

RoboMorph: Evolving Robot Morphology using Large Language Models

Kevin Qiu^{1,2}, Władysław Pałucki¹, Krzysztof Ciebiera¹, Paweł Fijałkowski¹,
 Marek Cygan^{1,3}, Łukasz Kuciński^{1,2,4}

¹University of Warsaw ²IDEAS NCBR ³Nomagic ⁴Polish Academy of Sciences
 kevinxqiu@gmail.com

Abstract—We introduce **RoboMorph**, an automated approach for generating and optimizing modular robot designs using large language models (LLMs) and evolutionary algorithms. Each robot design is represented by a structured grammar, and we use LLMs to efficiently explore this design space. Traditionally, such exploration is time-consuming and computationally intensive. Using a best-shot prompting strategy combined with reinforcement learning (RL)-based control evaluation, **RoboMorph** iteratively refines robot designs within an evolutionary feedback loop. Across four terrain types, **RoboMorph** discovers diverse, terrain-specialized morphologies, including wheeled quadrupeds and hexapods, that match or outperform designs produced by Robogrammar’s graph-search method. These results demonstrate that LLMs, when coupled with evolutionary selection, can serve as effective generative operators for automated robot design. Our project page and code are available at <https://robomorph.github.io>.

I. INTRODUCTION

Designing robot morphologies remains a fundamental challenge in modern robotics [1]. Traditional engineering approaches are time-consuming, heavily dependent on human intuition, and ill-suited to exploring the vast space of possible designs. Automating this process could significantly reduce development time and uncover solutions that surpass those conceived by human designers.

Large language models (LLMs) have recently transformed multiple areas of robotics, from generating robot policies as code [2] and enhancing semantic reasoning [3] to designing reward functions and domain randomization [4], [5]. Here, we apply these capabilities to the robot design problem itself.

Previous attempts to automate the design process have been computationally expensive, requiring extensive search over large design spaces. One class of such approaches is inspired by evolutionary algorithms (EAs) [6]. However, EA-based methods require hundreds of generations and struggle to scale with increasing design complexity [7]. Other methods, such as Robogrammar [8], introduce structured grammars and heuristic search, while recent latent optimization methods [9], [10] focus on guiding exploration with learned representations. Despite their successes, these approaches remain limited by fixed heuristics, narrow exploration, or computational bottlenecks.

We introduce **RoboMorph**, a framework that integrates LLM-based generation, evolutionary algorithms, robot grammars, and RL-based control evaluation. Unlike prior work,

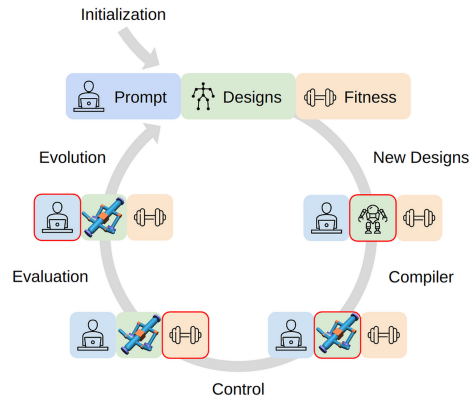


Fig. 1: One iteration of RoboMorph. At each stage, the prompt, robot design, and fitness score are displayed. The element modified in each step is highlighted in red. This closed-loop cycle enables progressive design refinement through evolutionary selection.

RoboMorph uses the compositional priors of LLMs as a mutation operator, enabling the discovery of morphologies that classical search heuristics overlook. By exploiting the reasoning capabilities of LLMs within the grammar, our method accelerates exploration of the design space and uncovers unconventional solutions beyond what human-engineered heuristics or latent models produce.

The grammar, adapted from Robogrammar [8], defines structural rules that ensure generated designs are physically plausible. At each generation, the LLM proposes new robot designs expressed in this grammar. These are compiled into XML models and simulated to train control policies, which determine each design’s fitness score. The highest-performing designs are retained and incorporated as few-shot exemplars in the next prompt, enabling iterative improvement. Figure 1 provides an overview of this process.

LLMs complement evolutionary algorithms in three ways. First, they explore the morphology space more efficiently than hand-coded mutation operators. Second, they provide a natural-language interface for specifying design constraints. Third, they generate structurally diverse candidates in a single step rather than through incremental perturbations.

While our experiments are conducted in simulation, manufacturing constraints and sim-to-real transfer remain open challenges. Nevertheless, the results demonstrate the feasi-




Fig. 2: Overview of the RoboMorph framework, which consists of a *design stage* and a *control stage*. The process is divided into six steps: (1) The input consists of the system prompt, user prompt, and few-shot examples. (2) Based on the input prompt, the LLM generates a new design and provides its reasoning. (3) A compiler converts the LLM-generated design into an XML file. (4) An RL control policy is trained in simulation for the robot in a given environment. (5) The robot is evaluated, and a fitness score is assigned to the design. (6) The population of robots is updated through evolution by pruning the worst-performing candidates and feeding the top designs back as examples for the next iteration. Together, these stages form an automated design loop that requires no human intervention after initialization.

bility of using LLMs as generative operators for morphology optimization. The designs discovered by RoboMorph exhibit distinctive traits that reflect the LLM’s capacity to recombine structural rules in unexpected ways.

