

Combining Large Language Models and Gradient-Free Optimization for Automatic Control Policy Synthesis

Carlo Bosio¹, Matteo Guarrera², Alberto Sangiovanni-Vincentelli², Mark W. Mueller¹

Abstract—Large Language models (LLMs) have shown promise as generators of symbolic control policies, producing interpretable program-like representations through iterative search. However, these models are not capable of separating the functional structure of a policy from the numerical values it is parametrized by, thus making the search process slow and inefficient. We propose a hybrid approach that decouples structural synthesis from parameter optimization by introducing an additional optimization layer for local parameter search. In our method, the numerical parameters of LLM-generated programs are extracted and optimized numerically to maximize task performance. With this integration, an LLM iterates over the functional structure of programs, while a separate optimization loop is used to find a locally optimal set of parameters accompanying candidate programs. We evaluate our method on a set of control tasks, showing that it achieves higher returns and improved sample efficiency compared to purely LLM-guided search. We show that combining symbolic program synthesis with numerical optimization yields interpretable yet high-performing policies, bridging the gap between language-model-guided design and classical control tuning. Our code is available at <https://sites.google.com/berkeley.edu/colmo>.

I. INTRODUCTION

Designing control policies that are both high-performing and interpretable remains a central challenge in robotics and control [1]. Classical control methods such as PID or optimal control provide guarantees, but they typically require expert modeling and significant manual effort. At the same time, modern machine learning techniques—most notably deep reinforcement learning (RL)—can discover highly performant policies from data, but often produce opaque neural networks that are difficult to interpret, verify, or deploy in safety-critical systems.

Recent work has explored the use of LLMs as generators of symbolic control policies, framing the problem as a form of program synthesis guided by reinforcement feedback. By iteratively generating and evaluating candidate programs in a genetic algorithm framework, it is possible to produce compact, human-readable controllers [2]. This line of work provides a computational tool to automatically synthesize interpretable and modular control strategies.

However, a key limitation of current LLM-driven approaches is their reliance on the LLM to simultaneously specify both structure and numerical parameters (e.g., control gains, thresholds). LLMs are good at producing syntactically correct and qualitatively meaningful code, but they are not

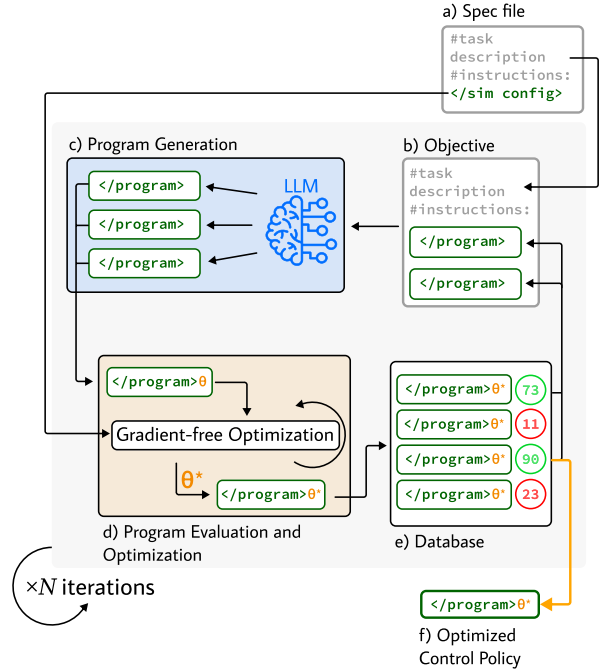


Fig. 1: Block diagram of the algorithmic infrastructure for our policy search method. The input to the algorithm is a specification file a) containing a task description, the implementation of an evaluation function to score programs, and some starter code to evolve. A prompt b) is constructed pasting the current best programs (the starter code at the beginning). The prompt is fed to a *Program Generation* block c) containing a pre-trained LLM, which produces more programs. The control policies generated by the LLM are fed to the *Program Evaluation and Optimization* block d), which optimizes them and scores them based on their closed-loop performance in simulation. The program-score pairs are stored in a *Database* e), from which they are sampled to be included in following prompts and improved. The output is a high-performance, programmatic control policy f).

well-suited for choosing precise numerical values. As a result, many generated policies that would perform well if associated to the right set of parameters, in reality perform poorly until enough evolutionary steps stumble upon suitable values—an inefficient and potentially brittle process.

With our hybrid symbolic-numeric search, we propose to decouple program structure generation from parameter tuning by introducing a numerical optimization stage in the loop with the program search process. In our approach (see Fig. 1), the LLM-generated programs are processed to replace numerical parameters with symbolic placeholders, and are then optimized using a zeroth-order numerical optimizer. This hybrid approach allows:

- The LLM to generate functional structure of candidate controllers.
- The optimizer to exploit gradient-free search to find a

UC Berkeley, Department of Mechanical Engineering¹ and Department of Electrical Engineering and Computer Sciences², Contact: {c.bosio, matteogu, alberto, mwm}@berkeley.edu

locally optimal set of parameters.

