Exploring and Evaluating Interplays of BPpy with Deep Reinforcement Learning and Formal Methods

Tom Yaacov¹, Gera Weiss¹, Adiel Ashrov², Guy Katz² and Jules Zisser¹ ¹Ben-Gurion University of the Negev, Israel

tomya@post.bgu.ac.il, geraw@bgu.ac.il, adiel.ashrov@mail.huji.ac.il, guykatz@cs.huji.ac.il, zisserh@post.bgu.ac.il

Keywords: Behavioral Programming, Deep Reinforcement Learning, Formal Methods.

Abstract: We explore and evaluate the interactions between Behavioral Programming (BP) and a range of Artificial Intelligence (AI) and Formal Methods (FM) techniques. Our goal is to demonstrate that BP can serve as an abstraction that integrates various techniques, enabling a multifaceted analysis and a rich development process. Specifically, the paper examines how the BPpy framework, a Python-based implementation of BP, is enhanced by and enhances various FM and AI tools. We assess how integrating BP with tools such as Satisfiability Modulo Theory (SMT) solvers, symbolic and probabilistic model checking, and Deep Reinforcement Learning (DRL) allow us to scale the abilities of BP to model complex systems. Additionally, we illustrate how developers can leverage multiple tools within a single modeling and development task. The paper provides quantitative and qualitative evidence supporting the feasibility of our vision to create a comprehensive toolbox for harnessing AI and FM methods in a unified development framework.

1 INTRODUCTION

It is commonly agreed that the future of software development, especially in reactive systems, relies on the use of models, advanced analysis techniques, and artificial intelligence (AI) (Naveed et al., 2024). However, while the current state-of-the-art involves various analysis tools and techniques, each still uses its own input language and modeling approach, requiring manual integration to combine results with system code. Finding a modeling framework that can safely interweave user code and machine-generated artifacts while supporting formal analysis is a challenging research problem. This paper examines the Behavioral Programming (BP) framework (Harel et al., 2012) and its potential as a unified modeling abstraction that glues tools without manual translation, showing its effectiveness in combining multiple tools throughout the software development process.

BP is a software engineering paradigm designed to allow developers to specify the behavior of reactive systems incrementally and intuitively, aligning with how they perceive the system's requirements (Harel et al., 2012). Its primary strength lies in its ability to break down intricate specifications into manageable components that interact through a unified protocol, enabling the creation of desired behavior. This compositional modeling capability can be applied across various domains (Bar-Sinai et al., 2018) and integrated with techniques from different disciplines (Bar-Sinai and Weiss, 2021).

The benefits afforded by BP to traditional software engineering have been studied extensively (Elyasaf et al., 2019; Elyasaf, 2021). In this work, we aim to complement existing research by focusing on the problem at hand — namely, the integration of BP with AI and formal analysis (FM) techniques. Specifically, we utilize the BPpy library (Yaacov, 2023), a framework for BP in Python, along with several Pythonbased libraries, to explore the interplay and benefits that arise from these combinations. Ultimately, we aim to lay the groundwork for constructing a comprehensive software engineering toolbox to harness modern AI and FM techniques.

Our contributions include:

Satisfiability-Modulo-Theory (SMT): We present an enhanced communication protocol among modules for solving complex constraint systems and show it promotes the efficiency of the execution mechanism.

Symbolic Model Checking: We developed a method that allows the verification of large BP models by analyzing each module separately and combining them symbolically.

Yaacov, T., Weiss, G., Ashrov, A., Katz, G. and Zisser, J.

Paper published under CC license (CC BY-NC-ND 4.0)

In Proceedings of the 20th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2025), pages 27-40 ISBN: 978-989-758-742-9; ISSN: 2184-4895

Proceedings Copyright © 2025 by SCITEPRESS – Science and Technology Publications, Lda

Exploring and Evaluating Interplays of BPpy with Deep Reinforcement Learning and Formal Methods. DOI: 10.5220/0013215200003928

Probabilistic Model Checking: We designed an approach that facilitates the analysis of probabilistic aspects of a behavioral system.

Deep Reinforcement Learning (DRL): We present a framework that improves the alignment of requirements with BP modules, delivering a more efficient execution mechanism.

While this is not the main focus of the paper, it is worth noting that benefits go both ways: BP's effective system components management can boost these tools. For instance, BP can serve as a theory for SMT solvers by dynamically generating candidate runs and incrementally revealing constraints, like other theories guide the solver's reasoning process. In the context of DRL, we leverage the flexibility offered by BP to specify a system's behavior only partially, leaving gaps for the DRL engine to resolve. This approach can help guide the learning process and safeguard the resulting execution mechanism from leading to undesirable outcomes. BP's flexible division of the model into smaller modules can also aid in symbolic or probabilistic verification process and facilitate more natural and direct modeling. This differs from current alternatives, focusing on a monolithic state-based system description.

This study not only focuses on new integrations. While the symbolic and probabilistic verification integrations presented here are novel, the combination of SMT and DRL with BP was discussed in prior research (Katz et al., 2019; Eitan and Harel, 2011; Elyasaf et al., 2019; Yaacov, 2023). To provide a comprehensive understanding of the benefits of these integrations, we incorporated ideas from previous work, introduced new tools, and conducted fresh evaluations comparing them against current alternatives in BP to assess the effectiveness of all tools in new dimensions.

As a glimpse into a future research direction, we also present a small experiment showing the applicability of combining SMT solvers, probabilistic modeling, and DRL in conjunction with BP. The result of this study provides compelling evidence, in our opinion, of the viability of using BP as a modeling approach that supports all the aforementioned methods.

The structure of the paper is as follows: We start with a short introduction to BP and BPpy in Section 2. Subsequently, each integration is discussed individually in sections 3 to 6. These sections include evaluations and analyses. We do not devote separate sections to related work and comprehensive conclusions, as relevant references and discussion are provided within each section. In Section 7, we showcase the combined use of these integrations. All supplementary mate-

rial is available at https://github.com/bThink-BGU/ Papers-2025-ENASE-BPpyEvaluation.