In summary, our contributions are as follows:

- We introduce RoboMorph, the first framework to integrate LLM-based generation, evolutionary algorithms, robot grammars, and RL control for automated robot morphology design.
- We demonstrate that treating the LLM as a mutation operator yields competitive or superior designs compared to Robogrammar’s graph-search method, while producing more structurally diverse morphologies.
- We analyze design outcomes across four terrains, highlight the strengths and limitations of our approach, and outline directions for bridging simulation with real-world deployment.

II. ROBOMORPH

The RoboMorph framework, shown in Figure 2, follows an evolutionary pipeline that maintains a population of robot designs generated by an LLM (specifically, GPT-4o). At each iteration, the LLM is prompted with few-shot examples and produces new designs expressed in a grammar-based representation [8] (Sections II-A and II-B). The grammar enables diverse morphologies while ensuring outputs are physically valid and fabricable.

Each design is compiled into an XML file (Section II-C) and simulated in MuJoCo [11] via MJX (Section II-D) to train a control policy and compute a fitness score (Section II-E). The population is updated by retaining the top- K designs, which also serve as few-shot examples for the next iteration (Section II-F). The process is initialized with randomly generated examples and summarized in Algorithm 1.

A. Prompt Structure

We generate robot designs using GPT-4o with three main prompt components:

Algorithm 1 Pseudocode for RoboMorph

Requires:

$\text{prompt}_{\text{system}}$	robot grammar rules
$\text{prompt}_{\text{user}}$	instructions for designing a robot
evolutions	number of evolutions
population	number of designs per population
K	number of few-shot examples
Q	min-queue $\{(\text{fitness}_i, \text{design}_i, \text{reason}_i)\}_{i=1}^K$

function RoboMorph

$\text{fewshots} \leftarrow$ generate K random designs by sampling grammar rules

for $i = 1$ **to** evolutions **do**

for $j = 1$ **to** population **do**

$\text{design}_{\text{new}}, \text{reason}_{\text{new}} \leftarrow$ $LLM(\text{prompt}_{\text{system}}, \text{prompt}_{\text{user}}, \text{fewshots})$

$\text{design}_{\text{xml}} \leftarrow \text{Compiler}(\text{design}_{\text{new}})$

if $\text{design}_{\text{xml}}$ is corrupted **then continue** // skip to next j

$\text{fitness}_{\text{new}} = \text{Evaluate}(\text{design}_{\text{xml}})$

$\gamma_{\text{min}} \leftarrow \arg \min_q \text{fitness}_q$ in Q

if $\text{fitness}_{\text{new}} > Q[\gamma_{\text{min}}].\text{fitness}$ **then**

$Q[\gamma_{\text{min}}] \leftarrow (\text{fitness}_{\text{new}}, \text{design}_{\text{new}}, \text{reason}_{\text{new}})$

end if

end for

$\text{fewshots} \leftarrow \emptyset$

for all $(\text{fitness}, \text{design}, \text{reason}) \in Q$ **do**

$\text{fewshots.append}(\text{str}(\text{design}, \text{reason}))$

end for

end for

end function

1) *System prompt*: Specifies the robot grammar rules [8], which constrain outputs to physically realizable designs. In effect, the system prompt describes the “building blocks” and rules for assembling them. The LLM is instructed to apply a reasoning trace detailing a step-by-step application of these rules during the assembly process. This not only clarifies the LLM’s design rationale but also helps the model avoid violating the grammar constraints.

2) *User prompt*: Provides a fixed high-level instruction: “Your task is to design a single robot. Use the examples provided to guide your reasoning and explain your design step by step.” This chain-of-thought style [12] encourages

sequential reasoning and improves the interpretability of the LLM’s outputs.

3) *Few-shot examples*: Include top-performing designs (with reasoning traces) from previous iterations. These serve as templates, allowing the LLM to reuse effective traits while maintaining diversity. We omit explicit fitness scores, as providing them caused the model to overfit and repeatedly propose similar designs, thereby reducing exploration.

B. Best-Shot Prompting

Our prompting strategy, inspired by in-context learning [13], augments the prompt with high-quality few-shot examples. We call this approach *best-shot prompting*, motivated by three factors:

Reducing invalid outputs. In zero-shot trials, the LLM often produced designs that violated grammar rules, yielding XML files that could not be compiled. Providing valid designs with reasoning steps grounds the model in the grammar and significantly reduces such errors [14]. In our experiments, most zero-shot outputs failed to compile, whereas with best-shot prompting and chain-of-thought reasoning, the vast majority of designs were grammar-valid.