We show that this combination leads to higher performances and improved sample efficiency (i.e., number of search iterations to reach a target performance) compared to pure LLM-driven search. Our results suggest that separating symbolic reasoning from numeric search is a powerful paradigm for automatic policy synthesis and could accelerate the adoption of programmatic controllers in real-world applications.

The contributions of this work are as follows:

- Proposing a novel symbolic-numeric method for automatic synthesis of programmatic control policies;
- Achieving over one order of magnitude fewer search steps in compared to purely LLM-driven search to achieve target performances;
- Conducting extensive testing on control tasks, including full-scale quadruped locomotion;
- Open-sourcing our framework based on low compute budget requirements.

II. RELATED WORK

In the following we present some of the previous work which motivated the development of this idea.

A. Program synthesis and Large Language Models

Program synthesis [3] is the task of automatically finding a program expressed in some underlying programming language and satisfying user intents and constraints. These methods have been limited by the large dimensions of the search space [4]. To make the search more tractable, a common strategy is to use manually designed Domain Specific Languages (DSLs) [5], [6], which are application-specific. Another approach is the use of sketches, where partial program structures are provided by the user and then completed by the synthesis algorithm [7]. LLMs for code [8]–[10], have been used as a DSL-free alternative to program synthesis. On top of LLMs for code generation, extensive system-level research has been proposed with the aim of integrating LLMs with additional components to achieve more complex tasks than with single prompt engineering. A popular example is the combination of LLMs with evaluators, i.e. programmatic ways of scoring their outputs, and evolution strategies such as genetic algorithms [11]. This paradigm has been applied in various contexts, such as automated reasoning [12], code debugging [13], and algorithm design [14], [15].

B. Interpretable and Programmatic Reinforcement Learning

There has been increasing interest within the machine learning community in intersections between program synthesis and data-driven methods [16]. In fact, programmatic and symbolic representations improve the interpretability of learned functions and policies. For example of these works are learning verifiable decision-tree policies [17], and programmatic RL [18], [19] where policies are synthesized from a domain-specific language. These methods offer interpretability but require carefully designed DSLs or templates, limiting expressivity. A DSL-free approach to control policy

synthesis using LLMs was proposed in [2], where a genetic search framework is used to iteratively synthesize controllers. However, all these LLM-based approaches bypass the numerical aspect of control design, relying on the discrete program search to guess suitable parameter values. This leads to inefficient exploration of the continuous parameter space and often yields suboptimal performance. Our method, by incorporating in-the-loop numerical optimization, introduces a systematic local search in the parameter space, complementing program synthesis.

C. Gradient-free Optimization in Control Design

Gradient-free optimization (GFO) [20] is a family of methods for function optimization in cases where derivatives are unavailable or unreliable. In control design, GFO methods have been widely applied to tasks such as auto-tuning of PID gains or model predictive control cost weights [21]–[23] in simulation-in-the-loop or hardware-in-the-loop settings. Such optimizers typically require a fixed control structure, and thus cannot discover fundamentally new strategies. Our work integrates these optimizers into a structure-search loop, enabling them to tune parameters for policy candidates automatically generated by an LLM-driven search framework. Our work draws inspiration from GFO-based control tuning to find local optima in the parameter space for the LLM-generated policies.

D. Hybrid Symbolic–Numeric Methods

Our work is related to recent efforts combining symbolic reasoning with numerical search. For example, symbolic regression systems (e.g., [24], [25]) discover analytic expressions from data and then optimize their parameters. In the context of LLMs, recent studies have explored LLMs generating differentiable code for physics simulation that is optimized via gradient-based methods [26]. Related work has also employed LLMs for automatic design of electronic circuits [27]. We extend this symbolic-numeric paradigm to control synthesis. In this work, we bring these symbolic-numeric ideas to develop an automatic synthesis tool for programmatic controllers.

III. METHODOLOGY

Our approach enables developing interpretable, high-performance policies for the control tasks of interest. We first briefly describe the problem fundamentals, following the formalism of [28], and then the algorithmic aspects of our approach.

A. Background

We wish to find a control policy for a discrete-time dynamical system with dynamics in the form

$$x_{t+1} = f(x_t, u_t), \quad (1)$$

where $t \in \mathbb{N}$ is the time index, $x_t \in \mathbb{R}^n$ is the state of the system at time step t , and $u_t \in U \subset \mathbb{R}^m$ is the control input. At each time step t , a stage reward

$$r_t = g(x_t, u_t) \quad (2)$$

is incurred. The general objective is to find a control policy $u_t = h(x_t)$ to maximize the cumulative reward

$$R = \sum_{t=0}^T r_t, \quad (3)$$

where T is a specified time horizon. In [2], the function $h(\cdot)$ is represented as a non-parametric program $\text{policy}(\cdot)$ in Python, i.e.

