials, and the robot randomly executes one of those skills. In the Baxter Blocks and Plates examples, the oracle determines which skills can be executed based on the AprilTags (Wang and Olson 2016) data, while the Kinova Vials example uses motion capture data. The oracle is created by the user when the skills are initially given to the robot and is only used during data collection for symbol generation. For example, in the Baxter Plates example, when the AprilTag on the blue plate is greater than 0.68 m from the center of the Baxter (in the set region) and the AprilTag on the green plate is less than 0.4 m from the Baxter (in either the dirty or clean region), the oracle determines the skill $a_{blue-set-to-dirty}$ can be executed. Note that while the oracle in the Baxter Plates example uses AprilTags (Wang and Olson 2016) to determine which skills can be executed, the symbol generation process uses only raw camera images. For each example, we collected multiple precondition-skill-postcondition sets of data, along with which actions could be executed at each precondition. Table 1 lists the number of skills and data collected for each of the three examples.

Baxter Blocks: After collecting the data, we generated symbols and partitioned the skills for each of the examples (Konidaris et al., 2018). As shown in Table 1, we generated 19 symbols and 20 partitioned skills for the Baxter Blocks example. The symbol generation process partitioned the skills that could move either the blue or red block from E or D to A or AT (and vice-versa) into two

skills each. The symbols generated refer to the x , y , or z position for each of the blocks. Figure 4(a) shows the approximate grounding of the propositions, grouped by the block they correspond to.

Baxter Plates: For the Baxter Plates example, we generated symbols directly from images taken by an external stationary USB camera, as shown in Figure 3(d). To generate the symbols, we first resize the images to 120×72 , convert them to grayscale and then apply independent component analysis (Hyvärinen and Oja 2000), keeping the top 5 components. The symbol generation process is then applied to these lower dimensional vectors—preconditions are estimated using a support vector machine ($C = 2$, $\gamma = 4$) (Cortes and Vapnik 1995), while effects are modeled using a kernel density estimator (Rosenblatt 1956; Parzen 1962) with a Gaussian kernel and bandwidth determined by 3-fold cross validation. This procedure generates 5 factors and 9 symbols, where each symbol is a subset of the low-dimensional representation of an image. Figure 4(b) shows each of the symbols.

Kinova Vials: In the Kinova Vials example, we generated 18 symbols and 48 partitioned skills. Each symbol corresponds to a vial being located in a different region. Figure 4(c) shows the approximate grounding of each of the symbols, grouped by the color vial they correspond to. Due to the nature of the state space and skills, the symbol generation process factored the state space such that each symbol is over both the x and y position of a vial, as opposed to only the x or y position, as in the Baxter Blocks example. Each skill in the Kinova Vials example changes both the x and y position of a vial; in the Baxter Blocks example the x , y , and z position of blocks do not always change together.

7.4. Skills-based specification

We automatically encode the symbols and skills in φ_{skills} for each example. We show selected parts from the specification for each example.

Baxter Blocks: In the Baxter Blocks example, we automatically encoded the symbols $\{\sigma_0, \dots, \sigma_{18}\} \in \Sigma$ and skills $\{a_{f-to-c}, \dots, a_{d-to-at-2}\} \in \mathcal{A}$ in φ_{skills} . In equation (7), we show part of the system safety formula $\varphi_{t,skills}^s$. We show part of the environment safety formula $\varphi_{t,eff}^e$ in equations (8) and (9). Figure 5(a)–(d) visualizes the result of applying skills $a_{e-to-a-2}$ and a_{f-to-c} .

The precondition requirements of a_{c-to-b} are encoded in $\varphi_{t,skills}^s$ in equation (7).

$$\square(\neg\bigcirc\sigma_1 \rightarrow \neg\bigcirc a_{c-to-b}) \quad (7)$$

$$\square(a_{e-to-a-2} \rightarrow (\bigcirc(\sigma_3 \wedge \sigma_4) \wedge \bigcirc(\neg\sigma_{11} \wedge \neg\sigma_{13}) \wedge (\sigma \leftrightarrow \bigcirc\sigma))) \quad (8)$$

$\sigma \in \sigma_{\text{eff}}^{\text{stay}}(a_{e-to-a-2})$

Table 1. Overview of the different robot demonstration environments. For each demonstration environment, we gave the robot a set of skills, collected data on the preconditions and postconditions, generated symbols, and automatically encoded the preconditions and postconditions in an LTL formula. Partitioned skills are those that have different effects from different preconditions. The number of formulas in φ_{skills} include constraints on the preconditions, postconditions, mutual exclusion of skills, and mutual exclusion of symbols.

	Number of given skills	Number of partitioned skills	Number of symbols generated ($ \Sigma $)	Number of pre/post pairs collected ($ D $)	Number of formulas in φ_{skills}
Baxter Blocks	9	20	19	1052	83
Baxter Plates	6	7	9	100	35
Kinova Vials	48	48	18	860	195

$$\begin{aligned} & \Box(a_{f-to-c} \rightarrow ((\bigcirc\sigma_0 \wedge \bigcirc\neg\sigma_{17} \wedge_{\substack{\sigma \in \sigma_{\text{stay}} \\ \text{eff}^1(a_{f-to-c})}} (\sigma \leftrightarrow \bigcirc\sigma))) \vee \\ & (\bigcirc(\sigma_0 \wedge \sigma_1 \wedge \sigma_2) \wedge \bigcirc(\neg\sigma_{15} \wedge \neg\sigma_{17} \wedge \neg\sigma_{16} \wedge \neg\sigma_{18})) \end{aligned} \quad (9)$$

$$\begin{aligned} & \wedge_{\substack{\sigma \in \sigma_{\text{stay}} \\ \text{eff}^2(a_{f-to-c})}} (\sigma \leftrightarrow \bigcirc\sigma))) \\ & \Box(\neg\bigcirc(\sigma_6 \wedge \sigma_3) \rightarrow \neg\bigcirc a_{\text{blue-clean-to-set}}) \end{aligned} \quad (10)$$

$$\begin{aligned} & \Box(a_{\text{blue-clean-to-set}} \rightarrow (\bigcirc\sigma_5 \wedge \bigcirc\neg\sigma_3 \wedge_{\substack{\sigma \in \sigma_{\text{stay}} \\ \text{eff}^1(a_{\text{blue-clean-to-set}})}} (\sigma \leftrightarrow \bigcirc\sigma))) \\ & \end{aligned} \quad (11)$$

Based on the data the robot has seen, it determines that it only needs to consider the value of y_3 in deciding if skill

This corresponds to block 2 moving from location E to A and blocks 1 and 3 not moving. A potential outcome of applying skill $a_{e-to-a-2}$ is visualized in Figure 5(b).

The part of $\varphi_{\text{t,skills}}^e$ referring to the nondeterministic effects of skill a_{f-to-c} is shown in equation (9) where $\sigma_{\text{eff}^1(a_{f-to-c})}^{\text{stay}} = \{\sigma_1, \dots, \sigma_{16}, \sigma_{18}\}$ and $\sigma_{\text{eff}^2(a_{f-to-c})}^{\text{stay}} = \{\sigma_3, \dots, \sigma_{14}\}$. Equation (9) encodes that when skill a_{f-to-c} is applied, either σ_0 becomes True and σ_{17} becomes False with symbols in $\sigma_{\text{eff}^1(a_{f-to-c})}^{\text{stay}}$ not changing (block 3 ends in G), or $\sigma_0, \sigma_1,$ and σ_2 become True and $\sigma_{15}, \sigma_{17}, \sigma_{16},$ and σ_{18} become False with symbols in $\sigma_{\text{eff}^2(a_{f-to-c})}^{\text{stay}}$ not changing (block 3 ends in C). This is visualized in Figure 5(c) and (d).

Baxter Plates: In the Baxter Plates example, the symbols correspond to images and the combination of symbols can be visualized together. For example, Figure 6 illustrates the preconditions and effects when the agent executes $a_{\text{blue-clean-to-set}}$. The precondition requirements of

$$\begin{aligned} & \Box(\neg\bigcirc((\sigma_7 \wedge \sigma_0 \wedge \sigma_{13}) \vee (\sigma_7 \wedge \sigma_0 \wedge \sigma_{14}) \vee (\sigma_7 \wedge \sigma_0 \wedge \sigma_{15}) \vee (\sigma_7 \wedge \sigma_0 \wedge \sigma_{17}) \vee \\ & (\sigma_8 \wedge \sigma_0 \wedge \sigma_{13}) \vee (\sigma_8 \wedge \sigma_0 \wedge \sigma_{15}) \vee (\sigma_9 \wedge \sigma_0 \wedge \sigma_{13}) \vee (\sigma_9 \wedge \sigma_0 \wedge \sigma_{14}) \vee \\ & (\sigma_9 \wedge \sigma_0 \wedge \sigma_{15}) \vee (\sigma_9 \wedge \sigma_0 \wedge \sigma_{17}) \vee (\sigma_{11} \wedge \sigma_0 \wedge \sigma_{14}) \vee (\sigma_{11} \wedge \sigma_0 \wedge \sigma_{15})) \rightarrow \\ & \neg\bigcirc a_{\text{green-right-bottom-to-rack-left}}) \end{aligned} \quad (12)$$

$$\begin{aligned} & \Box(a_{\text{green-right-bottom-to-rack-left}} \rightarrow \\ & (\bigcirc\sigma_4 \wedge \bigcirc(\neg\sigma_0 \wedge \neg\sigma_1 \wedge \neg\sigma_2 \wedge \neg\sigma_3 \wedge \neg\sigma_5) \wedge_{\substack{\sigma \in \sigma_{\text{stay}} \\ \text{eff}^1(a_{\text{green-right-bottom-to-rack-left}})}} (\sigma \leftrightarrow \bigcirc\sigma))) \end{aligned} \quad (13)$$

a_{c-to-b} can be performed. There is only one symbol falling inside the precondition set so $\sigma_{\text{pre}(a)}[a_{c-to-b}] = \{\{\sigma_1\}\}$. Therefore, equation (7) states that if σ_1 is not True, that is, block 3 is not at approximately $y = -0.3$ m, skill a_{c-to-b} cannot be applied.

The part of $\varphi_{\text{t,skills}}^e$ pertaining to the effect of skill $a_{e-to-a-2}$ is shown in equation (8) where $\sigma_{\text{eff}^1(a_{e-to-a-2})}^{\text{stay}} = \{\sigma_0, \sigma_1, \sigma_2, \sigma_5, \dots, \sigma_{10}, \sigma_{12}, \sigma_{14}, \dots, \sigma_{18}\}$.