Encouraging diversity. Without diverse exemplars, the LLM tended to repeat small variations of one concept. We seed prompts with randomly generated designs (Figure 3) to span different structures, and update them with top performers each generation. This maintains variety in limb count, joint types, and symmetries, preventing premature convergence.

Iterative improvement. Updating few-shot examples with the current best designs introduces an implicit elite-selection effect. The prompt gradually shifts toward stronger exemplars, guiding the LLM to propose increasingly effective morphologies. We exclude numeric fitness values because including them led to overfitting and reduced exploration. Instead, the model infers successful traits from qualitative patterns in the examples, often producing unconventional designs that are then filtered by evolution (Section II-F).




Fig. 3: Examples of randomly generated initial few-shot robot designs used to seed Algorithm 1. These serve as diverse seed morphologies to prime the evolutionary process.

C. Robot Design

Our design space is defined by a context-free grammar adapted from Robogrammar [8], which specifies how robots can be constructed from a small set of modules. The grammar enforces bilateral symmetry, meaning any limb added to one side of the body must be mirrored on the opposite side. This constraint simplifies locomotion, but also reduces the

effective diversity of possible morphologies. Each expansion rule in the grammar maps directly to a physically realizable module.

As shown in Robogrammar, such a grammar can describe hundreds of thousands of designs while guaranteeing that every generated robot can be assembled from modular parts. In our framework, this ensures that the LLM’s proposals always correspond to feasible designs, though the space is limited by the chosen module library.

The LLM generates textual representations of a robot design in the grammar formalism, which avoids the need for explicit tree search as in Robogrammar. This significantly accelerates design generation, as the LLM proposes complete structures directly. A compiler converts the text into an XML description compatible with MuJoCo (via MJX). If the design violates the grammar, the compiler fails to produce a valid XML and the output is discarded. In practice, invalid outputs are uncommon, and most designs compile successfully once the few-shot examples are in place.

After obtaining a valid XML model of the robot, we disable self-collisions for the duration of policy training to accelerate simulation, a trade-off also adopted in prior work [15]. This simplification may allow physically unrealistic interpenetration. However, our tree-structured grammars produce mostly open morphologies where self-collisions are infrequent, and we apply this condition uniformly across all designs. In Sections III-B and III-C, we re-enable collisions for final evaluation. The top designs retained their effectiveness under full collision physics, confirming physical plausibility.

D. Simulation

A key challenge in our approach is that each new robot design must be evaluated by training a controller in simulation, which is computationally expensive. We use MuJoCo XLA (MJX), a JAX-based implementation of MuJoCo [11] capable of running on accelerated hardware, combined with Brax [16] for RL algorithms such as SAC [17] and PPO [18]. MJX uses GPUs to parallelize experience collection across thousands of environment instances. We further parallelize RL training across multiple GPUs, allowing several designs to be trained simultaneously. This combined speedup makes the evolutionary loop practical by significantly increasing both the number of iterations and designs evaluated per generation.

E. Control

For each valid design, we train a locomotion controller in simulation using the SAC algorithm, and the resulting policy is evaluated to compute a fitness score. We also experimented with PPO, but found SAC to be more consistent across a diverse range of robot morphologies. We did not tailor the RL parameters to each morphology or terrain, as each design was evaluated using the same set of hyperparameters and for the same number of environment steps.

F. Evolution

We treat the LLM as a stochastic design generator analogous to a mutation operator. In each generation, we sample the LLM multiple times to obtain a diverse set of candidate designs. A grammar verifier filters out any invalid outputs, ensuring that only structurally sound candidates enter the evaluation pipeline. This combination of LLM stochasticity and grammar-based verification balances broad exploration with physical plausibility.

After all designs in a generation have been evaluated, we rank them by fitness and retain the top- K candidates, favoring the “survival” of better designs. This introduces an elite-selection effect, and the top-performing designs form the few-shot examples for the LLM in the next generation. Over successive generations, this evolutionary loop of variation and selection improves the population’s overall fitness. Three mechanisms maintain exploration during exploitation. First, LLM sampling stochasticity ensures each call produces a distinct design. Second, excluding fitness scores from prompts prevents overfitting and preserves structural variety (Section III-D). Third, randomly generated initial examples establish a broad starting distribution.

III. EXPERIMENTS

A. Experimental Setup

We evaluate RoboMorph using 8 random seeds over 50 evolutionary generations, with a population size of 32 designs per evolution, $K=3$ few-shot examples, and GPT-4o at its default temperature of 1.0. Given the stochastic nature of the LLM’s output, our initial grid search over these hyperparameters proved ineffective, as each combination yielded varying results across seeds. We therefore chose reasonable values for these hyperparameters based on similar work in [19]. For completeness, we also conduct ablation studies on (i) LLM variants and (ii) evolution strategy, which are presented in Section III-D. In total, 1,600 robot designs are evaluated per seed. During RL policy training, we evaluate 8 robot designs simultaneously in a computing cluster equipped with 8 NVIDIA A100-40GB GPUs. The total runtime of RoboMorph for a single seed amounts to approximately 200 GPU hours. We provide further details on the computational time of our framework in Section III-E.