2 BEHAVIORAL PROGRAMMING & BPpy

To introduce the reader to the language of BP, we begin by describing the BPpy package and the general BP principles it implements. In BP, developers specify scenarios, named b-threads, which are simple sequential threads of execution that represent behaviors the system should include or avoid. Each scenario is standalone and is usually self-contained, concerning itself with a specific aspect of the system-typically, a single requirement. During runtime, an applicationagnostic execution mechanism interprets and interweaves these b-threads to generate a cohesive system behavior. Specifically, the mechanism is based on a synchronization protocol presented by (Harel et al., 2010). It consists of each b-thread submitting a statement before selecting each event that the b-program produces. When a b-thread reaches a point where it is ready to submit a statement, it synchronizes a statement with its peers, declaring which events it requests, which events it waits for (but does not request), and which events it blocks (prevents from occurring). After submitting the statement, the b-thread is paused. Once all b-threads have submitted their statements, we say the b-program has reached a synchronization point. At this point, an event arbiter selects a single event that has been requested and is not blocked. It then resumes all b-threads that either requested or waited for the selected event. The other bthreads remain paused, and their statements are used in the next synchronization point. This process is repeated throughout the execution of the program. A formal definition of BP semantics is available as an appendix in the supplementary material.

To make BP's core concepts more concrete, we begin with an illustrative example of a *b-program* (a set of b-threads) implemented in BPpy. The example is an adaptation of one of the sample b-programs presented in the work of (Harel et al., 2012) describing a system responsible for controlling the mixing of hot and cold water from two separate taps. Listing 1 depicts three b-threads. Each b-thread is implemented as a Python generator—a function that can pause itself and pass data back to its caller at any point, using the yield idiom. It can then be resumed when reinvoked with the send method. The statements submitted by each b-thread are structured as sync class instances containing events or event sets labeled by the arguments request, block, waitFor. The execution mechanism in BPpy starts by independently invoking each b-thread generator and awaiting its statement yield. Once all the statements have been collected, an event is selected, and the program resumes its execution based on the aforementioned synchronization protocol.

```
@thread
def add_hot():
    for i in range(3):
        yield sync(request=BEvent("HOT"))
@thread
def add_cold():
    for i in range(3):
        yield sync(request=BEvent("COLD"))
@thread
def control():
    while True:
        yield sync(waitFor=BEvent("HOT"))
        yield sync(waitFor=All(), block=BEvent("HOT"))
```

Listing 1: The HOT/COLD b-program (Harel et al., 2012).

The first two b-threads, add_hot and add_cold, request the event of pouring a small amount of hot and cold water, respectively, three times. Unlike many other programming paradigms, BP offers developers the flexibility not to be bound by a single predefined behavior for the system. Instead, the system has the freedom to select any behavior that aligns with all the defined b-threads. For instance, a b-program consisting of the two b-threads shown in Listing 1 does not impose a specific order for pouring hot and cold water. Consequently, its execution can generate all sequences that include exactly three occurrences of the HOT event and three occurrences of the COLD event.

To illustrate further, consider that after running the initial version of the system for some time, a safety concern arises, and a new requirement is introduced, stating that it is undesirable to have two consecutive HOT events. While it is possible to modify the add_hot b-thread by adding new conditions and statements, the BP paradigm encourages us to maintain the alignment between existing b-threads and their respective requirements and add a new bthread. This approach promotes an incremental and modular development style, where developers can add or remove behaviors independently without affecting other b-threads. Thus, we introduce the control bthread in Listing 1, which repeatedly waits for the occurrence of HOT and then blocks HOT while waiting for any following event using the All event set.

3 BP \Leftrightarrow **SMT SOLVERS**

BPpy implements SMT solver integration following the concepts outlined in (Katz et al., 2019). The implementation is based on the Z3-solver (De Moura and Bjørner, 2008) package, although other solvers may also be used. In this integration, events are represented as an assignment over a set of SMT variables. At each synchronization point, b-threads specify request/block/waitFor constraints over the variables in the form of logical statements to be satisfied. Once all constraints are collected, the execution mechanism invokes the solver to find a satisfying assignment to the variables embedded in these constraints. Specifically, the solver finds an assignment requested by at least one b-thread that is not blocked.

To introduce this integration, we begin with an illustration of the *Cinderella-Stepmother* problem (Bodlaender et al., 2012). This problem involves a two-player game with a system of water buckets. Initially, there are N empty buckets arranged in a circle, each with a capacity of B water units. In each turn, Cinderella's stepmother distributes A water units across the buckets as she chooses. Subsequently, Cinderella empties C adjacent buckets. This cycle of the stepmother pouring and Cinderella emptying repeats. The stepmother's objective is to fill one bucket with B units, while Cinderella aims to prevent any bucket from becoming full.

The code in Listing 2 exemplifies a solver-based version of the Cinderella-Stepmother example. It consists of the main b-thread, which is responsible for changing the assignments to the buckets list of integers based on the constraints generated by the stepmother and cinderella functions. These functions use the last selected assignment, which is returned through the b-thread's yield command. The bucket_limit b-thread ensures that the variables will not exceed the buckets' capacity B. We note that bucket_limit does not specify request or wait-for constraints. Thus, its blocking constraint remains invariant throughout the execution of the b-program.

The SMT solver integration allows b-threads to communicate versatile variable assignments, going beyond discrete events. Along with its advantages in enhancing expressiveness, extensively discussed in (Katz et al., 2019), it can potentially improve BP's computational capabilities during execution. This section aims to validate this assumption by comparing the conventional discrete event mechanism and the presented SMT-based event mechanism. To assess this, a set of empirical experiments was carried out to measure the runtime execution and memory efficiency of a b-program tasked with solving a multipleconstraint problem using both approaches. The complete code developed for this evaluation is available in the supplementary material.

```
buckets = [Int(f"b{i}") for i in range(N)]
def stepmother(prev):
 added=Sum([b-prev.eval(b) for b in buckets])
non_neg=And([b-prev.eval(b)>=0 for b in buckets])
return And(added == A, non_neg)
def cinderella(prev):
    = list(range(N)) + list(range(N))
 def empty(rng):
  cs = [] # constraints list
for j in range(N):
    if j in rng:
     cs.append(buckets[j] == 0)
    else:
     cs.append(buckets[j] == prev.eval(buckets[j]))
 return And(cs)
return Or([empty(r[i:i+C]) for i in range(N)])
@thread
def bucket_limit():
    while True:
  yield sync(block=Or([b > B for b in buckets]))
@thread
def main():
e=yield sync(request=And([b==0 for b in buckets])
for i in range(STEPS):
    e = yield sync(request=stepmother(e))
    e = yield sync(request=cinderella(e))
```

Listing 2: A BP solver-based implementation to the Cinderella-Stepmother program.

The evaluation included three use cases. The first is the Cinderella-Stepmother problem, for which we developed an equivalent discrete implementation and compared its performance with the solver-based implementation presented in Listing 2. This implementation generates all the possible states of the buckets as individual events. We tested the programs with increasing parameter values of B and N. Since the complexity increases very rapidly in the discrete case with N, we only show here the growth with B.