$a_{\text{blue-clean-to-set}}$ are encoded in $\varphi_{\text{t,skills}}^s$ in equation (10). The postconditions of $a_{\text{blue-clean-to-set}}$ are encoded in $\varphi_{\text{t,skills}}^e$ as shown in equation (11) where $\sigma_{\text{eff}^1(a_{\text{blue-clean-to-set}})}^{\text{stay}} = \{\sigma_0, \sigma_1, \sigma_2, \sigma_4, \sigma_6, \sigma_7, \sigma_8\}$.

Kinova Vials: For the Kinova Vials example, we encode the preconditions and postconditions of the skills to move the vials in φ_{skills} . We show part of the specification involving $a_{\text{green-right-bottom-to-rack-left}}$. The preconditions of

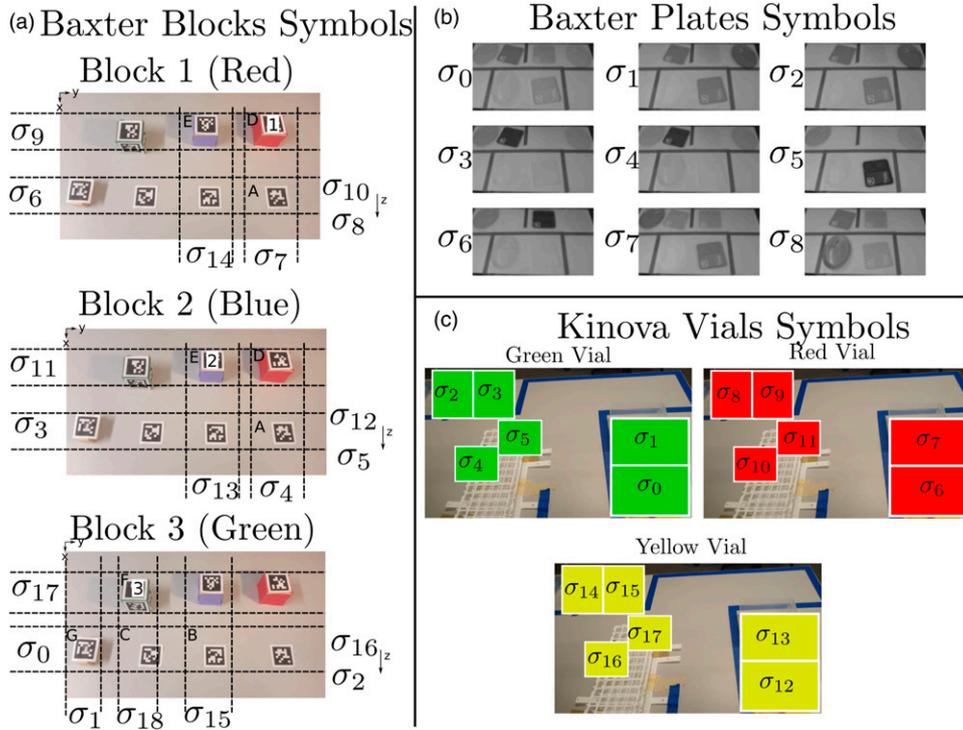


Figure 4. The approximate grounding of the symbols for the Baxter Blocks, Baxter Plates, and Kinova Vials examples. (a) For the Baxter Blocks, each symbols grounds to one dimension of the state space. (b) For the Baxter Plates, each symbol corresponds to a lower dimensional representation of the raw image. (c) For the Kinova Vials, each symbol grounds to two dimensions of the state space.

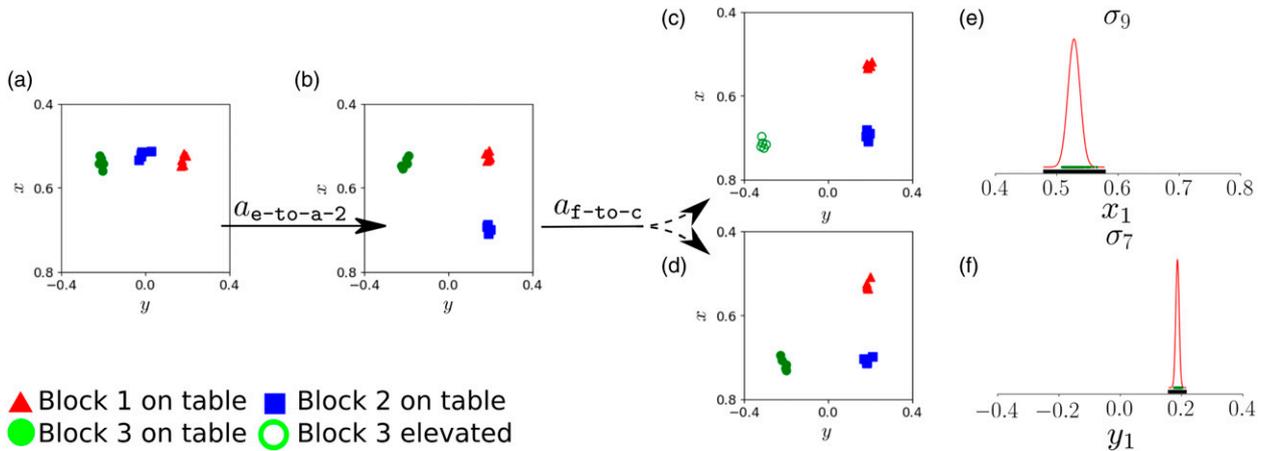


Figure 5. Visualization of symbol combinations (a) $\sigma_9 \wedge \sigma_7 \wedge \sigma_{10} \wedge \sigma_{11} \wedge \sigma_{13} \wedge \sigma_{12} \wedge \sigma_{17} \wedge \sigma_{18} \wedge \sigma_{16}$, (b) $\sigma_9 \wedge \sigma_7 \wedge \sigma_{10} \wedge \sigma_3 \wedge \sigma_4 \wedge \sigma_{12} \wedge \sigma_{17} \wedge \sigma_{18} \wedge \sigma_{16}$, (c) $\sigma_9 \wedge \sigma_7 \wedge \sigma_{10} \wedge \sigma_3 \wedge \sigma_4 \wedge \sigma_{12} \wedge \sigma_0 \wedge \sigma_1 \wedge \sigma_2$, (d) $\sigma_9 \wedge \sigma_7 \wedge \sigma_{10} \wedge \sigma_3 \wedge \sigma_4 \wedge \sigma_{12} \wedge \sigma_0 \wedge \sigma_{18} \wedge \sigma_{16}$. All other symbols were False. Ten samples were drawn from the intersection of the grounding sets of each symbol combination. Possible transitions are shown between the subfigures, corresponding to transitions in equations (8) and (9). Applying skill $a_{e-to-a-2}$ to (a) results in (b). Applying skill a_{f-to-c} in (b) results in (c) or (d). Examples of symbol groundings are shown in (e) and (f) as black lines. The raw data is shown in green and the Gaussian fit to it in red. Figure adapted from Pacheck et al. (2019).

$a_{green-right-bottom-to-rack-left}$ are combinations of different locations of the red and yellow vials when the green vial is in right-bottom and the rack-left location is clear as shown in equation (12). For all of the Kinova Vials skills, the skills could have up to 16 preconditions, as the precondition classifier does not necessarily learn that the vials not being moved in the skill cannot be in the same physical location. In this example, there are only 12

preconditions, with one ($\sigma_7 \wedge \sigma_0 \wedge \sigma_{13}$) having both the red and yellow vials at the right-top location.

The postconditions of $a_{green-right-bottom-to-rack-left}$ are that the green vial is in the rack-left location and not in the right-bottom location while the position of the red and yellow vials are unchanged as shown in equation (13) where

$$\sigma_{\text{eff}}^{\text{stay}}(a_{green-right-bottom-to-rack-left}) = \{\sigma_6, \sigma_7, \dots, \sigma_{16}, \sigma_{17}\}.$$

	Symbolic Expression	Image Representation				
Preconditions	$\sigma_6 \wedge \sigma_3$		\wedge		$=$	
Positive Effect	σ_5					
Negative Effect	σ_3					
Postcondition	$\sigma_6 \wedge \sigma_5$		\wedge		$=$	

Figure 6. Symbolic representation of the precondition and effect for $a_{\text{blue-clean-to-set}}$. The outcome of the action is computed by adding the positive effect to the precondition, and then removing the negative one. Since our representation is factorized, the preconditions and effects depend only on a subset of the factors that constitute the symbolic state space. In this case, σ_3 and σ_5 both refer to the same factor, while σ_6 refers to a different factor. The location of the plates is entangled in the individual symbols, so we need to view combinations of symbols to know the location of the plates. In general, the darker a plate is at a location, the more it needs to be there for the precondition or result. The lighter gray versions of plates represent the skill not being dependent on that plate location or not affecting that plate. The precondition for the skill represents states where the blue plate is in the clean position and the green plate is elsewhere. This is shown by the intersection of σ_6 and σ_3 showing the blue square plate located in the top-right (clean) position being very dark. The green oval plate can be located in any of its three positions as it shows up in all of those positions. The effect of the skill is that the blue plate is now in the clean position; the green plate remains unaffected. This is shown by the location of the blue square plate in the bottom right (set) position being very dark and the green oval plate being very light in all locations. Note that the raw images were greyscaled during the symbol generation process.

7.5. Realizable base task specifications

We introduce additional environment variables $v_u \in \mathcal{R}$ that the user controls. Using \mathcal{R} , we write reactive tasks. We describe the user-defined task for each example and show parts of the specification in this section. We show the full φ_{skills} for each example in Appendix A.

Baxter Blocks: For the Baxter Blocks example, we introduce an additional environment variable $\mathcal{R} = \{\text{Switch}\}$. The task liveness specifications are shown in Figure 7. When $\text{Switch} = \text{True}$, the red block (block 1) should eventually be in A, the blue block (block 2) in AT, and the green block (block 3) in F as shown in Figure 7(a)–(c) and encoded in the LTL formula above the subfigures. When $\text{Switch} = \text{False}$, the red block (block 1) should be in AT, the blue block (block 2) in A, and the green block (block 3) in C as shown in Figure 7(d)–(f) and encoded in the LTL formula above the subfigures.

We include a fairness assumption on the environment that the green block (block 3) will eventually be placed in location C when $a_{g\text{-to-c}}$ is applied, as shown in equation (14). Without this, the specification is unrealizable because in the worst case, skill $a_{g\text{-to-c}}$ always results in the green block (block 3) ending in G.

$$\varphi_{g,\text{task}}^c = \square\Diamond(a_{g\text{-to-c}} \rightarrow (\sigma_0 \wedge \sigma_1 \wedge \sigma_2)) \quad (14)$$

Baxter Plates: For the Baxter Plates example, we introduce two additional user-defined variables $\mathcal{R} = \{\text{BluePerson}, \text{GreenPerson}\}$. The task liveness specification is to *set* the blue plate when $\text{BluePerson} = \text{True}$, *set*

the green plate when $\text{GreenPerson} = \text{True}$, make sure the blue plate is not in set when $\text{BluePerson} = \text{False}$, and the green plate is not in set when $\text{GreenPerson} = \text{False}$, as shown and encoded in Figure 8(a)–(d), respectively. Note that in the Baxter Plates example, we need to specify the location of both plates in the liveness guarantees due to the entanglement of the symbols.