We assess the fitness of each design within a custom MJX environment and use the Brax implementation of the SAC algorithm. The system state is represented by a vector comprising the position and velocity of each joint, while the action space consists of continuous joint torques in the range $(-1, 1)$. The per-timestep reward is defined in Equation 1 as the instantaneous forward velocity of the robot’s head (the most forward body part) as it transitions from state s to s' under action a :

$$r_t(s, a, s') = \vec{v}_x(s, a, s'), \quad (1)$$

where \vec{v}_x denotes the robot head’s center-of-mass velocity in the x (forward) direction. Rollouts are terminated early if

the agent becomes unhealthy (based on z -height) or when the episode length is reached.

The overall fitness score of a design, given in Equation 2, is defined as the average reward achieved by its learned policy. We estimate this via Monte Carlo simulation across $N = 128$ rollouts:

$$F(\text{design}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{j=1}^{T_i} r_{ij}, \quad (2)$$

where T_i is the length of the i -th rollout and r_{ij} is the reward obtained at the j -th step of that rollout.

For policy learning, both the critic and policy networks are 2-layer MLPs with 256 units per layer. We employ the Adam optimizer [20] and use the ReLU activation function across all layers in each network. Each policy is trained for 10^6 environment steps with a batch size of 256, which we found provides a good balance between computational efficiency and training stability. All other RL-related hyperparameters are kept constant across all experiments for fairness and are listed in our repository.

We evaluate designs on four terrain environments that present distinct locomotion challenges. These are adapted from Robogrammar, with modifications where simulator limitations prevented exact replication. Together, they test robustness across smooth, uneven, low-friction, and constrained settings. Figure 5 provides visualizations of each terrain.

Ridged terrain. A series of periodic ridges that require climbing or crawling strategies for effective navigation.

Flat terrain. A featureless, high-friction plane suitable for a broad range of locomotion modes.

Frozen terrain. A low-friction surface that demands stability and efficient traction for forward movement.

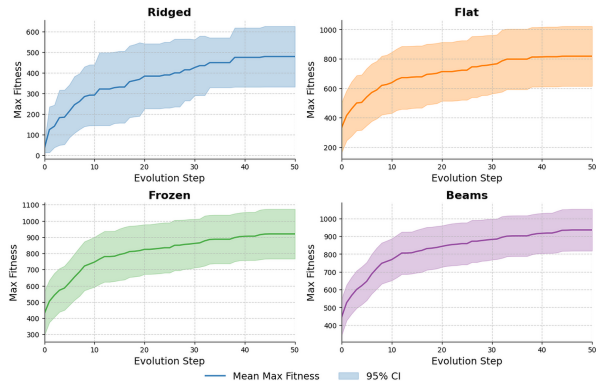
Beams terrain. A series of platforms with overhead beams that penalize tall or upright morphologies and reward compact, low-profile designs.

B. Main Experiment

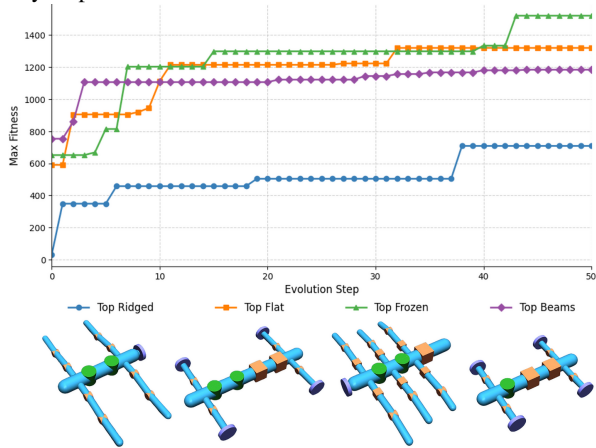
We first evaluate RoboMorph over the course of evolution. Figure 4a shows the average of the maximum fitness scores for the top design in the population at each evolutionary generation, across 8 seeds (shaded region denotes 95% confidence). A positive trend indicates that successive iterations of RoboMorph produce progressively better designs. The maximum fitness is monotonically non-decreasing by construction due to elite retention across generations.

Figure 4b depicts the evolution of the best-performing seed on each terrain, along with the highest overall fitness attained. This represents the exploitation phase of our evolutionary algorithm (see Section II-F), showing the best result RoboMorph discovered for each environment. The top-performing designs for each of the four terrains are visualized in Figure 5. Each terrain yields a distinct morphology tailored to the challenges of that environment.

Ridged terrain. A quadruped with elongated limbs and knee joints that can swing upward, allowing the robot to step over ridges and clear obstacles with high foot clearance.



(a) Average maximum fitness with 95% confidence intervals across 8 seeds for the best robot design in the population at each evolutionary step.



(b) Evolution of the best-performing seed for each terrain.

