
Evolution Gym: A Large-Scale Benchmark for Evolving Soft Robots

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Both the design and control of a robot play equally important roles in its task perfor-
2 mance. However, while optimal control is well studied in the machine learning and
3 robotics community, less attention is placed on finding the optimal robot design.
4 This is mainly because co-optimizing the robot design and control is character-
5 ized as a challenging problem, and more importantly, a comprehensive evaluation
6 benchmark for robot co-optimization does not exist. In this paper, we propose
7 Evolution Gym, the first large-scale benchmark for evolving the design of soft
8 robots. In our benchmark, each robot is composed of different types of voxels (e.g.,
9 soft, rigid, actuators), resulting in a modular and expressive robot design space.
10 Our benchmark environments span a wide range of tasks, including locomotion on
11 various types of terrains and manipulation. Furthermore, we develop several robot
12 evolution algorithms by combining state-of-the-art design optimization methods
13 and deep reinforcement learning techniques. By evaluating the algorithms on our
14 benchmark platform, the results demonstrate that the performance of robots contin-
15 uously increases as the evolution progresses. Additionally, even though the robots
16 are evolved autonomously from scratch without prior knowledge, they often grow
17 to resemble existing natural creatures while outperforming hand-designed robots.
18 Nevertheless, all tested algorithms fail to find robots that succeed in our hardest envi-
19 ronments. This suggests that more advanced algorithms are required to explore the
20 high dimensional design space and evolve increasingly intelligent robots - an open
21 problem for future research. Our website with code, environments, and videos is
22 available at <https://sites.google.com/view/evolution-gym-benchmark>

23 1 Introduction

24 One of the main goals of artificial intelligence is to develop effective approaches for the creation
25 of embodied intelligent systems. Inspired from real organisms, where body structure and brain are
26 two key factors for completing any task in a real environment, a successful intelligent robot typically
27 requires concurrently optimizing its structure design and control mechanism. Such a co-design
28 problem has been a long-standing key challenge in the robotics and machine learning communities.
29 Surprisingly, despite its importance, most previous research works still either only develop complex
30 control algorithms for existing robot structures [1, 2, 15, 27], or conduct co-optimization over robot
31 morphology and control for only a few simple tasks (e.g., running, jumping) [7, 28, 29]. The primary
32 reasons behind the under-exploration of co-design algorithms in sophisticated problems are: (1) the
33 underlining complex bilevel optimization scheme of a co-design algorithm, where the inner control
34 optimization loop leads to a long iteration cycle of the whole optimization process; (2) the lack of a
35 well-established benchmark platform providing the researchers with a suite to evaluate and compare
36 different algorithms.

37 Digital benchmark environments have proven to be successful at promoting the development of ad-
38 vanced learning techniques via providing a comprehensive evaluation suite to make fair comparisons
39 among different algorithms [5, 11, 33]. However, to our best knowledge, all existing benchmark plat-
40 forms constrain their domains within control optimization problems, and the space of co-optimization
41 environment suites is still rarely explored.

42 To fill this gap, in this work we propose Evolution Gym, a large-scale benchmark for evolving both the
43 shape structure and controller of soft robots. The body of each robot in Evolution Gym is composed
44 of various types of primitive building blocks (*e.g.*, soft voxels, rigid voxels, actuator voxels), and the
45 control of the robot includes the action signals applied on the actuator voxels. We choose to use this
46 multi-material voxel-based structure as the representation of robot body since it provides a general
47 and universal representation for various categories of robot designs, and at the same time results in a
48 modular and expressive structure design space. We adopt a mass-spring dynamics system [24] with
49 penalty-based frictional contact as the underlining physics engine. Such a light-weight simulator
50 allows the co-design algorithms to significantly reduce the simulation cost and thus accelerate the
51 develop-evaluate iteration cycle [3, 13, 21]. The back-end simulator is fully developed in C++ to
52 provide further computing efficiency. Another feature of Evolution Gym is its large variety of tasks
53 categorized by varying difficulty levels, which offer an extensive evaluation benchmark for comparing
54 approaches. The benchmark is currently comprised of more than 30 tasks, spanning locomotion on
55 various types of terrains and manipulation. Moreover, Evolution Gym is easy to use. In order to have
56 user-friendly interfaces, we build a Python wrapper outside the C++ simulator and carefully design
57 our APIs off of the well-received APIs of OpenAI Gym with minimum modifications. Evolution
58 Gym will be released fully open-source under the MIT license.