$$u_t = \text{policy}(x_t). \quad (4)$$

The goal is then to find a controller $\text{policy}^*(\cdot)$ by approximating the solution to the following optimization problem:

$$\begin{aligned} \text{policy}^*(\cdot) = \max_{\text{policy}(\cdot)} R \\ \text{s.t. } x_{t+1} = f(x_t, u_t), \forall t \in \{0, \dots, T\} \\ u_t = \text{policy}(x_t), \end{aligned} \quad (5)$$

where R is defined in eq. 3. The search in this case does not happen in a parameter space encoding a mathematical structure (such as, for instance, a neural network), but directly in the space of programs, which is infinite-dimensional and complex to search over for the reasons discussed in Section II-A. An LLM is used in [2] to guide the program search.

B. Parametric program approach

The use of an LLM shifts the optimization to the space of tokens, i.e. the elements LLMs use to decompose a string of text in elementary units. This formalization leads to some results in simple settings. However, it does not provide the numerical stability and robustness that are needed in control, where a small deviation in parameters can lead to numerical instabilities. Furthermore, LLMs have notoriously struggled with simple math tasks and multi-step numerical reasoning, often producing incorrect results. Additionally, LLMs are not designed to operate on floating point numbers, which can make this vanilla search approach brittle. These limitations motivate separating structural synthesis from numerical optimization.

We propose an improvement of the formulation (5), in which the control program is parametric, i.e.

$$u_t = \text{policy}(\theta; x_t), \quad (6)$$

where the parameters θ are all the continuous parameters that can be found in a controller, such as control gains, weights, thresholds. The values produced by an LLM are likely to be suboptimal, or not tuned for the specific task of interest. We alleviate this problem by adding an inner GFO loop optimizing the parameters θ :

$$\begin{aligned} \text{policy}^*(\cdot) = \max_{\text{policy}(\cdot) \in \Phi} \max_{\theta} R \\ \text{s.t. } x_{t+1} = f(x_t, u_t), \forall t \in \{0, \dots, T\} \\ u_t = \text{policy}(\theta; x_t), \end{aligned} \quad (7)$$

Compared to (5), the new formulation (7) consists of a nested optimization problem. The outer layer optimizes

over the functional structure of programs, while the inner layer finetunes the programs finding the optimal numerical parameters. This enables to reduce the number of LLM-generated samples required to find high-performing policies. The practical realization of these high-level ideas is inherently nontrivial, owing to several challenges that we detail in the following sections.

IV. IMPLEMENTATION DETAILS

The overall algorithmic infrastructure is shown in Fig. 1. The input is a specification file, where a description of the task, some starter code for the function to evolve and the implementation of the evaluation function are provided. At each iteration, a *Program Generation* block (containing the LLM) produces candidate control programs for the task of interest. The proposed control policies are then fed to a *Program Optimization* block, which simulates the underlying system in closed loop and optimizes the candidate program's parameters. The best performing programs are stored in a *Database* and added to the prompts for the subsequent iterations, where the LLM is instructed to improve upon the previously generated programs. In the following we explain in more detail the functioning of this framework. In the following sections, we briefly describe the algorithm components, with particular emphasis on the novel parts, and we refer to [2] for additional details on the base infrastructure.

A. Specification and Prompt Construction

The input to our control synthesis framework is a specification file (Fig. 1a), composed by three main parts:

- A description of the control task to solve;
- An initial candidate control program in the form of starter code;
- The evaluation function used by the program optimization block (Fig. 1d).

An example of a general specification structure for our control synthesis method is shown in Fig. 2.

At each iteration, a prompt (Fig. 1b) is constructed by concatenating two previously generated high-performing programs to the task description from the specification file (see Fig. 3). The LLM is thus instructed to improve upon the two provided programs. During the initial phases of the algorithm execution, only the starter code (from the specification file) is available, and is directly pasted into the prompt. As the evolution progresses and higher-performing programs are generated, these are used in the prompt instead of the starter code. An example of a general prompt structure is shown in Fig. 3.

B. Program Generation

A pre-trained LLM is the generation engine of the *Program Generation* block (Fig. 1c), in which candidate control programs are sampled. The generation step from the provided prompt can be thought of as a genetic operator [29], since, among all the possible mechanisms that describe an LLM output based on its inputs, these possibilities can occur:

```

"""Finds a policy for the control task."""

# import libraries needed
import numpy as np

# evaluation function for the control policy
def evaluate(params):
    """Returns optimized params and score."""
    opt_params, score = Optimize(objective_fn)
    return opt_params, score

def objective_fn(params):
    """Returns the simulation reward."""
    environment = initialize_env()
    obs = environment.get_observation()
    total_reward = 0.0
    for _ in range(1000):
        action = policy(params, obs)
        reward, obs = environment.step(action)
        total_reward += reward
    return total_reward

# function to evolve
def policy(params, observation):
    """Returns a control action."""
    action = np.random.uniform()
    return action

```

Fig. 2: Example pseudo-code template for a control synthesis specification. The `objective_fn` function is wrapped into a GFO loop (in orange).