The second use case is an adaptation of the *Lights Out* puzzle game (Anderson and Feil, 1998), which we refer to as the *bit-flip* problem. The problem domain is a Boolean matrix of dimensions $N \times M$. At first, the matrix is randomly initialized, and in each move, a single row or column values are flipped. Also, in each move, one row or column cannot be flipped, and it is therefore blocked. A potential objective is to find a sequence of bit-flips to transition the matrix from one configuration to another. This problem is relevant as Boolean matrices can effectively represent reactive systems such as communication networks.

For the third example, we implemented the *circled polygon* example (see Figure 1). The program's objective is to find a coordinate outside a regular polygon's area but inside the circumscribing unit circle. We tested this problem with an increasing number of polygon edges. The problem relevance stems from its multifaceted nature, common in reactive systems with geometrical constraints involving real-valued parameters that BP is well-suited to solve (Katz et al., 2019; Elyasaf, 2021). The discrete implementation

solves this multi-constraint problem by discretizing the product of the two-dimensional continuous intervals $[-1,1]^2$ (the area enclosing the circle centered at (0, 0)) incrementally. As the number of edges in the polygon grows, the area between the polygon and the circle diminishes, necessitating finer discretization and an increase in the number of events to find a solution.



Figure 1: A circled polygon problem with 4 edges. The colored area between the polygon and the circle represents the area of potential solutions to the problem.

The memory and runtime results for the three use cases are shown in Figure 2. For the Cinderella-Stepmother and bit-flip problem, the runtime and memory performance of the discrete implementation exponentially increases with B and $N \times M$, respectively. This behavior aligns with expectations, as increasing parameter values increase the number of events, which means that more events are being examined at the synchronization point, and hence, additional runtime and memory are needed. The runtime and memory usage of the solver-based Cinderella-Stepmother implementation remained constant as the problem size increased, while the solver-based bit-flip implementation increased exponentially but with a significantly smaller slope.

For the circled polygon problem, both implementations exhibit a significant increase in runtime as the complexity rises, contrasting with previous results. We attribute this latter increase to the fact that we now have many equations over real-valued variables; this may be because, in this model, the SMT engine runs the theory module many times. Interestingly, the runtime of the discrete implementation remains similar or lower until n = 164, after which it spikes significantly. This behavior stems from varying complexities, where some instances are straightforward for the discrete implementation, while others require finer discretization. Regarding memory usage, the solver-based implementation maintains a distinct advantage, stabilizing around n = 100. Nevertheless, the memory values surpass those of previous examples, underscoring the problem's complexity for the solver. Conversely, the discrete version memory usage increases with finer discretization as complexity rises, with spikes similar to the runtime analysis, indicating the discussed unexpected behavior.



Figure 2: The runtime and memory of the discrete and solver-based implementations for the three problems. The values on the y-axis were converted to log_{10} scale due to the variance in the original scales.

In summary, our study encompassed a performance evaluation contrasting discrete and solverbased implementations across the domains of Boolean, Integer, and Real variables, with various types of constraints and different levels of complexities. The experiments confirmed the anticipated superiority of the solver-based approach. Additionally, we showcased the adaptability of BP's SMT solvers integration in tackling various problem scenarios.

The potential use of BP to enhance SMT capabilities has been explored in prior studies (Harel et al., 2013). While SMT solvers are crucial for various applications like software verification and program synthesis, their incompleteness can hinder their effectiveness. Integrating BP with SMT can enhance system robustness and reliability, potentially leading to more advanced automated theorem-proving tools for specific theories like arithmetic and arrays. As mentioned in the introduction, BP can serve as a theory for deciding behavioral constraints in a setting where we try to find runs that satisfy a set of constraints modeled as a composition of b-threads.

Qin et al. utilized SMT in conjunction with BIP models, employing it to prune unsatisfiable transitions in Open Automata derived from Open pNets (Qin et al., 2020). Their work emphasizes formal analysis rather than enhancing the underlying mechanism. In contrast, our approach integrates SMT into BP to introduce a decision mechanism for dynamically selecting events, thereby increasing its expressiveness and adaptability.

4 BP ⇔ SYMBOLIC MODEL CHECKING

This section introduces a new method for verifying bprograms by translating them into the SMV symbolic specification language (McMillan, 1993). A significant advantage of this method is its ability to verify software written in languages like Python, where the states of the b-threads, modeled as generators, are not readily clonable (Yaacov, 2023). We evaluate the efficiency of symbolic model checking in BP programs and compare it with the current best practices in BP verification, which rely on explicit model checking, involving an exhaustive enumeration of the state space.

As in other modeling paradigms, verification is a subject of considerable work. Some introduced model checking methodologies and accompanying tools for verifying safety and liveness properties in behavioral programs (Harel et al., 2011; Bar-Sinai and Weiss, 2021). Other studies have explored the potential benefits of BP's compositionality in scalable verification and its applications (Harel et al., 2013). See also (Qin et al., 2020) for a similar approach using the BIP modeling framework. However, previous work on BP verification primarily concentrated on explicit model checking. As systems grow in size and complexity, the limitations of explicit model checking become apparent, necessitating the adoption of symbolic verification techniques as an imperative alternative.

BPpy's symbolic model checker takes a unique approach by utilizing the inherent compositionality of behavioral programs. It independently explores the state space of each b-thread and analyzes the product space symbolically, avoiding the explicit enumeration of all states. More specifically, when applying symbolic model checking in BPpy, b-programs are automatically translated to an SMV model (McMillan, 1993) and verified symbolically using PyNuSMV (Busard and Pecheur, 2013). The translation maps each b-thread to an SMV module, which reflects the space explored using Depth First Search (DFS). For instance, consider the module presented in Listing 3, which represents the add hot bthread in Listing 1. Within each b-thread module, there exists a local variable called state, capturing the current state of the b-thread. Additionally, each module incorporates local Boolean variables for each event that the associated b-thread may either request or block. For example, since the add_hot b-thread requests the HOT event, its corresponding module contains the Boolean variable HOT_requested, which dynamically changes based on the current state.



Listing 3: The add_hot b-thread translation to SMV.

Listing 4 depicts the main module of the translated SMV model derived from HOT/COLD b-program discussed in Section 2. In the translated model, this main module functions as the event arbiter for the bprogram and implements BP semantics. It initiates an enumerated variable event, representing the currently selected event. The event variable can take any of the events requested across all b-threads, in addition to two auxiliary events, BPROGRAM_START and BPROGRAM_DONE, marking the program's start and end of execution, respectively. Further, the main module activates all the b-threads of the program as module instances. It tracks each b-thread instance, requested and blocked variables, by using the DEFINE operator. To capture the b-program's dynamics, the system defines a transition relation that allows the next event value to be any currently enabled event. If no such event exists, the b-program terminates, and the system transitions to a sink state where event = BPROGRAM_DONE. This setting facilitates the detection of possible violations related to the termination of the b-program, such as deadlocks or early termination scenarios.