59 In addition, we develop several baseline algorithms by integrating the state-of-the-art design opti-
60 mization approaches and reinforcement learning techniques. In our baseline algorithms, specifically,
61 the design optimization methods are served in the outer loop to evolve the physical structures of
62 the robots and the reinforcement learning algorithms are applied in the inner loop to compute the
63 optimal control of a given proposed structure design. We conduct extensive experiments to evaluate
64 all baseline algorithms on Evolution Gym. The experiment results demonstrate that intelligent robot
65 designs can be evolved fully autonomously while outperforming hand-designed robots in easier tasks,
66 which reaffirms the necessity of jointly optimizing for both robot structure and control. However,
67 none of the baseline algorithms are capable enough to successfully find robots that complete the task
68 in our hardest environments. Such insufficiency of the existing algorithms suggests the demand for
69 more advanced robot co-design techniques, and we believe our proposed Evolution Gym provides a
70 comprehensive evaluation testbed for robot co-design and unlocks future research in this direction.

71 In summary, our work has the following key contributions: (i) We propose Evolution Gym, the
72 first large-scale benchmark for soft robot co-design algorithms. (ii) We develop several co-design
73 algorithms by combining state-of-the-art design optimization methods and deep reinforcement
74 learning techniques for control optimization. (iii) The developed algorithms are evaluated and
75 analyzed on our proposed benchmark suite, and the results validate the efficacy of robot co-design
76 while pointing out the failure and limitations of existing algorithms.

77 2 Related work

78 **Robot co-design** Co-designing the structure (*i.e.*, body) and control (*i.e.*, brain) of robots is a long-
79 standing key challenge in the robotics community. As the earliest work in this space, Sims [28]
80 represents the structure of a rigid robot as a directed graph and proposes an evolutionary algorithm
81 defined on graphs to optimize the robot design. Subsequently, the co-design of rigid robots is
82 formulated as a graph search problem where more efficient search algorithms are applied [35, 37] to
83 achieve increasingly interesting results. However, with the restriction of having rigid components
84 only, these algorithms are unable to produce optimal or even feasible designs for many challenging
85 tasks where a compliant joint or robot component is required to achieve the goal.

86 On the contrary, soft components offer much more flexibility to represent arbitrary shapes, making the
87 design of more complex, agile, and high-performing robots possible. Inspired by this, some work has
88 been conducted to co-design robots composed of soft cells. Cheney et al. [7, 8]; Van Diepen and Shea
89 [34]; Corucci et al. [10] propose evolutionary algorithms to co-optimize the structure and control of
90 voxel-based robots. However those algorithms typically parameterize the control as an open-loop

91 periodic sequence of actuation, which prevents robots from learning complex non-periodic tasks
 92 such as walking on uneven or varying terrains. Speilberg et al. [29] and Medvet et. al. [21] jointly
 93 optimize the spatial-varying material parameters and the neural network policy for soft robots but
 94 leaving the shape of the robot fixed. Our proposed benchmark shares a similar expressive structure
 95 design space as Cheney et al. [7], but allows the control to be parameterized by a sophisticated neural
 96 network feedback policy. To handle such sophisticated joint optimization of the robot structure and
 97 high-dimensional neural network control policy, we develop several baseline co-design algorithms by
 98 combining the state-of-the-art design optimization strategies and reinforcement learning techniques
 99 for control optimization.