- The two programs provided in the prompt are mixed (i.e., *crossover*);
- One of the programs provided is modified (i.e., *random mutation*);
- An entirely new program is produced.

C. Program Evaluation

The evaluation block (Fig. 1d) is the core innovation of our method. Candidate control programs are automatically extracted from the LLM output and undergo string-level preprocessing operations. The numerical parameters in a program are likely not very accurate, since they are LLM-generated. To alleviate this, all the numerical values are parsed through a Regular Expression (RegEx) [30] and replaced with the substrings `params[0]`, `params[1]`, `params[2]`, etc. with progressive indices, and the policy signature is modified so to accept the `params` variable as an input. These operations make the program parametric. The program parametrization is then fed to the GFO optimizer, which iteratively evaluates the candidate program with refined sets of parameters. At each GFO iteration, the candidate controllers are simulated in closed loop with the task of interest. The performance is quantified by the reward (3) provided in the specification file. Syntactically incorrect or incomplete programs are discarded, while promising program-score pairs are stored in the *programs database*, from which they are sampled to be added to the subsequent prompts, and evolved. We found that in most cases the dimension of the parameter vector is less than 20, which falls within a suitable regime for GFO algorithms [20].

```

"""Finds a policy for the control task.
On every iteration, improve policy_v1 over the
policy_vX methods from previous iterations.
"""
import numpy as np

def policy_v0(observation):
    """Returns a control action."""
    action = np.random.uniform()
    return action

def policy_v1(observation):
    """Returns a control action."""
    action = 0.0
    return action

def policy_v2(observation):
    """ Improved version of policy_v1."""

```

Fig. 3: Example pseudo-code template for a prompt. The LLM generates a body for the provided function signature trying to improve upon previously generated functions. `policy_v0` and `policy_v1` are sampled from the database, the best performing policies are sampled more frequently.

D. Programs Database

The generated high-performing programs are stored in the *programs database* (Fig. 1e). Each program is assigned a unique identifier derived solely from its symbolic representation to make sure that programs differing only in their parameters but sharing the same functional structure are not redundantly stored. In this way, we avoid feeding back to the LLM two programs that are structurally equivalent. To discourage getting stuck in local optima, an island approach is implemented, where different instances of the program search are run independently [31]. Thus, when constructing a prompt, a two stage sampling procedure happens: first, an island is sampled, then programs contained within the selected island are sampled to be added to the prompt. More details on the database implementation can be found in [2] and [11].

E. Computational Considerations

A number of hyperparameters affect the generation performances of an LLM [2]. As a general rule of thumb, we found that a larger batch size (i.e., the number of output programs generated in parallel, from the same prompt) is always better, since it leads to more exploration. The LLM inference time scales sublinearly with the batch size thanks to GPU parallel inference, therefore it is always recommended to work with the largest batch sizes allowed by the available hardware. Adding a GFO optimization in the loop may sound like an increase in computational cost. However, given the significant time required for the LLM inference (in the order of tens of seconds), the GFO evaluation can be shadowed through multithreading: during LLM inference (on GPUs), the evaluation block running on CPUs evaluates and optimizes the previously generated program batch (see Fig. 4). This overlapping utilization ensures that no resource remains underused. As long as the evaluation is faster than the generation, the queue of programs remains bounded. Depending on the specific GFO algorithm used, the evaluation time can vary. We tested Bayesian Optimization




Fig. 4: Parallel implementation of program generation and evaluation. The program generation (in blue) happens on the GPUs and evaluation (in orange) happens on the CPUs. In this way, the computational cost of optimization is completely shadowed.

TABLE I: State and input space dimensions for each tasks. Performances of the synthesized policies for each task using only LLM-driven search and adding the in-the-loop GFO.

	State dim.	Action dim.	Reward (no GFO)	Reward (GFO)
Pendulum swing-up	2	1	533	591
Ball in cup	8	2	575	755
Cheetah	18	6	159.1	252
Quadruped	36	12	62.1	181.3
Unitree A1	36	12	33.8	152.5

(BO) [32] and Evolutionary Strategies (ES) [33]. BO is more sample efficient but slower, while ES approaches require more iterations to converge but are computationally cheaper. For our use case, we obtained the best results with ES.

V. RESULTS

In the following we provide more detail about the practical aspects of the implementation of our method. We then present the results obtained on a set of simulated control tasks.

A. Setup

All experiments were conducted using two NVIDIA GeForce RTX 3090 GPUs and a single AMD Ryzen Threadripper 3990X 64-Core Processor. The model used in the program generation block (Fig. 1c) is the 7-billion-parameter model from the Qwen 2.5 family of models for code [34]. To fully take advantage of the available hardware and maximize the batch sizes per LLM inference, we used 4-bits weight quantization [35] and model pipelining [36]. All GFO algorithm implementations were taken off-the-shelf from the open-source library Nevergrad [37]. The results we show are obtained running 100 iterations of the (1+1)-ES algorithm [33] in the GFO block (Fig. 1d). All tasks are simulated through the open-source simulator MuJoCo [38], either directly or through the Control Suite [39] library.