Listing 4: The translated main module of the HOT/COLD b-program discussed in Section 2.

We now turn to a performance evaluation for the aforementioned approach. All programs and tools used for evaluation in this study are available in the supplementary material. The evaluation involved several b-programs, including 1) The HOT/COLD example presented in Section 2 with an increasing number of portions (N) and cold b-threads (M); 2) The Dining Philosophers b-program presented in (Elyasaf et al., 2023); and 3) The Tic-Tac-Toe game presented in (Elyasaf et al., 2023). For this evaluation, all the specifications we considered hold (i.e., no violations are discovered through model checking) to examine how quickly the various approaches can traverse the entire search space. We compared the results against similar b-programs verified using BPjs (Bar-Sinai and Weiss, 2021), a Java-based tool for running and analyzing behavioral programs written in JavaScript, which is the current best practice in BP verification. The symbolic verifier in BPpy can function in two modes: Binary Decision Diagrams (BDD) and SATbased Bounded Model Checking (BMC), where the model is unrolled for a fixed number of steps and is checked for violations that can occur within that number of steps or fewer. We compared these modes to the BPjs verifier in unbounded and bounded modes, respectively. The results are available in Table 1.

In terms of memory consumption, both symbolic model checkers outperformed the explicit verifier, especially the BDD-based verifier, exhibiting improved efficiency. This result aligns with expectations, given that the program's state space is symbolically represented through BDDs or logical formulas rather than being explicitly enumerated. Regarding verification time, we observed that the BDD-based model Table 1: The average time and memory (over 10 repetitions) required to verify b-programs in BPpy and BPjs. The Binary Decision Diagrams (BDD) and SAT-based Bounded Model Checking (BMC) modes implemented in BPpy are compared against the BPjs verifier in unbounded and bounded modes, respectively.

					Unbo	Bounded					
				tiı	me ³	men	nory ⁴	time ³		memory ⁴	
	N,M	$ E ^1$	$ S ^{2}$	BDD	BPjs	BDD	BPjs	BMC	BPjs	BMC	BPjs
HOT/	30,1	2	121	2.9	5.1	0.08	0.24	3.8	5.2	0.09	0.23
	60,1	2	122	3.5	5.3	0.08	0.29	5.0	5.3	0.13	0.29
	90,1	2	182	4.2	5.6	0.08	0.34	7.9	5.4	0.19	0.36
	30,2	3	1022	3.5	10.0	0.08	0.78	5.2	8.4	0.10	1.03
	60,2	3	3842	4.4	50.9	0.09	1.21	18.3	22.6	0.18	1.64
COLD	90,2	3	8462	5.3	215.6	0.11	1.32	99.1	102.0	0.37	1.65
	30,3	4	11437	3.9	398.8	0.08	1.19	5.4	223.9	0.10	1.61
	60,3	4	158901	5.1	t.o.	0.11	t.o.	22.7	t.o.	0.20	t.o.
	90,3	4	519151	7.4	t.o.	0.13	t.o.	121.5	t.o.	0.46	t.o.
Distant	3,-	12	106	4.6	6.1	0.12	0.29	5.5	4.5	0.08	0.25
Dining Phil.	6,-	36	30862	t.o.	56.8	t.o.	0.97	8.0	5.0	0.09	0.30
	9,-	54	3299501	t.o.	<i>o.m</i> .	t.o.	<i>o.m</i> .	11.2	5.8	0.10	0.36
TTT	3,-	21	69502	t.o.	1852.1	t.o.	10.01	26.1	750.6	0.15	7.98

¹ the number of program events, ² the number of program states, ³ in seconds, ⁴ in GB, *o.m.* out of memory (16GB), *t.o.* timeout (60 minutes)

checker excelled in the examples where the total number of events remained relatively small. However, as the number of events increased, the verification time for the BDD-based model checker experienced a significant surge. Conversely, when considering the bounded option, BMC demonstrated greater resilience to a growing number of events in the bprogram, positioning itself as a more robust approach.

The translation of B-programs into SMV models, utilizing their inherent compositionality, shows promise for achieving scalability and efficiency in verification. Our future research will focus on enhancing the translation process to improve the verification of programs with a larger number of events. Also, we aim to extend the symbolic model checking support to various execution mechanisms and protocols, including the one outlined in Section 3.

5 BP ⇔ PROBABILISTIC MODELING

This section discusses BP as a modular language for modeling and analyzing probabilistic systems, enhancing its semantics for compositional modeling of probabilistic and non-deterministic behaviors. Traditionally, randomness in BP is introduced through a random function for event selection, with the sync node's standard event selection protocol being an example. However, this is limited as it only allows uniform distributions. Since we want to allow non-uniform distribution, we take the alternative approach of allowing a random behavior of the bthreads. We want to define randomness in a way that allows both execution (sampling) and model checking. This poses a non-trivial challenge because, during model checking, the analysis needs to consider all possible paths with their probabilities. Thus, we introduce the choice idiom, illustrated in Listing 5, where a dictionary maps values to probabilities for defining categorical distributions in b-threads. These distributions can be sampled during the execution of a b-program or transformed for analysis.

def coin_flip(): side = yield choice({"heads":0.4,"tails":0.6}) yield sync(request=BEvent(side))

Listing 5: Uneven coin flip using the choice statement.

The proposed BP model combines nondeterminism and probability for different modeling purposes. Non-determinism is handled through sync points with multiple events, while probabilities are specified using choice. Non-determinism arises in systems when there is a decision to make. Probabilities are used to model uncontrolled random events. Thus, during analysis, non-determinism is typically addressed by analyzing contingencies, while probabilities are computed for each path.

Executing programs with choice statements involves random selection based on a specified distribution, facilitating naive sampling. Model analysis is done by translation into the PRISM language (Kwiatkowska et al., 2011), a tool for analyzing probabilistic systems with a module-based approach. The translation process, largely similar to the one discussed for SMV in Section 4, has been automated and integrated into BPpy. Supplementary material includes technical details about this process, code, and documentation. We utilize the mdp model type in PRISM to analyze systems that involve a combination of non-determinism and probabilistic behavior.

We compared sampling and formal analysis to evaluate translation performance. Experiments timed the analysis of parametrized versions of three models. Sampling involved running models repeatedly and monitoring outcomes. Formal verification translated the model into PRISM format and verified it using Storm probabilistic model checker (Hensel et al., 2022). All programs and tools used for this evaluation can be found in the supplementary material.