100 **Benchmark environments for robotics learning** Present research in robotics learning is largely
 101 facilitated by emerging benchmark environments. For instance, OpenAI Gym [5], DeepMind
 102 Control Suite [33], rllab [11], and Gibson [36] have been developed to benchmark RL algorithms for
 103 controlling rigid robots. At the same time, PlasticineLab [14] is specifically designed for soft robot
 104 learning. However, the existing benchmark environments are all constructed for learning the control
 105 only. To enable the possibility of evolving the structure of a robot, the existing co-design work has to
 106 either implement their own testing environment [29, 7, 8, 10, 34], or make substantial changes on
 107 the underlying code of the existing control-only environments [26]. The independent development
 108 of testing beds requires non-trivial workload, and as a result existing co-design works mainly focus
 109 on evaluating the robot on a few simple tasks such as walking on a flat terrain [7, 6, 8, 34, 29, 21],
 110 or swimming along a single direction [9, 35]. An unintended consequence of such independency is
 111 an indirect comparison among different algorithms. Evolution Gym fills this gap by presenting a
 112 large variety of tasks with different difficulty levels that span from locomotion to manipulation. The
 113 proposed benchmark suite can be effectively used to test the generalizability of the algorithms on
 114 different tasks, potentially accelerating research in robot co-design.

115 3 Evolution Gym

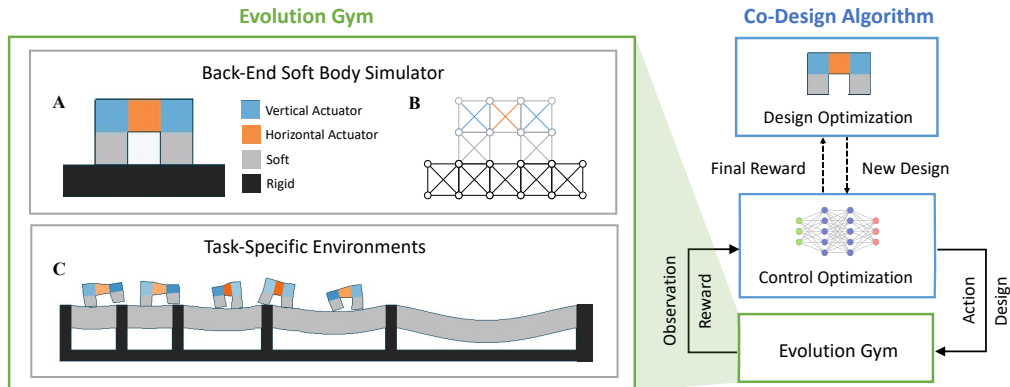


Figure 1: **Overview of Evolution Gym and its integration with the co-design algorithms.** Evolution Gym is comprised of a back-end soft body simulator (A, B) and task-specific environments (C). A user-customized co-design algorithm can be plugged in to optimize for both robot structure and control through interacting with Evolution Gym on a certain task.

116 3.1 Overview

117 In this section, we present Evolution Gym, a large-scale benchmark for co-design of voxel-based
 118 soft robots. Evolution Gym is featured by its versatile and expressive multi-material voxel-based
 119 structure design space, flexibility of the controller parameterization, wide spectrum of tasks of various
 120 difficulty levels, fast back-end soft-body simulation support, and user-friendly Python interfaces.

121 As shown in the overview in Figure 1, Evolution Gym is comprised of a task-specific environment
 122 and a back-end soft-body simulator. The gym suite provides seamless interfaces with a user-defined
 123 co-design algorithm. The co-design algorithm typically consists of a design optimizer and a control
 124 optimizer. The design optimizer can propose a new robot structure to the control optimizer, then the
 125 control optimizer will compute the optimal controller for the given structure through interactions with
 126 Evolution Gym and finally return the maximum reward that this robot structure can achieve. In this

127 way, Evolution Gym provides an easy-to-use platform for co-design algorithms to evolve both robot
128 structure and control to optimize for robots’ task performances. Evolution Gym is designed to be the
129 first comprehensive testbed for benchmarking and comparing different co-design algorithms with the
130 hope to facilitate the development of more novel and powerful algorithms in the co-design field.

131 3.2 Multi-material voxel-based representation

132 Evolution Gym employs a unified multi-material voxel-based representation for all the components
133 in the environment (*e.g.*, robot, terrain, object) as shown in Figure 1A. Specifically, each robot in
134 our gym is composed of rigid voxels, soft voxels, horizontal/vertical actuator voxels, and empty
135 voxels. Such a multi-material structure of robots provides a general and universal representation for
136 various categories of robot designs and results in a modular structure design space. Furthermore,
137 different combinations of the various types of voxels result in a combinatorial robot design space for
138 the co-design algorithms. For terrain and object, we use the same voxel-based structure but only with
139 passive voxel types (*i.e.*, soft/rigid voxels).