B. Case Studies

We test our framework on a set of standardized control tasks of increasing difficulty and spanning a wide range of dimensions of state and input spaces. We applied our method to find programmatic policies for the *pendulum swing-up*, *ball-in-cup* tasks, as well as three locomotion tasks:




Fig. 5: Visualizations of the systems used as benchmarking tasks. Each task’s state and action dimensions are reported in parentheses as S and A respectively.

dm_control cheetah, *dm_control quadruped*, and *Unitree A1* (see Fig. 5). The respective dimensions of state and input spaces are reported in Tab. I.

Pendulum swing-up: This task consists of controlling an input-constrained pendulum to reach the upright equilibrium. The maximum applicable torque is $1/6^{\text{th}}$ as required to lift it from a horizontal motionless position. Even if low-dimensional, this task is not easily solvable. The pendulum has to oscillate and accumulate enough energy to be able to swing to the upright configuration. For this task, the step reward (2) is defined as

$$r_t = \begin{cases} 1 & \text{if } |\theta| < 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where $\theta \in [-\pi, \pi)$ is the angle with respect to the upward equilibrium configuration.

Ball-in-Cup: The *ball-in-cup* system is composed of a two-dimensional double integrator (the cup) and a ball attached to it through a string (a unilateral distance constraint between the two objects). The system is planar. The task consists of finding a policy for the cup to catch the ball. In this case, the policy outputs reference horizontal and vertical positions, which are tracked by a lower-level linear controller. For this task, the step reward (2) is defined as

$$r_t = \begin{cases} 1 & \text{if ball inside cup} \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Cheetah: This is the simplest of the locomotion tasks. The task consists of controlling a planar biped to move forward by outputting joint angles, also tracked by lower-level linear controllers. The step reward (2) is defined as

$$r_t = \max(0, \min(1, v/10)), \quad (10)$$

where v is the horizontal velocity of the biped torso.

Quadruped: This is a three-dimensional locomotion task. The system to control is a four-legged robot, with three actuated joints per leg. The step reward function is the same as (10). The goal is to make the quadruped move forward.

Unitree A1: We also tested our framework to generate a gait for a simulated Unitree A1 robot. We used the official model from the MuJoCo Menagerie [40]. The Unitree A1 is one of the most widely used quadrupeds in robotics research, and is therefore the most realistic of the locomotion tasks. In this case, we decided to generate a gait that would make the robot walk backward. The locomotion reward is the same as (10), using the appropriate direction and sign for the velocity.




Fig. 6: Best-so-far curves of policy search runs with in-the-loop GFO and without, respectively in blue and red. The solid lines represent the best-so-far score across all islands, while the dashed lines represent the best-so-far score within each island.

When simulated, the dynamics equations are discretized through a semi-implicit Euler method (details in [38]). The controls are injected every 20 ms for the *pendulum swing-up*, *ball-in-cup* and *quadruiped* tasks, every 10 ms for the *cheetah*, and every 1 ms for the *Unitree AI*. All episodes are 1000 control steps long. In all cases, our approach:

- Shows higher performance compared to optimization-free program search, in some cases by margins of 100% or more;
- Automatically finds high performing static feedback policies in the form of (6) that are lightweight and interpretable.

We report best-so-far curves for the locomotion tasks in Fig. 6, and a quantitative comparison of the rewards achieved by the policies found for all tasks with in-the-loop GFO and without in Tab. I. For each task The performances in Tab. I for the *pendulum swing-up* and *ball-in-cup* tasks are averaged across 1000 episodes in which the initial conditions are randomized. In fact, these two tasks are more sensitive to initial conditions than the locomotion tasks. For *Cheetah*, *Quadruiped*, and *Unitree AI*, the performance is averaged across 10 episodes, which is the same number of episodes used to compute the programs’ score during the synthesis process. For these tasks, the policies achieve average forward torso velocities of $2.5 \text{ m} \cdot \text{s}^{-1}$, $1.8 \text{ m} \cdot \text{s}^{-1}$, and $-1.52 \text{ m} \cdot \text{s}^{-1}$ (backward walk), respectively.

With reference to Fig. 6, all the runs were carried out with a fixed wall-clock time budget: 2 hours for *pendulum swing-up* and *ball-in-cup*, 5 hours for *cheetah* and *quadruiped*, and 10 hours for *Unitree AI*. Had we allowed for more wall-clock time, even better policies would probably have been found by the search process. One more way to interpret the best-so-far curves is in a sample efficiency sense: taking a target score, i.e. a horizontal line in the graphs, how many generations will be needed approximately to achieve it. It is then clear that introducing the GFO loop greatly improves the sample efficiency of the search process. It is also interesting to note that the GFO seems to provide greater benefit for higher-dimensional state and action spaces, thus bringing these methods closer to be deployed in real-world systems.