The symbol generation process was not able to fully determine the effects of two of the skills, due to the lossy nature of the compressed state representation. The symbol generation process learns that for the skills $a_{\text{green-clean-to-set}}$ and $a_{\text{blue-clean-to-set}}$, a possible outcome is that no plates move. We add a fairness assumption that the skills should always eventually succeed in moving the plates shown in equation (15).

$$\begin{aligned} &\square\Diamond(a_{\text{green-clean-to-set}} \rightarrow \sigma_8) \wedge \\ &\square\Diamond(a_{\text{blue-clean-to-set}} \rightarrow \sigma_5) \end{aligned} \quad (15)$$

Kinova Vials: For the Kinova Vials example, we introduce the additional environment variable $\mathcal{R} = \{\text{React}\}$. The task is to arrange the vials in one configuration when $\text{React} = \text{True}$ and another when $\text{React} = \text{False}$. When $\text{React} = \text{True}$, the Kinova should arrange the vials such that the green vial is in the right-bottom position, the red vial is in the top-left position, and the yellow vial is in the top-right position (Figure 9(a) and (b)). When $\text{React} = \text{False}$, the Kinova should arrange the vials such that the green vial is in the top-left position, the red vial is in the right-bottom position, and the yellow vial is in the right-top position. The LTL formula encoding the

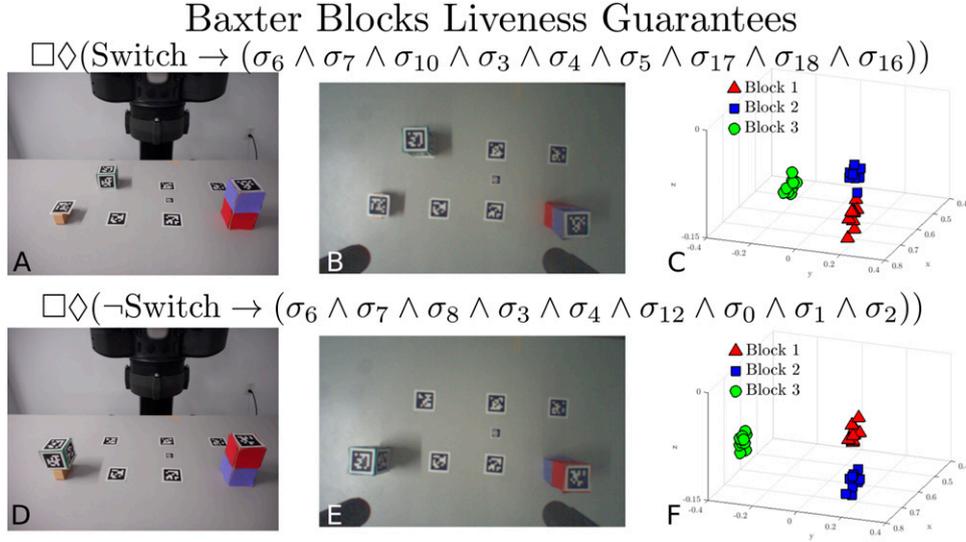


Figure 7. In the Baxter Blocks example, the robot reacts to the value of $\text{Switch} \in \mathcal{R}$. When $\text{Switch} = \text{True}$, the robot needs to stack the blue block on the red block in location A and the green block in location F. When $\text{Switch} = \text{False}$, the robot needs to stack the red block on the blue block in location A and the green block in location C. Figure from [Pacheck et al. \(2019\)](#).

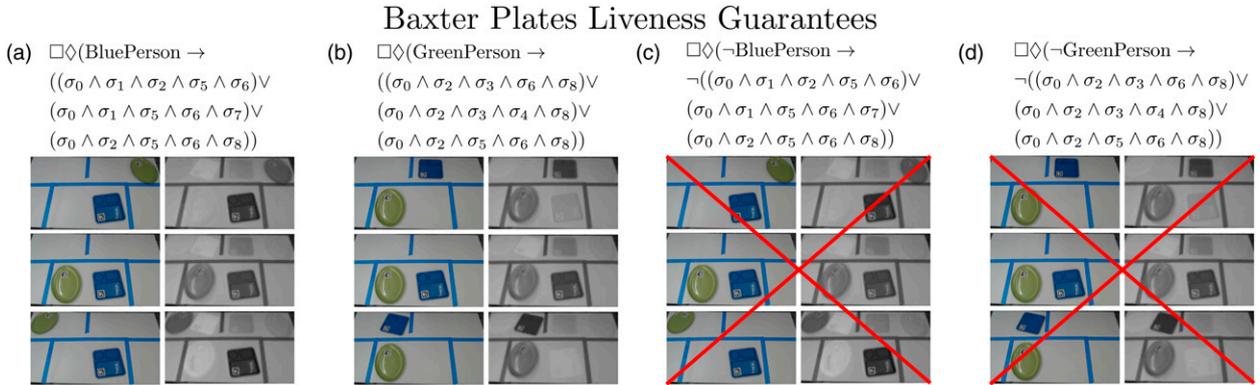


Figure 8. In the Baxter Plates example, (a) when $\text{BluePerson} = \text{True}$ the blue plate should be set, (b) when $\text{GreenPerson} = \text{True}$ the green plate should be set, (c) when $\text{BluePerson} = \text{False}$ the blue plate should not be set, and (d) when $\text{GreenPerson} = \text{False}$ the green plate should not be set. The left column of each subfigure shows the physical interpretation of the liveness guarantee and the right column shows the corresponding image generated based on the symbols.

task liveness specification is shown above the visual interpretation of the task in [Figure 9](#).

We add an additional constraint to $\varphi_{t, \text{task}}^s$ that the red and green vials should never be in the same yellow rack as shown in equation (16) and [Figure 16\(a\), \(b\), \(d\) and \(e\)](#).

$$\begin{aligned} \varphi_{t, \text{task}}^s = & \Box\neg(\sigma_0 \wedge \sigma_7) \wedge \Box\neg(\sigma_1 \wedge \sigma_6) \wedge \\ & \Box\neg(\sigma_2 \wedge \sigma_9) \wedge \Box\neg(\sigma_3 \wedge \sigma_8) \wedge \\ & \Box\neg\bigcirc(\sigma_0 \wedge \sigma_7) \wedge \Box\neg\bigcirc(\sigma_1 \wedge \sigma_6) \wedge \\ & \Box\neg\bigcirc(\sigma_2 \wedge \sigma_9) \wedge \Box\neg\bigcirc(\sigma_3 \wedge \sigma_8) \end{aligned} \quad (16)$$

7.6. Synthesis and execution

For each example, we are able to find a strategy to fulfill the original task. Throughout this paper, computation

times refer to running Slugs ([Ehlers and Raman 2016](#)) and our algorithms on an Ubuntu 18.04 machine with 12 GB RAM. For the Baxter Blocks example, we synthesize C with 256 states in 1 s. For the Baxter Plates example, it took 1 s to synthesize C with 124 states. The strategy for the Kinova Vials example took 22 s to synthesize and had 232 states.

We demonstrate the strategy for the Baxter Blocks example. We controlled the value of $\text{Switch} \in \mathcal{R}$ through a user interface. We sampled the current state $x \in X$ to find out which symbols were True. A symbol σ_{a, j, f_q} was True if the state was in the grounding set for the symbol, $\mathcal{G}(\sigma_{a, j, f_q})$. All other symbols were False. We show an example execution of C for the Baxter Blocks example in [Figure 10](#). We show an example sequence of states in C for the Baxter Plates example in [Figure 11](#).

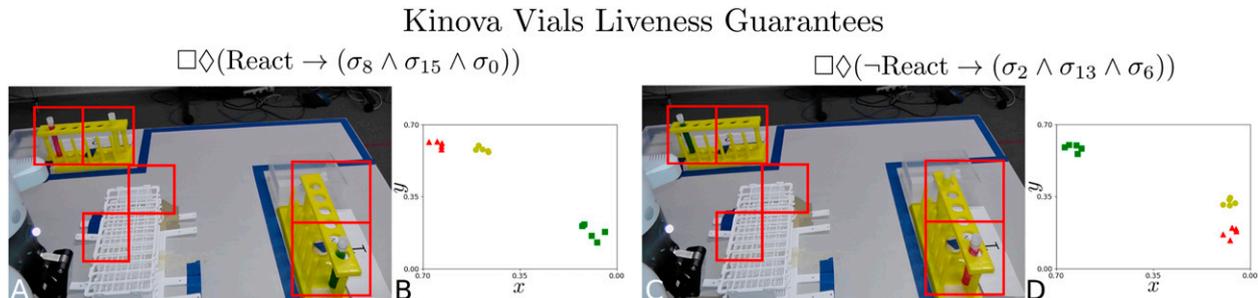


Figure 9. In the Kinova Vials example, the robot reacts the truth value of $\text{React} \in \mathcal{R}$. When $\text{React} = \text{True}$, the red vial should be in the top-left location, the yellow vial in the top-right location, and the green vial in the right-bottom location. When $\text{React} = \text{False}$, the red vial should be in the right-bottom location, the yellow vial in the right-top location, and the green vial in the top-left location.

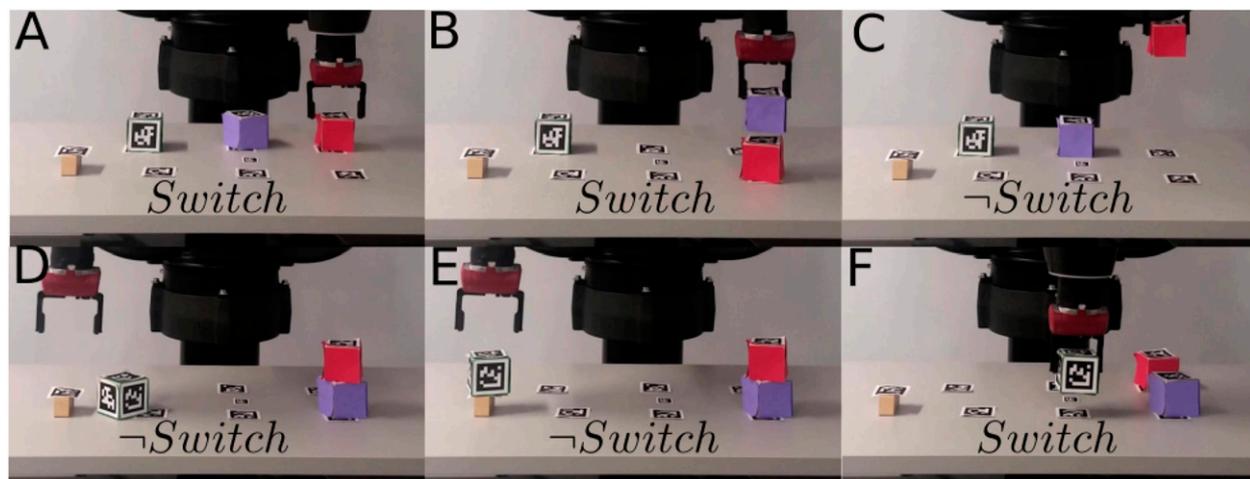


Figure 10. The Baxter robot executing a strategy to fulfill the desired liveness guarantees in Figure 7 without any additional user provided constraints. The value of Switch was controlled through a user interface. Time progresses from A to F. Figure adapted from Pacheck et al. (2019).