The first model, the classic Monty Hall problem, which is infamous for its counter-intuitive solution, is described as follows: "A game host hides a prize behind one of three doors. The contestant guesses which door has the prize. The host then opens one door with no prize. The contestant can stick with their choice or switch." Our evaluation parameterized the total number of doors, prizes, and doors ruled out, as proposed in (Depuydt and Gill, 2012). The model presented in Listing 6 consists of three b-threads. The first b-thread, named hide_prizes, determines where to hide prizes, requests the hiding events, blocks the opening of the doors with hidden prizes until the opening phase is complete, and waits for the contestant to open a door. It checks if the chosen door contains a prize and accordingly announces a win or lose event. We added the repeat, replace , and sorted parameters to the choice statement to ease the use of the idiom in cases where repeated sampling is required. The make_a_guess b-thread waits for the end of the hiding phase, guesses a door, and then blocks its opening. The open_doors b-thread waits for the guess and then requests to open doors while blocking the already opened doors. Then, it requests the event that marks the end of the host door opening phase and the event for the door the contestant eventually opens.

Listing 6: A BPpy model of the Monty Hall problem.

The experiment included generating instances of the b-program from Listing 6 for all parameter combinations up to 10 doors. For sampling, we executed the program 10,000 times to track the win event occurrence and used PRISM and Storm for model checking to calculate the probability of reaching the win state. The analysis time is the entire duration of Storm.

The results of applying two analysis methods to the Monty Hall model with different parameter values are shown in Figure 3. The blue line and halo depict the mean and standard error of collected samples over time. The dashed orange line shows the exact computation value, and the orange circles mark the time of translation to PRISM and the subsequent analysis. These findings indicate that exact analysis yields results before sampling variance reaches acceptable error margins in all examined cases.

Table 2 provides a breakdown of the runtime for the exact analysis in the proposed approach for the Monty Hall model with 10 doors, distinguishing between PRISM model construction and Storm analysis. An expected correlation exists between the time spent on model translation and overall computation time. The primary source of complexity appears to be the number of doors opened, as the modeling process times out entirely after 6 doors. This is likely due to numerous different paths that do not merge, unlike the situation with the number of prizes.

Table 2: Runtime in seconds for exact analysis with various parameter combinations for the Monty Hall model with 10 doors. The table presents PRISM model construction times on the left and Storm analysis times on the right.

	Opened									
Prizes	1	2	3	4	5	6				
1	0.6/0.1	1.1/0.5	9.4/8.4	114.7/536.3	4120.6/t.o.	t.o./t.o.				
2	2.2/1.3	4.1/3.8	8.7/46.0	112.0/1401.4	3798.0/t.o.	t.o./t.o.				
3	6.5/5.9	7.6/18.3	13.2/136.3	122.1/t.o.	4376.5/t.o.	t.o./t.o.				
4	13.9/19.	15.5/52.8	21.0/252.7	123.3/t.o.	3874.2/t.o.					
5	21.1/29.4	22.2/64.2	28.6/197.2	135.5/t.o.						
6	17.8/19.9	18.2/34.5	24.7/77.5							
7	10.1/6.4	10.9/8.4								
8	8.3/1.4	PI				15				

t.o. timeout (two hours)

Our second evaluation involves modeling a fair *n*-sided dice using fair coins, following Knuth's algorithm (Knuth, 1976) featured in PRISM. This algorithm uses rejection sampling to minimize the expected coin tosses by simulating a uniform distribution over powers of two with fair coins. For example, in the case of a six-sided die (see Figure 4), three coin tosses yield 8 possibilities. If the resulting value of the tosses is smaller than six, the algorithm accepts it. Otherwise, the algorithm repeats the process as if starting with the remainder. More complex cases may involve multiple trees for simulation.

The BP model for this algorithm tracks coin flips through b-threads representing nodes in the sequence. Each node waits for its parent. Then, depending on its place in the tree, it either returns the dice value or performs a toss and requests the corresponding next node. This results in a simple model where all bthreads can be generated by a single function shown in Listing 7.

Results from the dice program analysis in Table 3 and Figure 5 align with our observations from smaller



Figure 3: A comparison of the sampling-based analysis and the exact analysis for the Monty Hall model with different values of doors (d), prizes (p), and doors opened (o).



Figure 4: An illustration of the process simulating a sixsided dice using coin flips.

```
@thread
def node(u, x): # u: layer size, x: index in layer
while True:
yield sync(waitFor=BEvent(f"n{u}_{x}"))
if u < n: # inner node
flip = yield choice({0:0.5, 1:0.5})
yield sync(request=BEvent(f"n{u*2}_{2*x+flip}")
else: # last layer
if x >= n:
yield sync(request=BEvent(f"n{u-n}_{x%n}"))
else:
yield sync(request=BEvent(f"result_{x}"))
```

Listing 7: B-thread definition of any single node in the BP model of Knuth's algorithm (Knuth, 1976).

Monty Hall versions, where the exact analysis outperforms sampling. One difference from the previous example was that translation, rather than computation, was the more time-consuming part of the exact analysis due to more b-threads and fewer commonly shared events. This example is notable due to the unlimited number of times it may repeat before yielding a result, making it particularly suited for exact analysis, which explores the full possibility space simultaneously. We can observe that in more complex cases like Figure 5b, the exact analysis still performs well due to oversampling advantages.



Figure 5: Comparison of the sampling-based and exact analysis for the *n*-sided dice model.

Table 3: Exact analysis of various sizes for a dice. Some require multiple trees which repeat the same structure.

n	States	Trans./Comp. time (s)	Result	n	States	Trans./Comp. time (s)	Result
6	64	0.129/0.21	0.1667	19	1868	46.931/6.056	0.0526
7	68	0.141/0.053	0.1429	20	512	3.179/0.418	0.05
8	38	0.132/0.053	0.125	21	440	3.009/0.399	0.0476
9	292	1.361/0.214	0.1111	22	1200	20.060/2.451	0.0455
10	216	0.821/0.147	0.1	23	968	12.796/1.522	0.0435
11	596	5.095/0.635	0.0909	24	376	1.592/0.233	0.0417
12	168	0.423/0.097	0.0833	25	2744	101.133/12.561	0.04
13	848	9.941/1.191	0.0769	26	1704	39.804/5.134	0.0385
14	144	0.440/0.097	0.0714	27	2668	96.446/12.325	0.0370
15	148	0.443/0.099	0.0667	28	368	1.664/0.239	0.0357
16	78	0.451/0.097	0.0625	29	4448	265.348/33.184	0.03448
17	752	8.300/0.999	0.0588	30	304	1.705/0.237	0.03333
18	592	5.272/0.663	0.0556				

Our last case study is a variation of the bit-flip scenario discussed in Section 3. In this variation, diagonals are initiated randomly, and the actions are taken at each step, where each consecutive row or column can be flipped. The goal here is to find a sequence that results in the board having the same (or as close as possible for odd-size boards) number of bits turned on and off.