140 3.3 Task representation

141 Each task in Evolution Gym contains a robot structure proposed by the co-design algorithm, environ-
142 ment specifications (*e.g.*, terrain, object), and a task-related goal (*e.g.*, locomotion or manipulation).
143 The tasks interface with the co-design algorithm through a few key elements including *robot structure*
144 *specification*, *observation*, *action*, and *reward*. We introduce each element in details in below.

145 **Robot structure specification** As described in Section 3.2, we construct each robot from primitive
146 building blocks arranged on a grid layout. In code, each robot is specified as a material matrix of
147 voxels \mathcal{M} and a connection link list \mathcal{C} . The value of entry $m \in \mathcal{M}$ is a label corresponding to a voxel
148 type from the set {Empty, Rigid, Soft, Horizontal Actuator, Vertical Actuator}. The connection link
149 list \mathcal{C} stores a list of connection pairs of adjacent voxels. The co-design algorithm can update the
150 robot structure in the environment through *initialization* function with \mathcal{M} and \mathcal{C} as arguments.

151 **Observation** The observation is composed in each step to inform the controller of state information
152 of the robot, terrain information of the environment, and goal-relevant information. More specifically,
153 let N be the total number of voxel corner points of the robot. Then the state information of the robot
154 in our tasks is a $(2N + 3)$ -D vector including the relative position of each voxel corner with respect to
155 the center of mass of the robot ($2N$ -D), and the velocity and orientation of center of mass (3-D). To
156 handle complex tasks, specifically those with varying terrain types, an additional observation vector
157 including terrain information is provided. We compile terrain information within a local window of
158 size $2W$ around the robot into a length- $2W$ vector observation that describes the terrain’s elevation.
159 Furthermore, goal-related information is offered to inform the controller of the execution status of the
160 current task. This goal-related observation is task-specific and is defined on each task separately. For
161 instance, in manipulation tasks where the robot interacts with some object O , we provide orientation
162 and velocity as well as the position of O ’s center of mass relative to the robot.

163 **Action** At each time step, an action vector from the robot’s controller is provided to step Evo-
164 lution Gym’s simulator. In Evolution Gym, each value in the action vector is associated with an
165 actuator voxel (either horizontal or vertical) of the robot, and instructs a deformation target of that
166 voxel. Specifically, the action value u is within the range $[0.6, 1.6]$, and corresponds to a gradual
167 expansion/contraction of that actuator to u times its rest length.

168 **Reward** Each task is equipped with a reward function measuring the performance of the current
169 robot and the control action. The value of reward is defined step-wise and is fed back to the agent
170 through *step* function. The reward function is highly task-specific and should be defined to precisely
171 characterize the robot’s completeness of the task. Please refer to Section 3.5 and Appendix for
172 detailed descriptions of the reward functions on each task.

173 3.4 Simulation engine

174 We model the dynamics of the underlying simulator as a 2D mass-spring system [24]. This simple,
175 flexible formulation allows us to efficiently model soft robots with a wide range of capabilities in a
176 wide range of environments. The simulation engine is written entirely in C++. We create Python
177 bindings of our simulator so it seamlessly interfaces with standard learning frameworks.




Figure 2: A visual overview of selected 10 environments from Evolution Gym.

178 The simulation represents objects and their environment as a mass-spring system in a grid-like layout
 179 (Figure 1B). Objects and their environments are initialized as a set of non-overlapping, connected
 180 voxels. On initialization, each voxel is a cross-braced square, but may undergo deformation as the
 181 simulation progresses. Each edge acts as an ideal spring obeying Hooke’s law, with a spring constant
 182 defined by one of five possible material types. We employ symplectic RK-4 integration to step
 183 forward the simulation.