7.7. Repair of unrealizable tasks

We demonstrate the repair process by finding skill suggestions for six unrealizable specifications. For each example, we made the specification unrealizable by either adding additional task constraints to $\varphi_{t, \text{task}}^s$ or modifying $\varphi_{t, \text{skills}}^s$. Each unrealizable specification shows different aspects of the repair process.

For the Baxter Blocks example, we investigate four unrealizable specifications. In two specifications, we add constraints to avoid a skill or set of states. In these two specifications, we find $\mathcal{A}_{n.o.t} \neq \emptyset$, allowing us to narrow the search space for new skills to those with the same preconditions as $a \in \mathcal{A}_{n.o.t}$. For two other specifications, we modify the skills available to the robot and find $\mathcal{A}_{n.o.t} = \emptyset$, requiring us to perform an exhaustive search for new skills over all current preconditions sets.

For the Baxter Plates and Kinova Vials examples, we investigate one unrealizable specification each. In the Baxter Plates example, we add a reactive task constraint and show the benefits of the enumeration-based repair approach over the synthesis-based approach. In the Kinova Vials example, we add a constraint to avoid certain states and

show the benefits of the synthesis-based repair approach over the enumeration-based approach. For these unrealizable specifications, we find $\mathcal{A}_{n.o.t} \neq \emptyset$. The Baxter Plates and Kinova Vials examples both required two skills to repair, which increased the complexity of the repair process.

We describe the constraints added to each example that make the specifications unrealizable in this section and show selected formula. We show the full unrealizable task specifications, φ_{skills} , in Appendix A.

We give an overview of the number of solutions and time taken to find those solutions in Table 2.

Baxter Blocks: The two Baxter Blocks unrealizable specifications for which we can narrow the search space of possible skills have the same φ_{skills} and φ_{task} as in Sections 7.4 and 7.5, with the addition of $\varphi_{t, \text{task}}^s$, as shown in Figure 12 and described below. For each specification, we add the constraint that only one block can move at a time as shown in Appendix A.1.

Unrealizable Specification 1: In Figure 12(a), we show the added constraint that the green block (block 3) never be in location B, $\varphi_{t, \text{task}}^s = \Box\neg(\sigma_0 \wedge \sigma_{15}) \wedge \Box\neg\bigcirc(\sigma_0 \wedge \sigma_{15})$. This

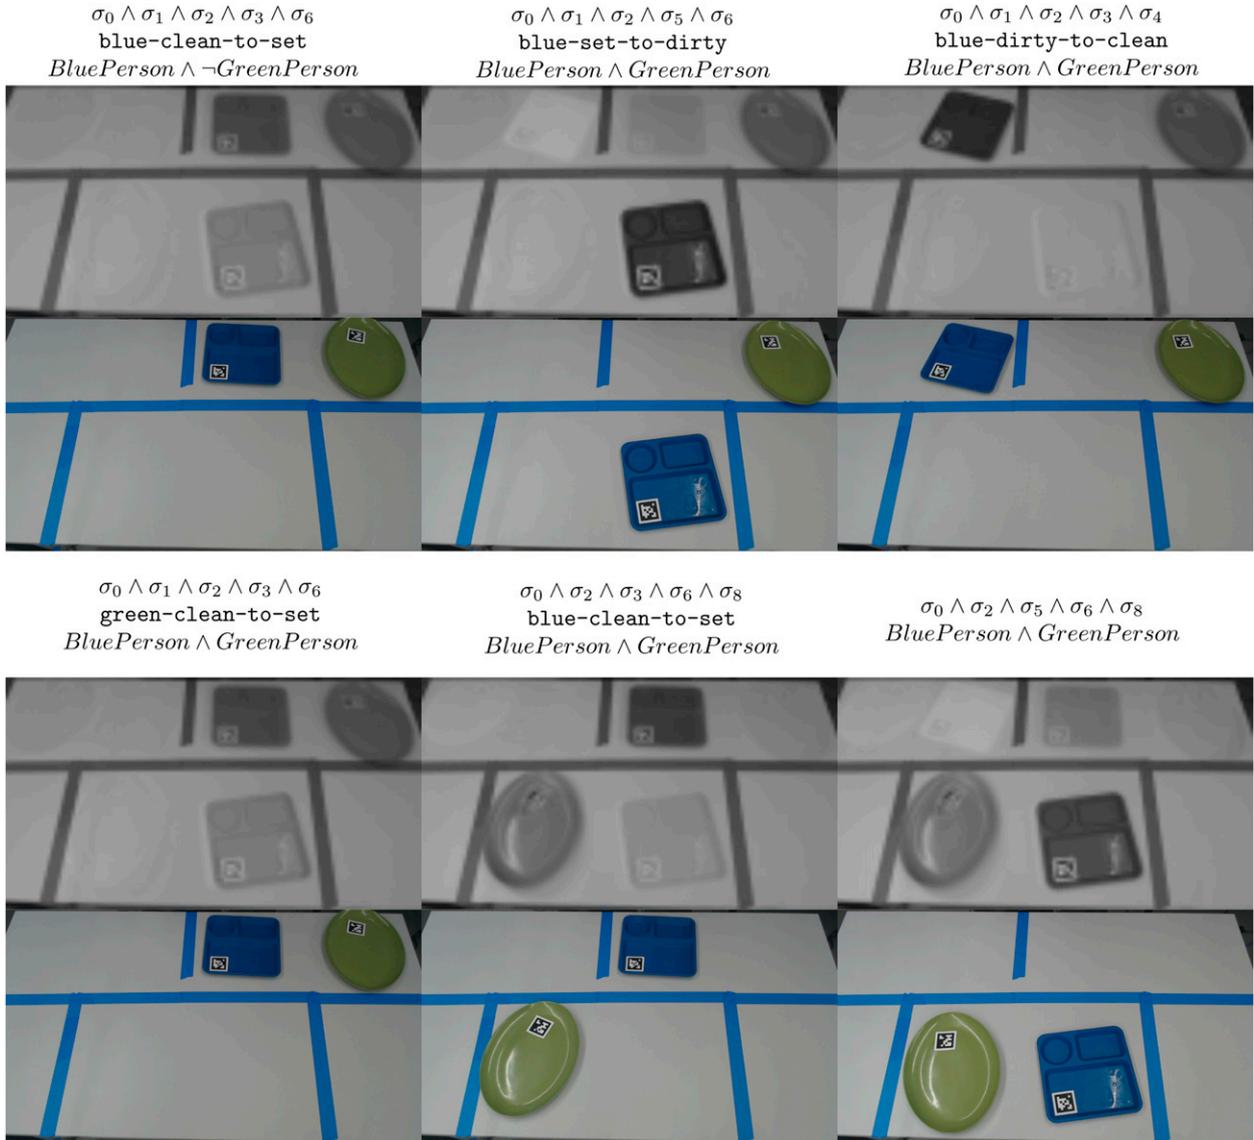


Figure 11. Sequence of states needed to achieve the liveness guarantees in Figure 8 for the Baxter Plates example. The strategy gives skills that react to the truth value of *BluePerson* and *GreenPerson*. We show the physical and symbolic interpretation of each state. The upper image of each pair is the visualization of the symbolic state, created by combining the symbols that are True. The lower image of each pair is the physical state. Above each pair of images, we show which symbols are True, which action should be taken, and the truth value of *BluePerson* and *GreenPerson*.

type of scenario could occur if there was an obstacle in location B.

Unrealizable Specification 2: In Figure 12(b), we show the added constraint that the robot never use skill a_{c-to-b} , $\varphi_{t,task}^s = \Box \neg a_{c-to-b} \wedge \Box \neg \bigcirc a_{c-to-b}$. This type of scenario could occur if a motor enabling skill a_{c-to-b} was damaged and the skill could not be performed.

Unrealizable Specifications 1 and 2 Repair: For both Unrealizable Specifications 1 and 2, the enumeration-based repair process found $\mathcal{A}_{n.o.t} = \{a_{c-to-b}\}$, corresponding to the precondition that the green block (block 3) be in location C. We only searched for one additional skill, so we did not need to find combinations of skills in Line 5 of Algorithm 1. The repair process searched through 62 skills to find six skill

suggestions for Unrealizable Specification 1 and three skill suggestions for Unrealizable Specification 2 in 2.7 min and 2.9 min, respectively. The user needs to determine which suggestion to implement. Some suggestions were not physically possible, making them impossible to implement, such as a suggestion with $\sigma_{pre(a_{new})} = \{\{\sigma_1\}\}$ and $\sigma_{eff(a_{new})}^{True} = \{\sigma_{17}, \sigma_1, \sigma_{16}\}$, corresponding to moving the green block (block 3) to the x position of location F and the y and z position of location C, which would leave the block floating in the air. If desired, the user could add additional constraints to the specification to remove the physically impossible suggestions. One skill suggestion for both Unrealizable Specifications 1 and 2, with $\sigma_{pre(a)}[a_{new}] = \{\{\sigma_1\}\}$ and $\sigma_{eff(a_{new})}^{True} = \{\sigma_{17}, \sigma_{18}, \sigma_{16}\}$, corresponding to

Table 2. Overview of the repair suggestions found. The synthesis-based repair approach takes substantially less time than the enumeration-based repair approach. For the Kinova Vials example, we were not able to run the enumeration-based repair approach until completion; the numbers listed are for stopping the repair process midway.