Fig. 4: Maximum fitness of the best robot design in the population at each evolutionary step. Fitness consistently improves across generations, confirming that RoboMorph progressively improves design quality.

Flat terrain. A low-profile body with shorter limbs ending in unactuated wheels spaced far apart. This wheeled quadruped exploits continuous rolling contact with the ground, minimizing energy loss and enabling faster locomotion than its legged counterpart. The low inertia of the short limbs, combined with free-spinning wheels, allows for smoother movements on flat surfaces.

Frozen terrain. A hexapod with rigid upper-leg joints that keep the legs splayed outward, creating a wide support base and low center of gravity. This configuration provides stability on slippery surfaces by resisting lateral sliding. Hexapods inherently offer greater static stability by always having at least three legs in contact with the ground, keeping the center of mass within the support polygon [21].

Beams terrain. Morphologically similar to the flat-terrain design, but with a lower profile and shorter limbs to maintain clearance under overhead beams. This similarity validates the framework, as identical ground conditions yield similar strategies, with adaptation only where the environment demands it (the overhead height constraint).




Fig. 5: Top robot designs from RoboMorph per terrain: ridged (top-left), flat (top-right), frozen (bottom-left), and beams (bottom-right). Each terrain yields a distinct morphology, including wheeled and hexapod designs not found by prior methods.

C. RoboMorph vs. Robogrammar

We compare RoboMorph against Robogrammar by reconstructing their best-performing designs for ridged, flat, and frozen terrains in our MJX environment (Figure 6). Robogrammar’s gapped terrain could not be replicated due to MJX collision geometry limitations, and our beams terrain has no Robogrammar baseline.

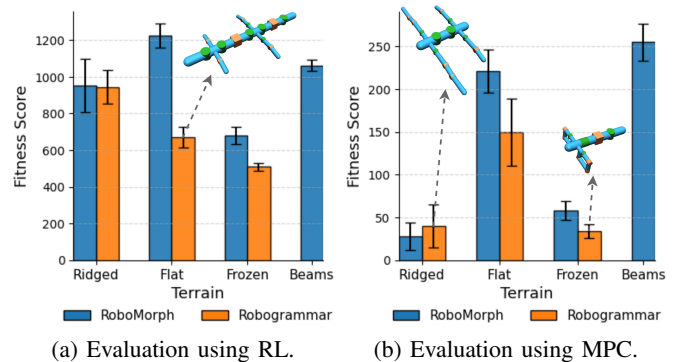


Fig. 6: Average fitness of top designs from RoboMorph and Robogrammar across 8 seeds per terrain. Under RL, RoboMorph matches or outperforms Robogrammar on all shared terrains. Under MPC, results are comparable on ridged terrain, while RoboMorph outperforms on flat and frozen.

For a fair comparison, we re-enabled full collisions for both methods’ designs (disabled during evolutionary search to reduce computation, as noted in Section II-C). We train an RL policy from scratch for each design and repeat training 8 times per design to account for variability¹. Results under RL control are shown in Figure 6a.

¹We take the highest-fitness design found by RoboMorph across all evolution runs and evaluate it with 8 independent RL trainings.

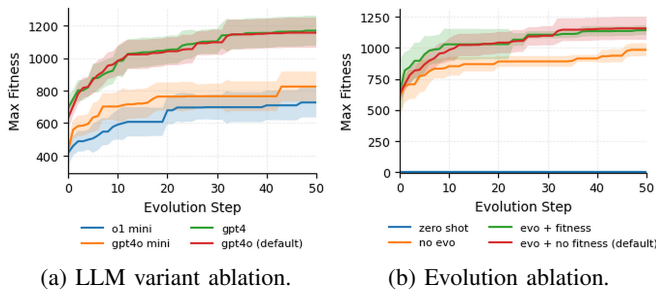


Fig. 7: Ablation on the flat terrain: average maximum fitness with 95% confidence intervals across 8 seeds. Larger LLMs and evolutionary prompting both yield higher fitness.

Since Robogrammar used a model-predictive control (MPC) planner, we also compare under MPC by replicating their MPPI-based controller (Figure 6b). This comparison is inherently asymmetric, as RoboMorph’s designs were evolved under RL-based fitness while Robogrammar’s were optimized with MPC. Evaluating under both controllers mitigates this bias, and consistent performance trends across both schemes support the validity of our conclusions.

Figure 6 compares performance of both methods on each terrain. On ridged terrain, both methods achieve similar fitness scores with overlapping error bars. This result is expected given that the two designs are morphologically similar. Both are quadrupeds featuring roll joints in the body and a comparable arrangement of leg joints. On flat terrain, RoboMorph significantly outperforms Robogrammar’s leg-only solution, as the discovered wheeled design yields faster and more efficient locomotion. Robogrammar’s graph search does not readily explore wheeled configurations because incremental structural mutations are unlikely to produce the specific joint and wheel combination that enables rolling contact.