The translation part in the exact analysis of the bit-flip model was more challenging than in previous cases, starting at relatively small matrix sizes. This is due to the large number of board configurations in this scenario. Each of the possible $2^{n \times m}$ possibilities adds both an event and a state in each b-thread, resulting in an exponential number of options that must be explored in all of the b-threads during translation. However, when considering all b-threads as a unified b-program during exact analysis, the number of possible outcomes is relatively small.

We conducted a sampling-based analysis by measuring 9,999 runs or as many as possible within an hour, as seen in Table 4. Sampling outperformed exact analysis, supporting double the translation's grid size before slowing down. We suspect the difficulties stem from the sync events being a product of relatively costly computations. Section 3 and Section 7 present an alternative analysis approach to the discrete event selection mechanism used here by integrating SMT solvers, which aims to address this issue.

Table 4: Sampling results of the discrete bit-flip model with: average, standard measurement error, number of samples (Samp.), and average time per sample in seconds (Time).

-			•	-					
Dim	Mean	SEM	Sam.	Time	Dim	Mean	SEM	Sam.	Time
2x2	0.263	0.004	9999	0.002	3x3	0.441	0.005	9999	0.023
2x3	0.264	0.004	9999	0.004	3x4	0.178	0.004	9999	0.295
2x4	0.309	0.005	9999	0.008	3x5	0.329	0.014	1062	3.39
2x5	0.309	0.005	9999	0.041	3x6	0.197	0.052	61	59.934
2x6	0.241	0.004	9999	0.237	3x7	0.4	0.25	5	885.183
2x7	0.124	0.007	2476	1.454	4x4	0.209	0.024	292	12.34
2x8	0.081	0.013	456	7.903	4x5	0.143	0.167	7	554.663
2x9	0.13	0.051	46	79.496					

The presented examples cover a wide range of parameter values and diverse model types, spanning from models with few states and minimal transitions, such as the dice problem, to highly interconnected ones in bit-flip. Our methods show promise in different problem domains, especially in smaller systems.

Modeling probabilistic systems with both nondeterministic and probabilistic behavior has been addressed in various ways previously. Some models focus on representing systems with probabilistic transitions (Puterman, 1990; Stoelinga, 2004), while others are designed to represent systems with non-deterministic choices (Rabin and Scott, 1959; Savitch, 1970). For systems exhibiting both nondeterministic and probabilistic behavior, models such as PRISM (Kwiatkowska et al., 2011) are commonly used. These models allow for a rich representation of system dynamics, enabling the analysis and verification of system properties. The advantages of modeling randomness in BP have not been explored before. The only reference to randomness we know of is in event selection strategies (Harel et al., 2010).

$6 \quad \mathbf{BP} \Leftrightarrow \mathbf{DRL}$

This section explores the interplay between BP and deep reinforcement learning (DRL), focusing on how it can enhance the alignment of requirements with BP modules for more effective execution. We illustrate how DRL's advanced data processing and interpretation abilities can aid in designing complex systems in BP. Subsequently, we discuss how BP's structured way of encoding rich behavioral specifications can enhance DRL-driven solutions. This is demonstrated using BPpy's integration with Gymnasium (Towers et al., 2023), a widely recognized API standard for reinforcement learning (RL) environments.

The combination of RL and BP was first introduced by (Eitan and Harel, 2011), which enhanced the semantics of live sequence charts (LSC), a visual language for BP, by incorporating reinforcements and applying learning algorithms. In recent years, this combination has gained growing interest, with several studies conducted with the assistance of BPpy. These studies explored different ways this learning mechanism can be combined with BP's execution (Elyasaf et al., 2019; Yaacov et al., 2024).

To illustrate how DRL can enhance the efficiency of execution mechanisms in BP, we begin with an evaluation of the Blueberry Pancake Maker example presented in (Bar-Sinai, 2020). In this example, a pancake batter is made from dry and wet mixtures using a mixer controlled by a b-program. The first two b-threads in Listing 8 prepare batter for *n* pancakes by adding *n* mixtures without specifying the order, which can lead to improper batter thickness. To prevent damage to the mixer, we need to control the order of mixture additions. The thickness_meter b-thread monitors thickness changes, while the range arbiter b-thread ensures the thickness stays within [-b,b]. Next, we add blueberries to the mix. To prevent the blueberries from bursting during the mixing process, 75% of the total mixture of batter should be added, and the batter itself should be relatively thin. The bthreads enough_batter and batter_thin_enough block the blueberry addition event until these conditions are met. Depending on the sequence in which the dry and wet mixtures are added, scenarios may arise where the AddBlueberries event, requested by the blueberries b-thread, remains blocked by batter thin enough throughout the program's execution. Such a situation would inevitably violate the system requirements, as it would result in a pancake mixture with no blueberries.

To address the risk of missing blueberries, one option is to define a specific sequence for their addition. However, this becomes progressively more complicated with larger systems and various requirements. For instance, if dry and wet mixtures are not consistently ready for addition, determining a strict sequence becomes more challenging and prone to errors. Thus, our aim is to design flexible b-programs to allow for all valid behaviors.

To address the complexity in specifying that blueberries must eventually be added, we used the integration between BPpy and Gymnasium introduced in (Yaacov, 2023). This integration encapsulates a bprogram as an RL environment, allowing developers to delineate the program's target objectives or optimization criteria by incorporating the localReward

```
@thread
def add_dry_mixture(n):
           range(n):
 yield sync(request=DryMixture())
@thread
def add_wet_mixture(n):
 for i in range(n):
 yield sync(request=WetMixture())
@thread
def thickness_meter():
while True:
    e = yield sync(waitFor=mixture_add)
  if e == DryMixture()
   yield sync(request=ThicknessUp(), block=mixture_add)
   yield sync(request=ThicknessDown, block=mixture_add)
@thread
def range_arbiter(b):
   thickness = 0
 while True:
    e = yield sync(waitFor=any_thick)
 thickness += 1 if e == ThicknessUp() else -1
if abs(thickness) >= b:
block_e = DryMixture() if thickness>0 else WetMixture()
   yield sync(block=block_e, waitFor=mixture_add)
  else:
  vield sync(waitFor=mixture add)
def blueberries():
yield sync(request = AddBlueberries())
def enough batter(n):
for j in range(int((n * 3) / 2)):
 vield sync(waitFor=mixture add, block=AddBlueberries())
@thread
def batter_thin_enough(n):
 thickness = 0
 while True:
  if thickness >= 0:
  e=yield sync(waitFor=any_thick, block= AddBlueberries())
   e=yield sync(waitFor=any_thick)
 thickness += 1 if e.name=="ThicknessUp" else -1
```

Listing 8: The *Blueberry Pancake Maker* b-program (Bar-Sinai, 2020).

parameter. This parameter can be integrated into any yield statement, complementing traditional requirements and facilitating learning through RL algorithms. Embedding these reward-based criteria within BPpy streamlines the modeling of intricate systems, fostering more intuitive and efficient development practices. Listing 9 demonstrates how rewards are added to the blueberries b-thread, reflecting the system's preference for eventually adding blueberries.