	Unrealizable spec	Enumeration-based		Synthesis-based	
		Number of suggestions	Time (sec)	Number of suggestions	Time (sec)
Baxter Blocks	1	6	163.2	50	70
	2	3	171.6	36	47
	3	68	6420	32	49
	4	17	2868	35	32
Baxter Plates	5	9	1214	1	1.1
Kinova Vials	6	22	25167	25	381

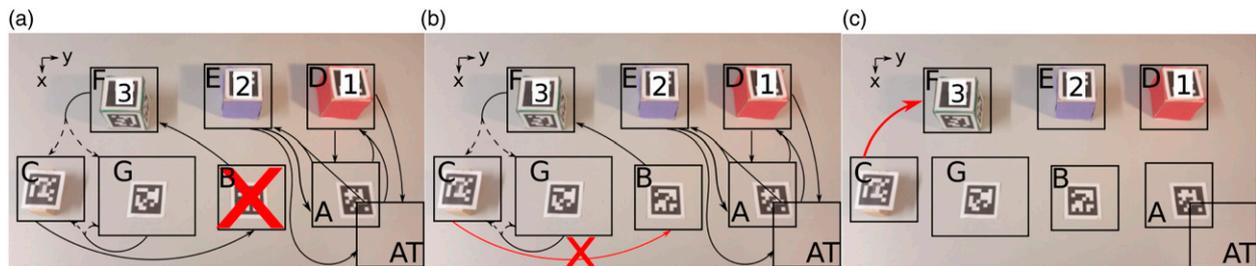


Figure 12. (a) The added constraint in Unrealizable Specification 1 that the green block (block 3) should never be in location B. (b) The added constraint in Unrealizable Specification 2 that the skill $c\text{-to-b}$ should never be executed. (c) A skill moving the green block (block 3) from location C to location F makes both Unrealizable Specification 1 and 2 (and 3 and 4) realizable. Figure from Pacheck et al. (2019).

moving the green block (block 3) from location C to F, is physically possible. When this skill is added to Unrealizable Specifications 1 and 2, the task is realizable.

The synthesis-based repair process found 50 suggestions in 70 s for Unrealizable Specification 1 and 36 suggestions in 47 s for Unrealizable Specification 2. Figure 13 shows selected suggestions to repair Unrealizable Specifications 1 and 2. The repair suggestions for Unrealizable Specifications 1 and 2 involve adding one or more new skills or relaxing the preconditions.

The synthesis-based repair process produces suggestions that exploit all possible preconditions of the existing skills. Suggestion 1 proposes 3 skills to repair Unrealizable Specification 1. For all of the skills in Suggestion 1, the red block (block 1) is in location D and blue block (block 2) is in location A. Each row in Figure 13(a) shows one skill. The first skill moves the green block (block 3) from location G to either floating above a new location with the x position of F and the y position of C (Postcondition 1), a new location with the x position of F and the y position of C (Postcondition 2), F (Postcondition 3), or D (Postcondition 4). The other two skills move the green block (block 3) from on the table in location C or stacked in location C.

Only Postcondition 3 was seen during the original symbol learning. However, the synthesis-based repair process exploits the existing preconditions and postconditions of other skills when it creates or modifies skills, enabling the robot to create behaviors not observed before.

The remainder of the suggestions to repair Unrealizable Specification 1 involving adding new skills are similar to Suggestion 1, with multiple skills that move the green block (block 3) to different locations that may not be physically possible, but still fall in the precondition of skill $a_{f\text{-to-c}}$. The other suggestions have different configurations of the red and blue blocks (blocks 1 and 2). Since we do not add constraints that different blocks cannot be in the same physical position or floating in the air, some suggestions may be physically impossible. Suggestion 2 for Unrealizable Specification 1 in Figure 13(b) shows a suggestion with the blue block (block 2) in location E, but with a z value that is above the table.

The synthesis-based repair process generates similar suggestions for Unrealizable Specification 2. Suggestion 1 proposes two skills to make the specification realizable. One skill moves the green block (block 3) from location G to either location B, E, or F. The other skill moves the green block (block 3) from location C to either location B, E, or F.

The suggestions that relax the preconditions of the skills for Unrealizable Specification 1 and 2 are physically impossible. These involve adding preconditions that allow for blocks to be floating in the air or at the same location as other blocks, which is not possible. We show examples of preconditions being relaxed in the Kinova Vials Unrealizable Specification 6. If desired, additional constraints could be added to the specification to generate suggestions without these physically impossible suggestions.

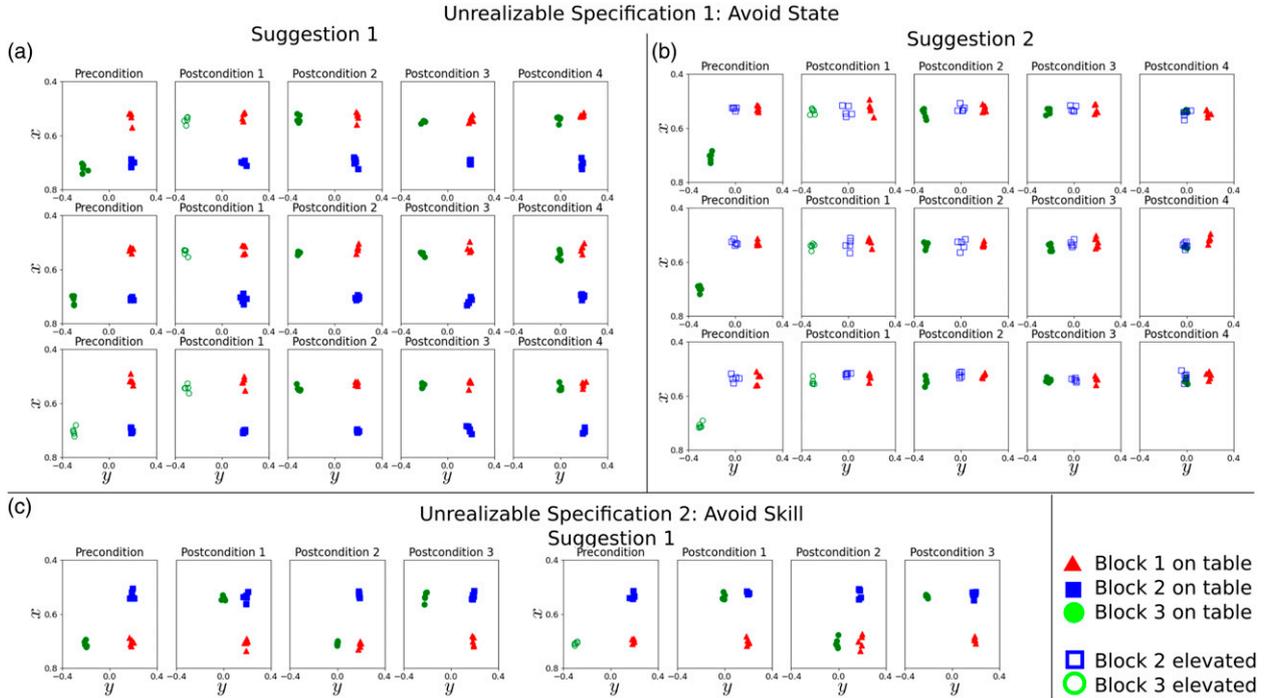


Figure 13. Selected suggestions from the synthesis-based repair process for Unrealizable Specifications 1 and 2. (a) Suggestion 1 for Unrealizable Specification 1 proposes three skills with nondeterministic postconditions. (b) Suggestion 2 for Unrealizable Specification 1 proposes three skills. These skills have the precondition and postcondition that block 2 is hovering over location E. This is physically impossible, but the specification does not disallow it. (c) Suggestion 1 for Unrealizable Specification 2 proposes two skills with nondeterministic postconditions.

Unrealizable Specifications 3 and 4 involved modifying the base specification.

Unrealizable Specification 3: We removed skill a_{c-to-b} from \mathcal{A} before writing the specification, using the same set of symbols Σ as in Section 7.3. The user-defined task was the same as in Figure 7.

Unrealizable Specification 4: We removed all data pertaining to skill a_{c-to-b} before the symbol generation process. This resulted in a different set of symbols, Σ . The user-defined task was the same as represented in Figure 7. There were no longer symbols corresponding to the green block (block 3) being in location B, as symbols are only generated from effect sets, so the subscripts of the symbols in the liveness guarantees shown in Figure 7 were different.

Unrealizable Specification 3 and 4 Repair: For both Unrealizable Specification 3 and 4, the enumeration-based repair process found $\mathcal{A}_{n.o.t} = \emptyset$, requiring an exhaustive search of the skill space. For Unrealizable Specification 3, the repair process searched through 1172 skills and found 68 possible new skills in 107 min. For Unrealizable Specification 4, the repair process searched through 788 skills and found 17 possible skills to repair the specification in 48 min. The repair process suggested a skill that would move the green block (block 3) from both locations C and G to location F for both specifications. With the fairness assumption in equation (14), this has the same result as giving the robot a skill moving the green block (block 3) from C to

F. When we added a skill that moved the green block (block 3) from C to F, both specifications were realizable.

The synthesis-based repair found 32 suggestions to repair Unrealizable Specification 3 in 49 s and 35 suggestions to repair Unrealizable Specification 4 in 32 s. Figure 14 shows selected suggestions from the synthesis-based repair process to repair Unrealizable Specifications 3 and 4.

We show one suggestion to repair Unrealizable Specification 3 in Figure 14(a). This suggestion is similar to those for Unrealizable Specifications 1 and 2. The repair process suggests two new skills that move the green block (block 3), one that has a precondition of location C and one that has a precondition of location G. Both skills have postconditions that move to either location B, E, or F. The remainder of the suggestions to repair Unrealizable Specification 3 are similar but have different configurations of the red and blue blocks (blocks 1 and 2). Again, the suggestions involving relaxing the preconditions are physically impossible.

One of the suggestions to repair Unrealizable Specification 4 is shown in Figure 14(b). This suggestion proposed two skills to move the green block (block 3) to location F from either C or G. Note that since we removed data to generate Unrealizable Specification 4, there is only a symbol pertaining to the green block (block 3) being elevated off the table.

Baxter Plates: For the Baxter Plates example, we investigate one unrealizable specification.