We report the best policies found for the five tasks in

Fig. 7. In all cases, the found policies are below 35 lines of code, and can be read, edited, and debugged. In all cases except *quadruiped*, the policies contain fewer than 10 parameters. The comments and variable names in the snippets are produced by the LLM. By identifying what the variables represent and their bounds, the policy could be simplified even more. The shared language (code) allows a user to iterate in the loop with our synthesis algorithm as well. Videos showing the systems’ behaviors with the respective policies running in closed-loop can be found at <https://sites.google.com/berkeley.edu/colmo>.

VI. CONCLUSIONS

We presented a hybrid symbolic-numeric approach to automatically synthesize interpretable control policies combining LLM-guided program search with in-the-loop gradient-free optimization. By decoupling the synthesis of a control program from the selection of its parameters, our method leverages the complementary strengths of symbolic representations and numerical optimization methods.

Across a set of control benchmarks, we demonstrate that our approach achieves higher performance and faster convergence than purely LLM-driven search while maintaining the interpretability and modularity of the resulting policies. The policies we find are also simpler and carry orders of magnitude fewer parameters than neural network policies solving the same tasks through reinforcement learning. This compactness in terms of parameter count not only improves interpretability but also facilitates deployment on resource-constrained systems where memory and compute efficiency are critical. These results suggest that jointly exploring symbolic and numeric spaces leads to reliable and efficient discovery of control strategies.

By demonstrating promising results without resorting to prohibitively expensive computational resources, we showed the accessibility and practicality of our approach, which brings us closer to fast programmatic control design iterations and more efficient turnaround times when developing control policies. We believe our method is an important step toward automatically synthesizing interpretable, verifiable, and high-performing controllers for real-world systems.

This work opens several promising directions. Future research could extend this approach to more complex policy representations, such as time-varying functions or policies

with memory (e.g. access not only to the current observation, but also to past history) as well as differentiable code generation and gradient-based in-the-loop optimization.

```
def policy(obs):
    """obs size is 3. return shape is ()"""
    x1 = np.arctan2(-obs[1], obs[0])
    # Adjust action
    if np.abs(x1) > np.pi / 6.512:
        action = np.sign(obs[2])
    else:
        action = 9.437 * np.sin(x1) - 2.746 * obs[2]

    return action
```

(a) Pendulum swing-up.

```
def policy(obs):
    """obs size is 17: position (8), velocity (9).
    return shape is (6,)"""

    pos = obs[:3]
    vel = obs[3:6]
    ang_vel = obs[6:9]

    # Adjust forward control gain
    action = np.zeros(6)
    action[0] = -12.241 * vel[0]

    action[1] = -0.603 * ang_vel[0] # Hip angle
    action[2] = 11.351 * ang_vel[1] # Knee angle
    action[3] = -8.106 * ang_vel[2] # Ankle angle

    # Introduce a new control term for balance
    action[4] = -12.268 * ang_vel[1]

    return action
```

(b) Cheetah

```
def policy(obs):
    torso_upright = obs[47]
    imu = obs[48:54]

    if torso_upright < 1.131 or any(abs(x) > 0.529 for x
    in imu):
        # Adjust actions for better stability
        actions = np.array([-0.4, -0.499, 1.321,
        -1.483, 0.774, -0.281, 1.049, -0.836, -0.195,
        0.846, 1.898, 2.066])
    elif abs(torso_upright) < 0.825:
        # Adjust actions for a balanced stance
        actions = np.array([
        0.2, -0.493, -0.287, 0.861, -2.279, 0.03,
        -1.584, -0.079, 0.427, 0.233, 0.801, 0.163])
    else:
        actions = np.array([0.0, -0.173, -0.024,
        2.506, -0.82, 0.029, -1.302, 0.009, -1.243,
        0.057, -1.264, -1.961])

    if obs[-1] < -0.275: # Performance adjustment
        actions *= -1.29
    elif obs[-1] > 1.442:
        actions *= -1.893

    # Minimal noise addition for fine-tuning
    actions += np.random.normal(scale=0.04, size=12)

    return actions
```

(c) Quadraped

```
def policy(obs)
    """obs size is 8. return shape is (2,)"""

    ball_pos = obs[:3]
    cup_pos = obs[3:]
    action = np.zeros((2,))

    # Further refine horizontal movement
    if abs(ball_pos[0] - cup_pos[0]) > -12.071:
        action[0] += (cup_pos[0] - ball_pos[0]) * -5.66

    # Fine-tune vertical position correction
    if abs(ball_pos[2] - cup_pos[2]) > -7.763:
        action[1] += (cup_pos[2] - ball_pos[2]) * 1.023

    # Add slight bias to vertical movement
    action[1] += 0.531

    return action
```