@thread def blueberries(): yield sync(request= AddBlueberries(), localReward=-0.0001) yield sync(waitFor=All(), localReward=1)

Our evaluation focused on the task of finding a deterministic event selection mechanism to generate a single execution trace adhering to system requirements. We compared two mechanisms. The first uses a DRL algorithm, wherein the state space is represented by local variables of the b-threads, serving as input to the agent's neural network. The learning algorithm was aborted upon successful task completion, producing a correct sequence of events in its predictions. Specifically, we employed the Maskable PPO algorithm (Huang and Ontañón, 2022) implemented in the Stable Baselines3 package (Hill et al., 2018), with a standard multilayer perception (MLP) network with two hidden layers of size 64. The second mechanism we evaluated uses an explicit program synthesis approach, in which a valid execution trace is found by exploring the program's state space. This is considered the current standard approach, as it has been discussed previously, and is supported by existing tools.

The runtime and memory usage of the two mechanisms for finding a single execution trace are presented in Table 5. Tests were conducted with $n \in$ {200, 300, 400, 500} and $b \in$ {25, 50, 75, 100}. As expected, the runtime and memory of the synthesis approach increase with the problem size, especially *n*. In contrast, the DRL approach's runtime and memory remain relatively stable. This can be attributed to the fact that increasing *n* and *b* doesn't change the number of program variables, keeping the neural network complexity constant. This demonstrates that DRL can be particularly beneficial when system complexity is due to variable value range rather than quantity.

Table 5:	Comparison	of average	(over 10	0 repetition	is) of
runtime a	nd memory of	f the DRL a	nd progr	am synthes	is ap-
proaches	for finding a s	single valid	executio	n trace.	

	Time ¹		Memory ²				Tin	ne ¹	¹ Memor		
n	b	DRL	Syn.	DRL	Syn.	п	b	DRL	Syn.	DRL	Syn.
200	25	17	123	1.79	0.06	400	25	16	555	1.79	0.11
	50	19	142	1.79	0.06		50	18	569	1.79	0.12
	75	19	149	1.79	0.06		75	16	558	1.79	0.12
	100	18	163	1.79	0.06		100	17	591	1.79	0.12
300	25	16	333	1.79	0.08		25	15	894	1.79	0.15
	50	20	312	1.79	0.08	500	50	18	927	1.79	0.15
	75	19	339	1.79	0.09	300	75	19	873	1.79	0.15
	100	20	333	1.79	0.09		100	16	903	1.79	0.16
	100	20	555	1.77	0.07		100	10	705	1.77	0.10

¹ in seconds, ² in GB

We also evaluated the *Cinderella-Stepmother problem* from Section 3. In this variation, the DRL mechanism needs to effectively handle a counterstrategy of the stepmother described in (Shevrin and Yossef, 2020), where in each round, the stepmother maintains two buckets, which Cinderella cannot empty at once, evenly filled. Once the stepmother reaches a round in which one of the buckets can be entirely filled, it does so and wins. The results of this example experiment are provided as an appendix in the supplementary material. The evaluation yielded results similar to the pancake maker example. The appendix also contains early results for the task of find-

Listing 9: The updated blueberries b-thread which the localReward parameter.

ENASE 2025 - 20th International Conference on Evaluation of Novel Approaches to Software Engineering



Figure 6: Mean reward of the PPO algorithm for the bit-flip two-player game on square matrices of sizes 3×3 , 4×4 , 5×5 , and 6×6 . The results are compared with random and greedy baseline strategies.

ing a non-deterministic event selection mechanism that adheres to system requirements for the two examples. The results show that in both cases, the algorithms converged to accurate strategies after a short training time.

The evolving DRL field presents challenges in the safety, robustness, and interpretability of policies. Previous research suggests using BP to encode expertise for enhancing training and system reliability (Yerushalmi et al., 2023; Ashrov and Katz, 2023). We note that BPpy supports knowledge encoding by modifying rewards, features, and actions before relaying them to the learning algorithm or environment, enhancing system safety and ease of training.

7 **BP** \Leftrightarrow (**DRL+Probabilities+SMT**)

In the previous sections, we saw four integrations of BP with other mechanisms. As our vision is to use BP as a unifying modeling language suited for combining these methods, we demonstrate the interplay of the different mechanisms described in the preceding sections. Specifically, we show how the integrations with SMT solvers, probabilistic modeling, and DRL can be used together to create a comprehensive environment that facilitates both modeling and analysis.

For this demonstration, we use a variation of the *bit-flip* problem described earlier: A two-player game over a Boolean matrix of dimensions $N \times M$. The initial state of the matrix resembles a chessboard with alternating positive and negative bits. Our opponent randomly flips rows and columns, while our objective is to devise a strategy that involves strategically flipping rows or columns to turn on as many bits as possible *simultaneously*. We can randomly flip bits or implement a greedy strategy that selects the row or column with the highest number of bits to turn on in each round. However, as we will demonstrate later, a better strategy that is more challenging to implement manually can be achieved using DRL.

Section 3 and Section 5 demonstrate how the discrete implementation of the problem can be challenging for execution and analysis as the matrix size increases. This section uses a solver-based implementation combined with the choice idiom presented in Section 5 to model the opponent's behavior. The complete program used for the following evaluation, along with a description of the code, is available in an appendix in the supplementary material.

We ran the PPO algorithm (Huang and Ontañón, 2022) to learn a strategy for the task. The learned strategy was compared against two baseline strategies. The first strategy randomly selects rows and columns to flip, while the second greedily selects the row/column with the highest number of bits to turn on in each round. The mean expected reward of these strategies was computed using the choice idiom. Results are presented in Figure 6. We observe that the PPO algorithm rapidly obtained a strategy that can plan ahead and achieve better (or equal in the 3×3 case) results than the greedy approach in all matrix sizes examined.