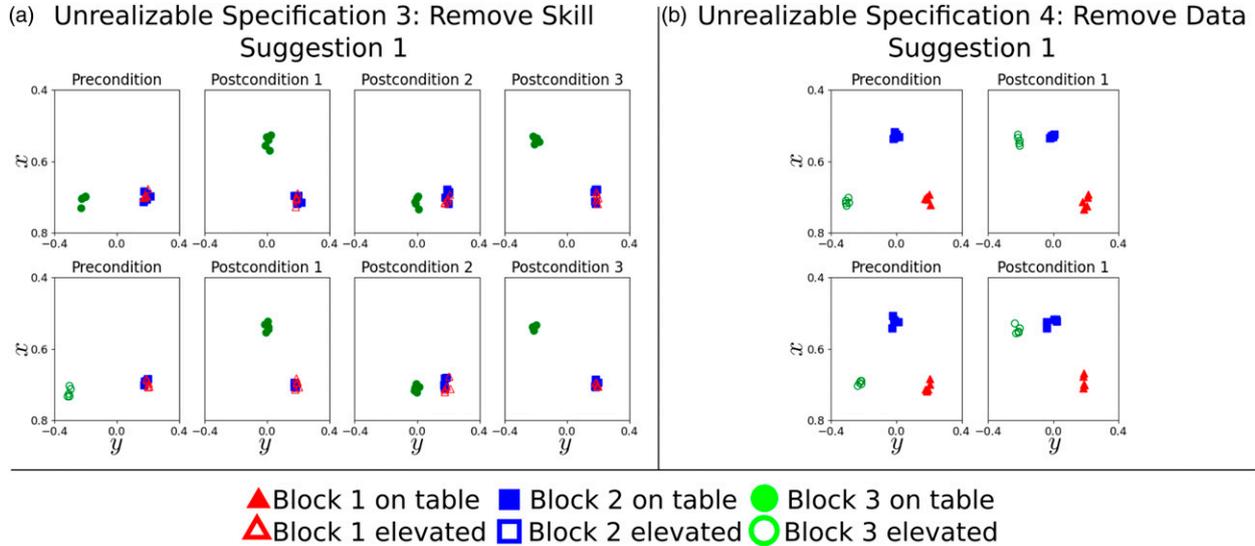


Figure 14. Selected suggestions from the synthesis-based repair process for Unrealizable Specifications 3 and 4. (a) Two additional skills to move block 3 that have nondeterministic postconditions of locations B, E, and F. One skill starts from location C and one from location G. (b) One suggestion to repair Unrealizable Specification 4 involves two new skills. One moves block 3 from location G to F and the other moves block 3 from location C to F.

Unrealizable Specification 5: The user-defined liveness guarantees are the same as in Figure 8. We add the additional constraints that the when either *BluePerson* or *GreenPerson* is True, the same colored plate should not be moved out of the set position. This is encoded in $\varphi_{t, \text{task}}^s$ as shown in Figure 15(a).

Unrealizable Specification 5 Repair: The enumeration-based repair process found $\mathcal{A}_{n.o.t} \neq \emptyset$. We considered repair suggestions that consisted of two additional skills. The repair process searched through 3840 potential new skills and found 9 suggestions in 20.23 min. We show one of the suggestions in Figure 15(b). The first skill has the precondition σ_5 and postcondition σ_8 . When this skill is executed in the repaired strategy, this corresponds to the blue plate moving from set to dirty when the green plate is set. The second skill has the precondition σ_8 and postconditions $\sigma_3 \wedge \sigma_4$. This corresponds to moving the green plate from clean to set when the blue plate is set.

The synthesis-based repair process found 1 suggestion in 1.1 s. This suggestion is to reduce the nondeterminism in skill $a_{\text{green-clean-to-set}}$ such that it does nothing. This suggestion is valid symbolically because it works to violate the liveness assumption that the skill $\square \diamond (a_{\text{green-clean-to-set}} \rightarrow \sigma_8)$.

We investigated if the synthesis-based repair would find suggestions for repair if the symbol generation process was able to determine the skill symbolic structure correctly. We modified the learned skills to remove the nondeterminism, removing the option for skills $a_{\text{green-clean-to-set}}$ and $a_{\text{blue-clean-to-set}}$ to not change any environment variables.

The synthesis-based repair was unable to find any other suggestions for repair. During the synthesis-based repair process, Algorithm 2, restrictPostconditions, finds

the required postconditions for new skills. The algorithm takes the current set of winning states, Z ; the postconditions of new skills need to be in Z . As found by the enumeration-based repair, we need two new skills: one with postconditions $\sigma_0 \wedge \sigma_2 \wedge \sigma_5 \wedge \sigma_6 \wedge \sigma_8$ and one with postconditions $\sigma_0 \wedge \sigma_2 \wedge \sigma_3 \wedge \sigma_4 \wedge \sigma_8$. However, during the synthesis process, all states with either postcondition are removed from Z before the repair process is started. As a result, we are unable to find a suggestion with the synthesis-based repair process.

When we use the synthesis-based repair as proposed in Pacheck et al. (2020) without modification, we are able to find suggestions to repair the specification. However, these suggestions only contain skills which change the truth value of the user-defined variables *BluePerson* and *GreenPerson*, which is not desired. Essentially, they let the robot do its task by enforcing restrictions on how the people are behaving. One such suggestion modifies the skills such that when the green plate and blue plate are both clean, if *BluePerson* = True, it must be False at the next step if $a_{\text{green-clean-to-set}}$ is not executed to make sure the green plate could be set if needed. The suggestion also includes extra skills that enforce when the green plate and blue plate are both set and *BluePerson* = False and *GreenPerson* = True, *GreenPerson* must become False so the robot can move the green plate, enabling it to then move the blue plate.

Kinova Vials: For the Kinova Vials example, we investigate one unrealizable specification.

Unrealizable Specification 6: The user-defined task was the same as in Figure 9. We added the constraint that the red and green vials should never be in the same rack (either white or yellow) at the same time. The additional constraint in $\varphi_{t, \text{task}}^s$ that the red and green vials not be in the white rack

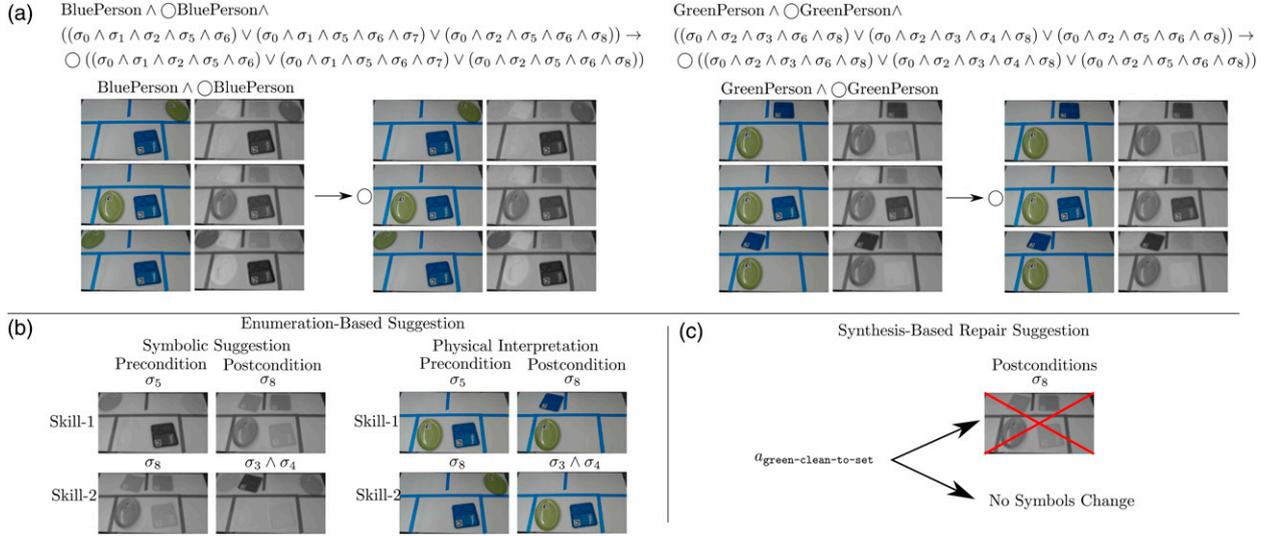


Figure 15. (a) The added constraint in Unrealizable Specification 5 for the Baxter Plates example. When *BluePerson* is True at both the current and next step and the blue plate is set, the blue plate should be set at the next step. Similarly, when *GreenPerson* is True at both the current and next step and the green plate is set, the green plate should be set at the next step. (b) One of the 9 suggestions found by the enumeration-based repair process. The suggestion proposes two new skills. The first skill moves the blue plate from set to dirty when the green plate is set. The second skill moves the green plate from clean to set when the blue plate is set. (c) The synthesis-based repair process finds one suggestion to repair the specification. This suggestion proposes modifying the skill $a_{\text{green-clean-to-set}}$ such that it does nothing to violate the added environment liveness assumption.

at the same time is encoded in equation (17) and the physical interpretation is shown in Figure 16(c) and (f). The physical interpretation that the red and green vials not be in the same yellow rack is shown in Figure 16(a), (b), (d) and (e).

$$\begin{aligned} & \square \neg (\sigma_4 \wedge \sigma_{11}) \wedge \square \neg (\sigma_5 \wedge \sigma_{10}) \\ & \square \neg \bigcirc (\sigma_4 \wedge \sigma_{11}) \wedge \square \neg \bigcirc (\sigma_5 \wedge \sigma_{10}) \end{aligned} \quad (17)$$

We also add the constraint to $\varphi_{l, \text{task}}^s$ that two vials cannot be in the same physical location at the same time. We show a portion of this constraint in equation (18) and the full constraint in Appendix A.3.

$$\square \neg (\sigma_0 \wedge \sigma_6) \wedge \square \neg (\sigma_0 \wedge \sigma_{12}) \wedge \square \neg (\sigma_1 \wedge \sigma_7) \wedge \dots \quad (18)$$

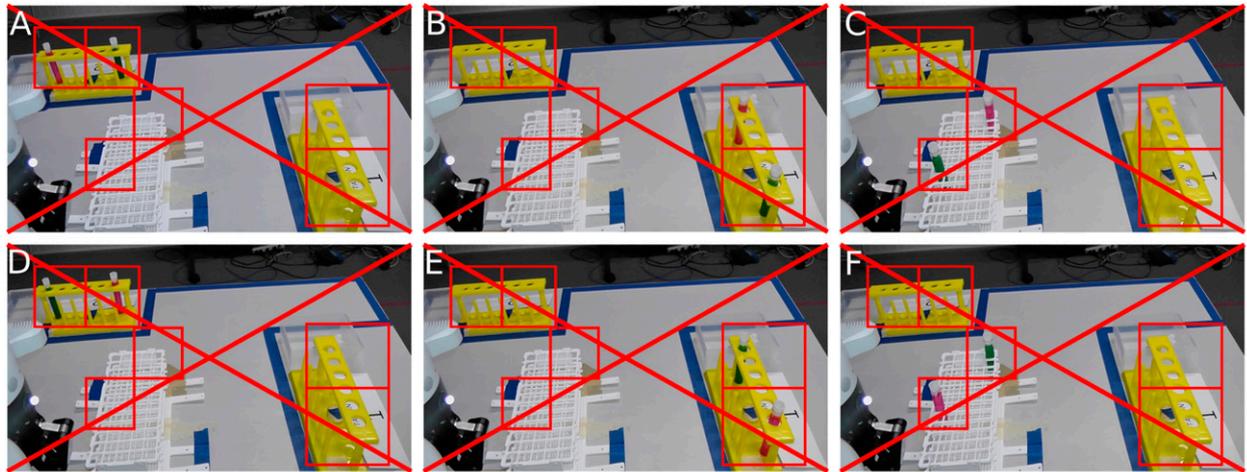
$$\begin{aligned} & ((\sigma_6 \leftrightarrow \neg \bigcirc \sigma_6) \vee (\sigma_7 \leftrightarrow \neg \bigcirc \sigma_7) \vee (\sigma_8 \leftrightarrow \neg \bigcirc \sigma_8) \vee (\sigma_9 \leftrightarrow \neg \bigcirc \sigma_9) \vee (\sigma_{10} \leftrightarrow \neg \bigcirc \sigma_{10}) \vee (\sigma_{11} \leftrightarrow \neg \bigcirc \sigma_{11})) \rightarrow \\ & ((\sigma_0 \leftrightarrow \bigcirc \sigma_0) \wedge (\sigma_1 \leftrightarrow \bigcirc \sigma_1) \wedge (\sigma_2 \leftrightarrow \bigcirc \sigma_2) \wedge (\sigma_3 \leftrightarrow \bigcirc \sigma_3) \wedge (\sigma_4 \leftrightarrow \bigcirc \sigma_4) \wedge (\sigma_5 \leftrightarrow \bigcirc \sigma_5)) \wedge \\ & ((\sigma_{12} \leftrightarrow \bigcirc \sigma_{12}) \wedge (\sigma_{13} \leftrightarrow \bigcirc \sigma_{13}) \wedge (\sigma_{14} \leftrightarrow \bigcirc \sigma_{14}) \wedge (\sigma_{15} \leftrightarrow \bigcirc \sigma_{15}) \wedge (\sigma_{16} \leftrightarrow \bigcirc \sigma_{16}) \wedge (\sigma_{17} \leftrightarrow \bigcirc \sigma_{17})) \end{aligned} \quad (19)$$