(d) Ball-in-cup

```
def policy(obs):
    """obs size is 27: position and angles (19), velocities
    (8). return shape is (12,)"""

    joint_angles = obs[0:12]
    joint_velocities = obs[12:24]
    body_velocity = obs[24:27]

    target_joint_angles = joint_angles.copy()

    # Adjust the hip joints
    target_joint_angles[0::3] += -27.697 * np.sign(
        body_velocity[0])
    target_joint_angles[1::3] += -6.726 * np.sign(
        body_velocity[0])

    # Adjust the knee joints
    target_joint_angles[1::3] += 6.09 * joint_velocities
    [1::3]

    # Adjusting the shoulder joints
    target_joint_angles[2::3] += -14.423

    # Add a simple spring-like mechanism
    balance_factor = 17.69
    target_joint_angles[0::3] += balance_factor * (
        body_velocity[1] + body_velocity[2])
    target_joint_angles[1::3] += balance_factor * (
        body_velocity[1] - body_velocity[2])

    # Add additional balancing adjustments
    if body_velocity[2] < -32.662:
        target_joint_angles[0::3] += 10.462
        target_joint_angles[1::3] += -22.429

    # Adjustments to ensure proper leg movement
    target_joint_angles[0::3] += 14.31
    target_joint_angles[1::3] -= -49.337

    # Return the target joint angles
    return target_joint_angles
```

(e) Unitree A1

Fig. 7: Policies found by our synthesis algorithm for the control tasks considered. The parameters optimized by the GFO loop are highlighted in orange.

ACKNOWLEDGMENT

This work has been partially supported by the Cal-Next Solar Center, <https://cal-next.berkeley.edu/> and the UC Berkeley College of Engineering. The authors would like to thank Orr Paradise for the insightful conversations and the members of HiPeRLab (<https://hiperlab.berkeley.edu/members/>) and SCALE Research Lab (<https://scale.berkeley.edu/people/>) for their feedback.