On frozen terrain, RoboMorph also produces a higher-performing design than Robogrammar’s bipedal structure, which relies on compact arms for ground contact while the rear body slides freely. Our approach instead discovered a hexapod with a wide base and continuous contact points, which is more efficient at generating forward thrust on a slippery surface. Overall, the results highlight that RoboMorph discovers novel, effective strategies that Robogrammar’s search did not find.

D. Ablation Studies

We conduct targeted ablation studies to isolate the impact of two main design choices in our pipeline: the LLM model variant and evolutionary prompting strategy. The results are summarized in Figure 7.

1) *LLM variant*: We compared several LLM back-ends for the design generation module, including GPT-4o (default), GPT-4o mini, GPT-4, and o1-mini. As seen in Figure 7a, both GPT-4o and GPT-4 produce significantly higher-fitness designs than the smaller variants. GPT-4 yielded only marginal improvements over GPT-4o but at substantially

higher cost (\$30/\$60 per million input/output tokens vs. \$2.50/\$10 for GPT-4o).

2) *Evolution strategy*: We compared four evolutionary strategies to understand the contribution of each component. We tested: (i) zero-shot baseline with no optimization, (ii) prompting without any evolution, (iii) evolution with fitness scores, and (iv) evolution without fitness scores. Results in Figure 7b show that incorporating evolution in the pipeline outperforms non-evolution baselines, confirming that iterative refinement is key. Explicit fitness scores provide only minor benefit and can reduce diversity by overfitting to top designs. Excluding scores maintains exploration while still improving performance, supporting our choice in Section II-B.

E. Computational Time

Overall compute. RoboMorph requires approximately 200 GPU hours per seed on A100 GPUs, compared to Robogrammar’s reported 31 CPU hours [8]. In total compute, RoboMorph is substantially more expensive, and these numbers correspond to different hardware classes (GPUs vs. CPUs), making direct comparison difficult. However, our compute is dominated by RL training, which is highly parallelizable. On 8 A100 GPUs, a full run completes in about 25 wall-clock hours. While Robogrammar invests most of its compute in sequential graph search, RoboMorph shifts effort to parallelizable RL evaluation. This parallelism allows RoboMorph to converge within about a day, making it practical for iterative development despite higher total resource usage.

Design generation. Each GPT-4o API call takes approximately 9.5s (1s time-to-first-token plus ~ 1500 output tokens at 175 tokens/s) [22]. For one seed with 1600 evaluations, cumulative design generation time is ~ 4.2 hours, confirming that RL evaluation dominates total compute. By comparison, Robogrammar allocates 11 of its 31 total hours to design generation.

F. Discussion

On the LLM’s contribution. Our ablation studies (Section III-D) demonstrate that iterative evolution with best-shot prompting yields consistent improvements over zero-shot and non-evolutionary baselines. However, we acknowledge that our experiments do not include a traditional evolutionary algorithm (e.g., random mutation and crossover on the same grammar) as a baseline. Such a comparison would help isolate whether the LLM’s structured generation provides an advantage beyond acting as a stochastic design sampler. Unlike random mutation, the LLM produces complete, grammar-valid designs with reasoning traces, which helps explain its effectiveness in navigating the combinatorial design space. Crucially, the LLM can make non-local jumps in design space, proposing an entirely different body plan (e.g., shifting from a quadruped to a hexapod) in a single generation, whereas traditional mutation operators are limited to incremental structural edits. Recent work [14] suggests that LLMs can outperform random operators by leveraging

learned priors. A direct comparison with a non-LLM EA remains an important direction for future work.

On morphological diversity. A notable outcome of RoboMorph is the structural diversity of the discovered designs. Across terrains, the framework converges on qualitatively distinct morphologies, including wheeled quadrupeds on flat ground, high-clearance legged walkers on ridges, wide-stance hexapods on ice, and compact crawlers under beams. This diversity arises without any explicit diversity objective. We attribute this diversity to the LLM’s broad prior over structural compositions, which allows it to propose fundamentally different body plans in a single generation, and to the terrain-specific fitness signal, which selectively reinforces the morphologies best suited to each environment. By contrast, grammar-based graph search typically converges on a narrower set of local optima because its mutation operators make only incremental structural changes. These diversity claims are qualitative rather than formally measured.

Failure modes. We observe three common failure modes: (1) grammar violations in early or zero-shot generations, which best-shot prompting largely mitigates, (2) kinematically degenerate designs (e.g., collapsed or overlapping joints) that compile but produce no useful motion, which are pruned by fitness-based selection, and (3) morphological convergence when prompt diversity is low, which our diverse initialization and fitness-score exclusion help counter.