The bit-flip demonstration shows how BP can serve as a "Swiss army knife" for modeling and analysis: Starting with a problem that was challenging to execute in a discrete program and upgrading it to utilize solvers. Subsequently, DRL can be applied to achieve a more effective strategy, and it can be evaluated and compared using probabilistic analysis.

The integration of SMT, probabilistic model checking, DRL, and BP holds great potential, as the first three are already closely linked. Many probabilistic model checking problems are solved using SMT-based constraint solving or linear optimization, while DRL algorithms address scalability issues in these problems. Tools like PRISM and STORM demonstrate this synergy, combining solvers and DRL. Extending this integration to BP could enable adaptive decision-making with formal guarantees for complex systems.

8 CONCLUSION

In conclusion, our study focused on integrating BP with a variety of techniques to establish a comprehensive framework for specifying and analyzing reactive systems. Moving forward, our future work will delve deeper into use cases involving multiple integrations and explore more complex and extensive examples. Additionally, we plan to focus on evaluating the usability of our proposed framework and assess how well it aids in programming to increase adoption and enhance accessibility.

ACKNOWLEDGEMENTS

This work of Weiss, Yaacov, and Zisser was partially supported by funding from the Israel Science Foundation (ISF) grant number 2714/19. The work of Ashrov and Katz was partially funded by the European Union (ERC, VeriDeL, 101112713). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

- Anderson, M. and Feil, T. (1998). Turning lights out with linear algebra. *Mathematics Magazine*.
- Ashrov, A. and Katz, G. (2023). Enhancing Deep Learning with Scenario-Based Override Rules: a Case Study. In *MODELSWARD*.
- Bar-Sinai, M. (2020). Extending Behavioral Programming for Model-Driven Engineering. PhD Thesis, Ben-Gurion University of the Negev, Israel.
- Bar-Sinai, M. and Weiss, G. (2021). Verification of Liveness and Safety Properties of Behavioral Programs Using BPjs. In Leveraging Applications of Formal Methods, Verification and Validation: Tools and Trends.
- Bar-Sinai, M., Weiss, G., and Shmuel, R. (2018). BPjs: an extensible, open infrastructure for behavioral programming research. In *MODELS*.
- Bodlaender, M. H. L., Hurkens, C. A., Kusters, V. J., Staals, F., Woeginger, G. J., and Zantema, H. (2012). Cinderella versus the wicked stepmother. In *Theoretical Computer Science*.
- Busard, S. and Pecheur, C. (2013). PyNuSMV: NuSMV as a Python Library. In NASA Formal Methods.
- De Moura, L. and Bjørner, N. (2008). Z3: An efficient SMT solver. In *TACAS*.
- Depuydt, L. and Gill, R. D. (2012). Higher Variations of the Monty Hall Problem (3.0 and 4.0) and Empirical Def-

inition of the Phenomenon of Mathematics, in Boole's Footsteps, as Something the Brain Does. *Advances in Pure Mathematics*, 02.

- Eitan, N. and Harel, D. (2011). Adaptive behavioral programming. In *ICTAI*.
- Elyasaf, A. (2021). Context-Oriented Behavioral Programming. Information and Software Technology, 133.
- Elyasaf, A., Sadon, A., Weiss, G., and Yaacov, T. (2019). Using Behavioural Programming with Solver, Context, and Deep Reinforcement Learning for Playing a Simplified RoboCup-Type Game. In *MODELS*.
- Elyasaf, A., Yaacov, T., and Weiss, G. (2023). What Petri Nets Oblige us to Say: Comparing Approaches for Behavior Composition. *IEEE Transactions on Software Engineering*, 49.
- Harel, D., Kantor, A., Katz, G., Marron, A., Mizrahi, L., and Weiss, G. (2013). On composing and proving the correctness of reactive behavior. In *EMSOFT*.
- Harel, D., Lampert, R., Marron, A., and Weiss, G. (2011). Model-checking behavioral programs. In *EMSOFT*.
- Harel, D., Marron, A., and Weiss, G. (2010). Programming Coordinated Behavior in Java. In *ECOOP*.
- Harel, D., Marron, A., and Weiss, G. (2012). Behavioral programming. *Communications of the ACM*, 55.
- Hensel, C., Junges, S., Katoen, J.-P., Quatmann, T., and Volk, M. (2022). The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer*, 24.
- Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. (2018). Stable Baselines.
- Huang, S. and Ontañón, S. (2022). A Closer Look at Invalid Action Masking in Policy Gradient Algorithms. In FLAIRS.
- Katz, G., Marron, A., Sadon, A., and Weiss, G. (2019). Onthe-fly construction of composite events in scenariobased modeling using constraint solvers. In *MODEL-SWARD*.
- Knuth, D. (1976). The complexity of nonuniform random number generation. Algorithm and Complexity, New Directions and Results.
- Kwiatkowska, M., Norman, G., and Parker, D. (2011). PRISM 4.0: Verification of Probabilistic Real-Time Systems. In CAV.
- McMillan, K. L. (1993). The SMV System. Symbolic Model Checking.
- Naveed, H., Arora, C., Khalajzadeh, H., Grundy, J., and Haggag, O. (2024). Model driven engineering for machine learning components: A systematic literature review. *Information and Software Technology*, 169.
- Puterman, M. L. (1990). Markov decision processes. In Stochastic Models, volume 2. Elsevier.
- Qin, X., Bliudze, S., Madelaine, E., Hou, Z., Deng, Y., and Zhang, M. (2020). Smt-based generation of symbolic automata. *Acta Informatica*, 57.
- Rabin, M. O. and Scott, D. (1959). Finite automata and their decision problems. *IBM Journal of Research and Development*, 3.

- Savitch, W. J. (1970). Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4.
- Shevrin, I. and Yossef, M. (2020). Spectra example: Cinderella-stepmother problem.
- Stoelinga, M. (2004). An Introduction to Probabilistic Automata. Bulletin of the EATCS, 78.
- Towers, M., Terry, J. K., Kwiatkowski, A., Balis, J. U., Cola, G., Deleu, T., Goulão, M., Kallinteris, A., KG, A., Krimmel, M., Perez-Vicente, R., Pierré, A., Schulhoff, S., Tai, J. J., Tan, A. J. S., and Younis, O. G. (2023). Gymnasium.
- Yaacov, T. (2023). BPpy: Behavioral programming in Python. *SoftwareX*, 24.
- Yaacov, T., Elyasaf, A., and Weiss, G. (2024). Keeping Behavioral Programs Alive: Specifying and Executing Liveness Requirements. In International Requirements Engineering Conference.
- Yerushalmi, R., Amir, G., Elyasaf, A., Harel, D., Katz, G., and Marron, A. (2023). Enhancing Deep Reinforcement Learning with Scenario-Based Modeling. SN Computer Science, 4.