REFERENCES

- [1] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint arXiv:1702.08608*, 2017.
- [2] C. Bosio and M. W. Mueller, "Synthesizing interpretable control policies through large language model guided search," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 583–590.
- [3] S. Gulwani, O. Polozov, R. Singh *et al.*, "Program synthesis," *Foundations and Trends® in Programming Languages*, vol. 4, no. 1-2, pp. 1–119, 2017.
- [4] S. Gulwani, "Dimensions in program synthesis," in *Proceedings of the 12th international ACM SIGPLAN symposium on Principles and practice of declarative programming*, 2010, pp. 13–24.
- [5] G. Cui, Y. Wang, W. Qiu, and H. Zhu, "Reward-guided synthesis of intelligent agents with control structures," *Proceedings of the ACM on Programming Languages*, vol. 8, no. PLDI, pp. 1730–1754, 2024.
- [6] N. Patton and K. Rahmani, "Program synthesis for robot learning from demonstrations," *arXivorg*, 2023.
- [7] A. Solar-Lezama, "The sketching approach to program synthesis," in *Asian symposium on programming languages and systems*. Springer, 2009, pp. 4–13.
- [8] D. Friedl, A. Aghajanyan, J. Lin, S. Wang, E. Wallace, F. Shi, R. Zhong, W.-t. Yih, L. Zettlemoyer, and M. Lewis, "Incoder: A generative model for code infilling and synthesis," *arXiv preprint arXiv:2204.05999*, 2022.
- [9] R. Li, L. B. Allal, Y. Zi, N. Muennighoff, D. Kocetkov, C. Mou, M. Marone, C. Akiki, J. Li, J. Chim *et al.*, "Starcoder: may the source be with you!" *arXiv preprint arXiv:2305.06161*, 2023.
- [10] A. Lozhkov, R. Li, L. B. Allal, F. Cassano, J. Lamy-Poirier, N. Tazi, A. Tang, D. Pykhtar, J. Liu, Y. Wei *et al.*, "Starcoder 2 and the stack v2: The next generation," *arXiv preprint arXiv:2402.19173*, 2024.
- [11] B. Romera-Paredes, M. Barekatin, A. Novikov, M. Balog, M. P. Kumar, E. Dupont, F. J. Ruiz, J. S. Ellenberg, P. Wang, O. Fawzi *et al.*, "Mathematical discoveries from program search with large language models," *Nature*, vol. 625, no. 7995, pp. 468–475, 2024.
- [12] E. Zelikman, Y. Wu, J. Mu, and N. Goodman, "Star: Bootstrapping reasoning with reasoning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 15 476–15 488, 2022.
- [13] P. Haluptzok, M. Bowers, and A. T. Kalai, "Language models can teach themselves to program better," *arXiv preprint arXiv:2207.14502*, 2022.
- [14] J. Lehman, J. Gordon, S. Jain, K. Ndousse, C. Yeh, and K. O. Stanley, "Evolution through large models," in *Handbook of Evolutionary Machine Learning*. Springer, 2023, pp. 331–366.
- [15] F. Liu, X. Tong, M. Yuan, X. Lin, F. Luo, Z. Wang, Z. Lu, and Q. Zhang, "An example of evolutionary computation+ large language model beating human: Design of efficient guided local search," *arXiv preprint arXiv:2401.02051*, 2024.
- [16] M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, and D. Tarlow, "Deepcoder: Learning to write programs," *arXiv preprint arXiv:1611.01989*, 2016.
- [17] O. Bastani, Y. Pu, and A. Solar-Lezama, "Verifiable reinforcement learning via policy extraction," *Advances in neural information processing systems*, vol. 31, 2018.
- [18] A. Verma, V. Murali, R. Singh, P. Kohli, and S. Chaudhuri, "Programmatically interpretable reinforcement learning," in *International conference on machine learning*. PMLR, 2018, pp. 5045–5054.
- [19] A. Verma, H. Le, Y. Yue, and S. Chaudhuri, "Imitation-projected programmatic reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [20] A. R. Conn, K. Scheinberg, and L. N. Vicente, *Introduction to derivative-free optimization*. SIAM, 2009.
- [21] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe, "Data-efficient autotuning with bayesian optimization: An industrial control study," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 3, pp. 730–740, 2019.
- [22] F. Berkenkamp, A. Krause, and A. P. Schoellig, "Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics," *Machine learning*, vol. 112, no. 10, pp. 3713–3747, 2023.
- [23] C. Sartore, M. Rando, G. Romualdi, C. Molinari, L. Rosasco, and D. Pucci, "Automatic gain tuning for humanoid robots walking architectures using gradient-free optimization techniques," in *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*. IEEE, 2024, pp. 996–1003.
- [24] M. Cranmer, "Interpretable machine learning for science with pysr and symbolicregression. jl," *arXiv preprint arXiv:2305.01582*, 2023.
- [25] S.-M. Udrescu and M. Tegmark, "Ai feynman: A physics-inspired method for symbolic regression," *Science advances*, vol. 6, no. 16, p. eaay2631, 2020.
- [26] P. Ma, T.-H. Wang, M. Guo, Z. Sun, J. B. Tenenbaum, D. Rus, C. Gan, and W. Matusik, "Llm and simulation as bilevel optimizers: A new paradigm to advance physical scientific discovery," *arXiv preprint arXiv:2405.09783*, 2024.
- [27] Y. Lai, S. Poddar, S. Lee, G. Chen, M. Hu, B. Yu, P. Luo, and D. Z. Pan, "Analogcoder-pro: Unifying analog circuit generation and optimization via multi-modal llms," *arXiv preprint arXiv:2508.02518*, 2025.
- [28] D. Bertsekas, *Reinforcement learning and optimal control*. Athena Scientific, 2019, vol. 1.
- [29] M. Mitchell, *An introduction to genetic algorithms*. MIT press, 1998.
- [30] J. Friedl, *Mastering regular expressions*. O'Reilly Media, Inc., 2006.
- [31] E. Cantú-Paz *et al.*, "A survey of parallel genetic algorithms," *Calculateurs paralleles, reseaux et systems repartis*, vol. 10, no. 2, pp. 141–171, 1998.
- [32] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.
- [33] N. Hansen, D. V. Arnold, and A. Auger, "Evolution strategies," in *Springer handbook of computational intelligence*. Springer, 2015, pp. 871–898.
- [34] B. Hui, J. Yang, Z. Cui, J. Yang, D. Liu, L. Zhang, T. Liu, J. Zhang, B. Yu, K. Lu *et al.*, "Qwen2. 5-coder technical report," *arXiv preprint arXiv:2409.12186*, 2024.
- [35] T. Dettmers, M. Lewis, Y. Belkada, and L. Zettlemoyer, "Llm.int8(): 8-bit matrix multiplication for transformers at scale," 2022. [Online]. Available: <https://arxiv.org/abs/2208.07339>
- [36] Y. Huang, Y. Cheng, A. Bapna, O. Firat, D. Chen, M. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu *et al.*, "Gpipe: Efficient training of giant neural networks using pipeline parallelism," *Advances in neural information processing systems*, vol. 32, 2019.
- [37] P. Bennet, C. Doerr, A. Moreau, J. Rapin, F. Teytaud, and O. Teytaud, "Nevergrad: black-box optimization platform," *Acm Sigevolution*, vol. 14, no. 1, pp. 8–15, 2021.
- [38] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
- [39] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq *et al.*, "Deepmind control suite," *arXiv preprint arXiv:1801.00690*, 2018.
- [40] K. Zakka, Y. Tassa, and MuJoCo Menagerie Contributors, "MuJoCo Menagerie: A collection of high-quality simulation models for MuJoCo," 2022. [Online]. Available: <http://github.com/google-deepmind/mujoco-menagerie>