IV. RELATED WORK

LLMs for evolutionary search. LLMs have recently shown promise for creative engineering tasks [14], but their application to automated robot design remains limited [1]. Early efforts introduced a human-in-the-loop framework that leverages ChatGPT to steer both conceptual and technical stages of design [23], while [24] demonstrated how GPT-4 can synthesize manufacturable mechanical assemblies through discrete composition. Subsequent studies coupled LLMs with principled search or prompting strategies. [25] showed that combining evolutionary algorithms with LLMs produces powerful prompt optimizers, and [26] introduced grammar prompting to ensure domain-specific language outputs remain syntactically valid. In parallel, LASER [27] adopts a bi-level framework, using an LLM as a mutation operator to diversify robot morphologies within an evolutionary loop. Similarly, Text2Robot [19] translates natural-language descriptions into quadruped designs given a fixed limb count template. In contrast to these works, our method operates directly in a 3D design space and places no a priori limits on the number of limbs, enabling exploration of a far larger and more complex morphology search space.

Robot design automation. Beyond LLMs, robot design is shifting from exhaustive combinatorial search to data-driven methods that learn task-specific morphologies. Early work showed that deep reinforcement learning can act as a heuristic to guide the search for optimal modular designs [28]. Later, morphology generation was framed as a conditional generative problem using GANs to map task descriptions

to diverse robot assemblies [29]. GLSO [9] extended this direction by embedding the discrete design space into a continuous latent space via a grammar-guided variational autoencoder. Beyond modular assemblies, diffusion models paired with learned dynamics networks have been used to generate custom end-effectors [30], and policy-gradient methods have been applied to voxel robots [31]. In parallel, graph search has been augmented with neural heuristics [32], and [10] proposed a multi-objective graph search that jointly optimizes morphology and control. Together, these advances have broadened the toolkit for automated robot design. Our work builds on this trend by incorporating an LLM as a generative model and an evolutionary pipeline to efficiently traverse the design space guided by performance feedback.

Legged locomotion. Classical control methods such as MPC [33], [34], as used in Robogrammar, require accurate models and extensive tuning. We instead adopt model-free RL [35], [36], training a separate SAC [17] policy per design. Recent morphology-agnostic controllers [37], [38] could replace per-morphology training in the future, and our modular framework is compatible with such advances.

V. LIMITATIONS AND FUTURE WORK

Expanding the design space. One future direction is to broaden the range of generated robot designs. This could be achieved by relaxing grammar constraints to permit a wider variety of materials, component types, and scalable link dimensions. Another promising avenue is to map grammar-defined components directly to hardware modules, enabling real-world assembly from a standardized kit of parts.

Diverse mix of environments. A natural next step is to extend the framework to robots that can traverse multiple terrain types, for example by initializing with few-shot examples tailored to different environments or incorporating textual terrain descriptions directly into the prompt. Achieving robustness across terrains would move closer to creating versatile, all-terrain robots and serve as a testbed for evaluating how well LLMs generalize design principles beyond single-task optimization.

Joint generation of a robot and its policy. A promising direction is jointly generating robot morphologies and corresponding control strategies with LLMs. Recent studies, such as [4], have shown that LLMs can generate reward functions without task-specific prompting or predefined templates. Building on this, LLMs could also propose high-level control heuristics or policy structures in parallel with morphology, which may help reduce the reliance on costly per-design RL training.

Baselines and robustness. Comparing against a traditional EA with random mutation and crossover on the same grammar and RL pipeline would isolate the LLM’s contribution more precisely. Studying prompt sensitivity across variations in phrasing, few-shot formatting, and chain-of-thought instructions would clarify robustness to prompt design choices. Evaluating open-source LLM alternatives and incorporating quantitative diversity metrics would further strengthen the evaluation.

VI. CONCLUSION

We presented RoboMorph, a framework that integrates LLMs, evolutionary algorithms, RL-based control, and robot grammars to automate the design of modular robots. Robot morphologies evolved by RoboMorph improve consistently over successive generations, and the final designs match or exceed the performance of Robogrammar’s graph-search method across multiple terrains. The discovered morphologies, spanning wheeled quadrupeds, hexapods, and low-profile crawlers, demonstrate the LLM’s capacity to explore beyond human-intuitive solutions while remaining physically feasible. Notably, these designs emerged without hand-crafted mutation operators or domain-specific heuristics, relying solely on the LLM’s prior knowledge guided by evolutionary selection pressure.

More broadly, our results show that LLMs can serve as effective generative operators within evolutionary design loops, complementing rather than replacing traditional search. The key advantage lies in the LLM’s ability to propose structurally complete and diverse candidates drawing on broad prior knowledge, reducing the need for manually designed variation operators. This paradigm extends naturally to other grammar-constrained design domains such as mechanisms, furniture, and architectural structures. In future work, we aim to extend RoboMorph beyond simulation by integrating real-world manufacturing constraints and rapid fabrication techniques to produce deployable, task-specific robots.

ACKNOWLEDGMENTS

This work was partially supported by the National Science Centre, Poland, under PRELUDIUM Grant No. 2024/53/N/ST6/03370 and Sonata Bis Grant No. 2024/54/E/ST6/00388.