Additionally, we add that only one vial can be moved at a time. We show the constraint that if the red vial moves, the green and yellow vials cannot move in equation (19). The constraints for the green and yellow vials are similar and we show them in Appendix A.3. Without these additional constraints, the suggestions returned are difficult to interpret and are unusable because the suggestions involve moving multiple vials at once, which a single arm cannot do.

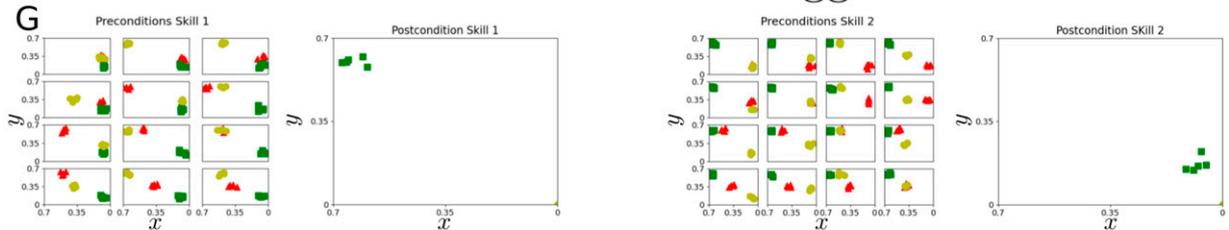
Unrealizable Specification 6 Repair: For the enumeration-based repair, we found that $\mathcal{A}_{n.o.t} \neq \emptyset$. However, at least two skills are required to repair the specification. This necessitates looping through combinations of the skills in $\mathcal{A}_{n.o.t}$ and the 48 other partitioned skills. We stopped the repair process after 7 h and found 22 suggestions. One of the suggestions was a skill to move the green vial from the right-bottom location to the top-left location and another skill to move the green vial from the top-left location to the right-bottom location. Figure 16(g) shows this suggestion.

The synthesis-based repair process returns 25 suggestions in 381 s. This is orders of magnitude faster than the enumeration-based repair process. One of the suggestions is shown in Figure 16(h). The suggestion proposes three new

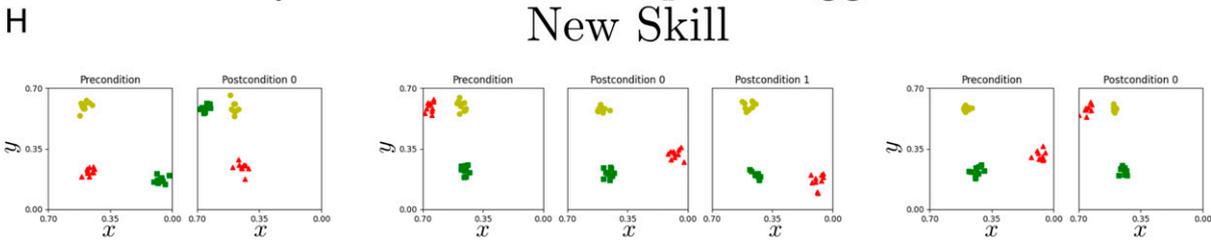
skills: one with preconditions $\sigma_0 \wedge \sigma_{10} \wedge \sigma_{15}$ and postcondition $\sigma_2 \wedge \sigma_{10} \wedge \sigma_{15}$, one with preconditions $\sigma_4 \wedge \sigma_8 \wedge \sigma_{15}$ and postconditions $(\sigma_4 \wedge \sigma_6 \wedge \sigma_{15}) \vee (\sigma_4 \wedge \sigma_7 \wedge \sigma_{15})$, and one with preconditions $\sigma_4 \wedge \sigma_7 \wedge \sigma_{15}$ and postcondition $\sigma_4 \wedge \sigma_8 \wedge \sigma_{15}$. Note that the synthesis-based repair suggestion contains three skills, even though the specification could be repaired with two, while the enumeration-based repair only contains two skills. This is because the synthesis-based repair has no restrictions on how many skills it can propose and



Enumeration-Based Suggestion



Synthesis-Based Repair Suggestions New Skill



Relax Preconditions

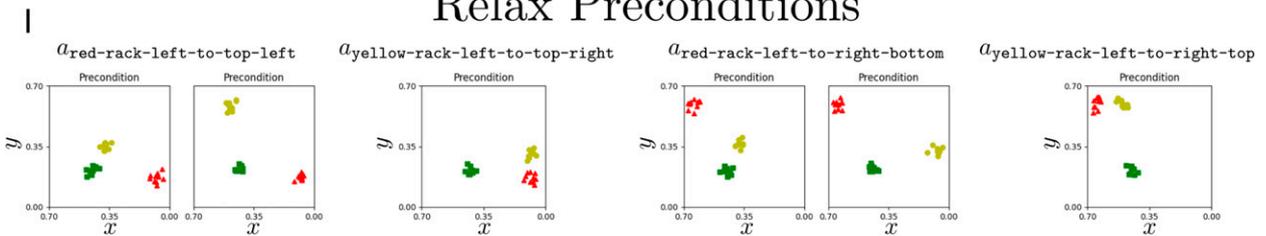


Figure 16. (a–f) Images showing the configurations of the red and green vials that should not occur based on the task specification. In the suggestion plots, the red vials are denoted by red triangles, the green vials by green squares, and the yellow vials by yellow circles. (g) One suggestion given by enumeration-based repair. There are two skills suggested. Note that some of the vials in the preconditions overlap because the preconditions are taken from existing skills. When determining which symbols are contained in the precondition classifier, we do not enforce mutual exclusion of symbols added in φ_{task} , these states are never visited, so the specification is not violated. (h) Suggestion from the synthesis-based repair. There are three skills suggested. (i) A suggestion from the synthesis-based repair process to relax the preconditions. The repair process suggests relaxing the preconditions of four skills.

does not attempt to minimize the number of skills suggested. The enumeration-based repair would consider two skills to repair a specification before considering three, as it would add substantial computation expense.

The synthesis-based repair process also produces suggestions that relax the preconditions of skills. We show one

such suggestion in Figure 16(i). The suggestion suggests relaxing the preconditions of $a_{red-rack-left-to-top-left}$ to include states when the red vial is in the right-bottom location, the green vial is in the rack-left, and the yellow vial is in the rack-right or top-right locations. It also suggests allowing $a_{yellow-rack-left-to-top-right}$ to be allowed when the yellow vial

is in the right-top location, the red vial is in the right-bottom location, and the green vial is in the rack-left location. The third precondition it suggests relaxing is that for $a_{red-rack-left-to-right-bottom}$ to include the state when the red vial is in the top-left location, the yellow vial is in the rack-right or right-top locations, and the green vial is in the rack-left location. It also suggests allowing $a_{yellow-rack-left-to-right-top}$ to be executed when the yellow vial is in the top-left location, the red vial is in the top-left location, and the green vial is in the rack-left location. Note that skills involving moving the yellow vial are not strictly necessary to repair the specification, but are included due to the choices made by the synthesis-based repair process when extracting a strategy.

7.8. Enumeration-based versus synthesis-based repair

Both the enumeration-based and synthesis-based repair are able to find symbolic suggestions for skills to repair unrealizable tasks. The quality and speed of the enumeration-based and synthesis-based suggestions differ depending on the task and current skills available to the robot.

Table 2 shows the synthesis-based repair is faster than the enumeration-based repair, especially as the number of skills and symbols increases. The difference in time to find suggestions was especially apparent in the Kinova Vials example, where we were not able to run the enumeration-based repair to completion. This disparity is pronounced because the enumeration-based repair needs to enumerate all possible combinations of skills, which does not scale well when there are multiple skills required to repair a specification. It is possible to terminate both repair processes early and only receive a portion of the suggestions; however, it is not possible to know at which point in the repair process the suggestion desired by the user will be found.

While the enumeration-based repair takes longer than the synthesis-based repair, the suggestions returned tend to be more interpretable due to both the number and type of skills suggested. The enumeration-based repair attempts to find skills that have the preconditions of existing skills and the effect masks of existing skills. As a result, the suggested skills will look similar to the existing skills. The synthesis-based repair process suggests skills with postconditions and preconditions that do not necessarily resemble the existing skills. There is no limit on the number of skills provided by the synthesis-based repair process in a single suggestion, which can make suggestions more difficult to interpret. For example, in the suggestion shown in Figure 16, the synthesis-based repair process suggests 3 skills, while the enumeration-based repair process only suggests 2 skills. For the synthesis-based repair process, the skills that are suggested are highly dependent on the choices the system makes during the process of finding a strategy. Changing the order of states visited by the system during the determination process will likely result in different suggestions.

The synthesis-based repair does not always find suggestions to repair the specification. As shown in Section 7

with the Baxter Plates example, there are certain specifications for which the only way the synthesis-based repair can provide suggestions is by suggesting skills that control the value of user-defined variables, essentially enforcing a behavior for uncontrolled events.

8. Conclusion

In this work, we present a framework for automatically encoding the skills of a robot in an LTL formula from sensor data. We provide a task to the robot and generate a strategy to accomplish the task if possible. If the task is not possible, we show two methods to repair the specifications by providing skill suggestions that would make the task possible. We demonstrate the process of encoding the skills in an LTL formula, an enumeration-based repair process, and a synthesis-based repair process on three examples.

We plan to extend this work in the future. Currently, our approach does not consider the state during skill execution. We plan to extend the synthesis-based repair process to address this limitation by allowing specifications with constraints on what occurs during skill execution. Additionally, we plan to handle a larger class of specifications, such as those in the Baxter Plates example. This work requires the user to implement the suggestions by hand. Future work will consist of automatically implementing controllers based on the symbolic suggestions, further reducing the burden on the user. Lastly, we plan to demonstrate the scalability of our algorithm on larger and more complicated specifications.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work is supported by the Office of Naval Research award N00014-17-1-2699.

ORCID iDs

Adam Pacheck  <https://orcid.org/0000-0002-4349-1888>

Steven James  <https://orcid.org/0000-0003-4366-4125>

Supplemental Material

Supplemental material for this article is available online.

Note

1. Note that σ are only generated from effect sets in Konidaris et al. (2018).

References

- Ahmetoglu A, Seker M, Sayin A, et al. (2020) DeepSym: deep symbol generation and rule learning from unsupervised

- continuous robot interaction for planning. ArXiv preprint arXiv:2012.02532.
- Alur R, Moarref S and Topcu U (2013) Counter-strategy guided refinement of gr (1) temporal logic specifications. In: Formal methods in computer-aided design, Portland, OR, 20–23 October 2013, pp. 26–33. IEEE.
- Ames B, Thackston A and Konidaris G (2018) Learning symbolic representations for planning with parameterized skills. In: IEEE/RSJ international conference on intelligent robots and systems, Madrid, Spain, 01–05 October 2018, pp. 526–533. IEEE.
- Asai M (2021) Unsupervised grounding of plannable first-order logic representation from images. *Proceedings of the International Conference on Automated Planning and Scheduling* 29: 583–591.
- Asai M and Fukunaga A (2018) Classical planning in deep latent space: bridging the subsymbolic-symbolic boundary *Proceedings of the AAAI Conference on Artificial Intelligence* 32.
- Bloem R, Jobstmann B, Piterman N, et al. (2012) Synthesis of reactive (1) designs. *Journal of Computer and System Sciences* 78(3): 911–938.
- Chatterjee K, Henzinger TA and Jobstmann B (2008) Environment assumptions for synthesis. *International Conference on Concurrency Theory*. Berlin, Germany: Springer, 147–161.
- Cortes C and Vapnik V (1995) Support-vector networks. *Machine Learning* 20(3): 273–297.
- Dantam NT, Kingston ZK, Chaudhuri S, et al. (2018) An incremental constraint-based framework for task and motion planning. *The International Journal of Robotics Research* 37(10): 1134–1151. DOI: [10.1177/0278364918761570](https://doi.org/10.1177/0278364918761570)
- DeCastro JA and Kress-Gazit H (2016) Nonlinear controller synthesis and automatic workspace partitioning for reactive high-level behaviors. In: International conference on hybrid systems: computation and control, Pittsburgh, PA, April 2016, 225–234. ACM.
- Ehlers R and Raman V (2016) Slugs: extensible gr (1) synthesis. *International Conference on Computer Aided Verification*. Berlin, Germany: Springer, 333–339.
- Ester M, Kriegel HP, Sander J, et al. (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proceedings of the second international conference on knowledge discovery and data mining, KDD'96, 2–4 August 1996, Portland, Oregon, pp. 226–231. AAAI Press.
- Fainekos GE (2011) Revising temporal logic specifications for motion planning. In: International conference on robotics and automation, 09–13 May 2011, Shanghai, China, pp. 40–45. IEEE.
- Fikes RE and Nilsson NJ (1971) Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2(3–4): 189–208.
- Finucane C, Jing G and Kress-Gazit H (2010) Ltllmop: experimenting with language, temporal logic and robot control. In: IEEE/RSJ international conference on intelligent robots and systems, 18–22 October 2010, Taipei, Taiwan, pp. 1988–1993.
- Fox M and Long D (2003) Pddl2. 1: an extension to pddl for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20: 61–124.
- Garrett CR, Chitnis R, Holladay R, et al. (2021) Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems* 4: 265–293.
- Ghallab M, Nau D and Traverso P (2004) *Automated Planning: Theory and Practice*. Amsterdam, Netherlands: Elsevier.
- Ghallab M, Nau D and Traverso P (2016) *Automated Planning and Acting*. Cambridge, UK: Cambridge University Press.
- He K, Lahijanian M, Kavraki LE, et al. (2019) Automated abstraction of manipulation domains for cost-based reactive synthesis. *IEEE Robotics and Automation Letters* 4: 285–292.
- He K, Wells A, Kavraki LE, et al. (2019) Efficient symbolic reactive synthesis for finite-horizon tasks. In: International conference on robotics and automation, Montreal, QC, 20–24 May 2019, pp. 8993–8999. IEEE.
- Hollander M, Wolfe DA and Chicken E (2014) *Nonparametric Statistical Methods*. 3rd edition. New York, NY: John Wiley & Sons, Inc.
- Hyvärinen A and Oja E (2000) Independent component analysis: algorithms and applications. *Neural Networks: The Official Journal of the International Neural Network Society* 13(4–5): 411–430.
- James S, Rosman B and Konidaris G (2020) Learning portable representations for high-level planning. In: International conference on machine learning, Baltimore, MD, 23–29 Jul 2023, pp. 4682–4691. PMLR.
- Jetchev N, Lang T and Toussaint M (2013) Learning grounded relational symbols from continuous data for abstract reasoning. In: Proceedings of the 2013 ICRA workshop on autonomous learning, Karlsruhe, Germany, 6–10 May 2013.
- Kim K, Fainekos G and Sankaranarayanan S (2015) On the minimal revision problem of specification automata. *The International Journal of Robotics Research* 34(12): 1515–1535.
- Konidaris G, Kaelbling LP and Lozano-Perez T (2018) From skills to symbols: Learning symbolic representations for abstract high-level planning. *Journal of Artificial Intelligence Research* 61: 215–289.
- Könighofer R, Hofferek G and Bloem R (2009) Debugging formal specifications using simple counterstrategies. In: Formal methods in computer-aided design, Austin, TX, 15–18 November 2009, pp. 152–159. IEEE.
- Kress-Gazit H, Fainekos GE and Pappas GJ (2009) Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics* 25(6): 1370–1381.
- Kress-Gazit H, Lahijanian M and Raman V (2018) Synthesis for robots: guarantees and feedback for robot behavior. *Annual Review of Control, Robotics, and Autonomous Systems* 1: 211–236.
- Lahijanian M, Andersson SB and Belta C (2012) Temporal logic motion planning and control with probabilistic satisfaction guarantees. *IEEE Transactions on Robotics* 28(2): 396–409.
- Lahijanian M, Maly MR, Fried D, et al. (2016) Iterative temporal planning in uncertain environments with partial satisfaction guarantees. *IEEE Transactions on Robotics* 32(3): 583–599.

- Lee J, Nam C, Park J, et al. (2021) Tree search-based task and motion planning with prehensile and non-prehensile manipulation for obstacle rearrangement in clutter. In: IEEE international conference on robotics and automation, Xi'an, China, 30 May 2021–05 June 2021, pp. 8516–8522. IEEE.
- Li W, Dworkin L and Seshia SA (2011) Mining assumptions for synthesis. In: ACM/IEEE international conference on formal methods and models for codesign, Cambridge, UK, 11–13 July 2011, pp. 43–50. IEEE Computer Society.
- Mazo M, Davitian A and Tabuada P (2010) Pessoa: a tool for embedded controller synthesis. In: International conference on computer aided verification, Berlin, Heidelberg, 17–22 July 2010, pp. 566–569. Springer-Verlag.
- McDermott D, Ghallab M, Howe A, et al. (1998) Pddl—the planning domain definition language. Yale Center for Computational Vision and Control, Technical report CVC TR98003/DCS TR1165.
- Mugan J and Kuipers B (2009) Autonomously learning an action hierarchy using a learned qualitative state representation. In: Proceedings of the 21st international joint conference on artificial intelligence, Pasadena, CA, 11–17 July 2009, pp. 1175–1180.
- Mugan J and Kuipers B (2012) Autonomous learning of high-level states and actions in continuous environments. *IEEE Transactions on Autonomous Mental Development* 4(1): 70–86.
- Muhayyuddin, Moll M, Kavraki L, et al. (2018) Randomized physics-based motion planning for grasping in cluttered and uncertain environments. *IEEE Robotics and Automation Letters* 3(2): 712–719.
- Pacheck A, Konidaris G and Kress-Gazit H (2019) Automatic encoding and repair of reactive high-level tasks with learned abstract representations. In: International symposium on robotics research, Hanoi, Vietnam, 6–10 October 2019.
- Pacheck A, Moarref S and Kress-Gazit H (2020) Finding missing skills for high-level behaviors. In: International conference on robotics and automation, Paris, France, 31 May 2020–31 August 2020, pp. 10335–10341. IEEE.
- Parzen E (1962) On estimation of a probability density function and mode. *The Annals of Mathematical Statistics* 33(3): 1065–1076.
- Pnueli A (1977) The temporal logic of programs. In: 18th annual symposium on foundations of computer science, Providence, RI, 31 October 1977–02 November 1977, pp. 46–57. IEEE.
- Raman V and Kress-Gazit H (2013) Explaining impossible high-level robot behaviors. *IEEE Transactions on Robotics* 29(1): 94–104.
- Rosenblatt M (1956) Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics* 27: 832–837.
- Srivastava S, Fang E, Riano L, et al. (2014) Combined task and motion planning through an extensible planner-independent interface layer. In: IEEE international conference on robotics and automation, Hong Kong, China, 31 May 2014–07 June 2014, pp. 639–646. IEEE.
- Ugur E and Piater J (2015a) Bottom-up learning of object categories, action effects and logical rules: from continuous manipulative exploration to symbolic planning. In: International conference on robotics and automation, Seattle, WA, 26–30 May 2015, pp. 2627–2633. IEEE.
- Ugur E and Piater J (2015b) Refining discovered symbols with multi-step interaction experience. In: 2015 IEEE-RAS 15th international conference on humanoid robots, Seoul, South Korea, 03–05 November 2015, pp. 1007–1012. IEEE.
- Wang J and Olson E (2016) AprilTag 2: efficient and robust fiducial detection. In: International conference on intelligent robots and systems, Daejeon, South Korea, 09–14 October 2016, pp. 4193–4198. IEEE.
- Wongpiromsarn T, Topcu U and Murray RM (2010) Receding horizon control for temporal logic specifications. In: ACM international conference on hybrid systems: computation and control. ACM, Milan Italy, 4–6 May 2022, pp. 101–110.
- Yoon S, 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, Providence, RI, 22–26 September 2007, pp. 352–359.