REFERENCES

- [1] F. Zeng *et al.*, “Large language models for robotics: A survey,” *arXiv preprint arXiv:2311.07226*, 2023.
- [2] J. Liang *et al.*, “Code as policies: Language model programs for embodied control,” 2023.
- [3] A. Brohan *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” 2023.
- [4] Y. J. Ma *et al.*, “Eureka: Human-level reward design via coding large language models,” 2024.
- [5] —, “Dreureka: Language model guided sim-to-real transfer,” in *Robotics: Science and Systems (RSS)*, 2024.
- [6] T. Bäck *et al.*, “An overview of evolutionary algorithms for parameter optimization,” *Evolutionary computation*, vol. 1, no. 1, pp. 1–23, 1993.
- [7] D. Thierens, “Scalability problems of simple genetic algorithms,” *Evolutionary computation*, vol. 7, no. 4, pp. 331–352, 1999.
- [8] A. Zhao *et al.*, “Robogrammar: graph grammar for terrain-optimized robot design,” *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–16, 2020.
- [9] J. Hu *et al.*, “Glso: Grammar-guided latent space optimization for sample-efficient robot design automation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1321–1331.
- [10] J. Xu *et al.*, “Multi-objective graph heuristic search for terrestrial robot design,” in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 9863–9869.
- [11] E. Todorov *et al.*, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
- [12] J. Wei *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [13] T. Brown *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [14] B. Romera-Paredes *et al.*, “Mathematical discoveries from program search with large language models,” *Nature*, vol. 625, no. 7995, pp. 468–475, 2024.
- [15] A. Gupta *et al.*, “Embodied intelligence via learning and evolution,” *Nature communications*, vol. 12, no. 1, p. 5721, 2021.
- [16] C. D. Freeman *et al.*, “Brax - a differentiable physics engine for large scale rigid body simulation,” 2021. [Online]. Available: <http://github.com/google/brax>
- [17] T. Haarnoja *et al.*, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International Conference on Machine Learning (ICML)*, 2018.
- [18] J. Schulman *et al.*, “Proximal policy optimization algorithms,” 2017. [Online]. Available: <https://arxiv.org/abs/1707.06347>
- [19] R. P. Ringel *et al.*, “Text2robot: Evolutionary robot design from text descriptions,” *arXiv preprint arXiv:2406.19963*, 2024.
- [20] D. P. Kingma *et al.*, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [21] R. B. McGhee *et al.*, “Adaptive locomotion of a multilegged robot over rough terrain,” *IEEE transactions on systems, man, and cybernetics*, vol. 9, no. 4, pp. 176–182, 2007.
- [22] Artificial Analysis, “Gpt-4o providers,” 2024, accessed: 2025-04-25. [Online]. Available: <https://artificialanalysis.ai/models/gpt-4o/providers>
- [23] F. Stella *et al.*, “How can llms transform the robotic design process?” *Nature machine intelligence*, vol. 5, no. 6, pp. 561–564, 2023.
- [24] L. Makatura *et al.*, “How can large language models help humans in design and manufacturing?” *arXiv preprint arXiv:2307.14377*, 2023.
- [25] Q. Guo *et al.*, “Connecting large language models with evolutionary algorithms yields powerful prompt optimizers,” 2024. [Online]. Available: <https://arxiv.org/abs/2309.08532>
- [26] B. Wang *et al.*, “Grammar prompting for domain-specific language generation with large language models,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 65 030–65 055, 2023.
- [27] J. Song *et al.*, “Laser: Towards diversified and generalizable robot design with large language models,” in *The Thirteenth International Conference on Learning Representations*, 2025.
- [28] J. Whitman *et al.*, “Modular robot design synthesis with deep reinforcement learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 06, 2020, pp. 10 418–10 425.
- [29] J. Hu *et al.*, “Modular robot design optimization with generative adversarial networks,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4282–4288.
- [30] X. Xu *et al.*, “Dynamics-guided diffusion model for robot manipulator design,” *arXiv preprint arXiv:2402.15038*, 2024.
- [31] M. Li *et al.*, “Reinforcement learning for freeform robot design,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8799–8806.
- [32] T. Wang *et al.*, “Neural graph evolution: Towards efficient automatic robot design,” *arXiv preprint arXiv:1906.05370*, 2019.
- [33] J. Di Carlo *et al.*, “Dynamic locomotion in the mit cheetah 3 through convex model-predictive control,” in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2018, pp. 1–9.
- [34] D. Kim *et al.*, “Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control,” *arXiv preprint arXiv:1909.06586*, 2019.
- [35] J. Hwangbo *et al.*, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [36] T. Haarnoja *et al.*, “Learning to walk via deep reinforcement learning,” *arXiv preprint arXiv:1812.11103*, 2018.
- [37] M. Shafiee *et al.*, “Manyquadrupeds: Learning a single locomotion policy for diverse quadruped robots,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3471–3477.
- [38] N. Bohlinger *et al.*, “One policy to run them all: an end-to-end learning approach to multi-embodiment locomotion,” in *Conference on Robot Learning*. PMLR, 2025, pp. 3356–3378.