184 Collision detection is performed using a bounding-box tree structure [12]. Penalty-based contact
 185 forces and frictional forces are computed proportionally to the depth of penetration of the correspond-
 186 ing voxels in contact, and are applied on the voxel vertices in the normal and tangential directions of
 187 the contact respectively. Please refer to Appendix A for more details of simulation.

188 3.5 Benchmark environment suite

189 We have developed over 30 unique tasks with Evolution Gym and select 10 tasks here to illustrate
 190 the diversity and comprehensiveness of our benchmark task set. All tasks are organized into two
 191 categories - locomotion and manipulation - though some tasks are kind of a mix of both. We further
 192 classify the tasks into the different difficulty level (*i.e.*, easy, medium, hard) based on the performances
 193 of the baseline algorithms on them, which we will mention in Section 4. We briefly introduce the
 194 selected tasks in this section. For more detailed description and visualizations of the tasks, please
 195 refer to our website or Appendix B. It is also worth mentioning that our gym is designed to be
 196 extendable and the user can easily create new tasks for their needs.

197 3.5.1 Locomotion tasks

198 **Walker (Easy)** This is a common standard task typically considered by previous works where the
 199 robot needs to walk on a flat terrain as fast as possible.

200 **Bridge Walker (Easy)** In this task, the robot traverses a series of soft “rope” bridges separated by
 201 fixed pillars, and similarly as before it needs to maximize its forward speed.

202 **Stairs Walker (Medium)** The agent walks up a fixed staircase with of steps of varying length.

203 **Climber (Medium)** The robot must climb two tall fixed walls on each side. The robot is rewarded by
 204 its upward climbing speed.

205 **Traverser (Hard)** In this hard task, the robot needs to traverse a pit of rigid blocks to get to the other
 206 side without sinking into the pit.

207 3.5.2 Object manipulation tasks

208 **Carrier (Easy)** The robot needs to catch a small, soft rectangular object initially dropped from above
 209 and then carry it along the forward direction. The robot is rewarded by the distance both the robot
 210 and the object have traveled.

211 **Thrower (Medium)** The robot throws a soft rectangular box as far as possible without moving itself
 212 significantly from its original position.

- 213 **Slider (Hard)** In this task, a beam sits on top of a set of spaced-out floating platforms. The robot is
 214 rewarded for moving to the beam and sliding it in the forward direction.
- 215 **Catcher (Hard)** The agent needs to catch a spinning object randomly falling from a high location.
- 216 **Lifter (Hard)** The robot has to manipulate an object and lift it out of a hole.

217 4 Evolving soft robots

218 The robot evolution/co-design algorithm is formulated as a two-level optimization problem, which
 219 involves a design optimization method that evolves physical structures of the robots in the outer loop
 220 and a control optimization algorithm that computes the optimal control of a given robot structure in
 221 the inner loop, as illustrated in Algorithm I. We briefly introduce several instantiations of design
 222 optimization methods and control optimization methods in Section 4.1 and 4.2 that we use for
 223 evaluation on our benchmark, and more details can be found in Appendix C.

Algorithm 1 Algorithmic framework of robot evolution

Inputs: Task specification T , number of generations n , population size p .
Outputs: The best robot design D^* and controller C^* .
 $S \leftarrow \emptyset$ // Dataset of robot designs, controllers and reward
for $i \leftarrow 1$ **to** n **do**
 $D_1, \dots, D_p \leftarrow \text{OPTIMIZEDDESIGN}(S, p)$ // Propose a population of robot designs to evaluate
for $j \leftarrow 1$ **to** p **do**
 $C_j \leftarrow \text{OPTIMIZECONTROL}(T, D_j)$ // Optimize the controller of given robot design
 $r_j \leftarrow \text{EVALUATEREWARD}(T, D_j, C_j)$ // Evaluate the reward of given design and controller
 $S \leftarrow S \cup \{(D_j, C_j, r_j)\}$ // Update the evaluation result to the dataset
 Find the best design D^* and controller C^* in dataset S with the maximum reward r^* .

224 4.1 Design optimization

225 Design optimization aims at evolving the robot structures to maximize the reward under two physical
 226 constraints: the body has to be connected, and actuators must exist. In this section, we introduce
 227 three instantiations of the design optimization algorithm (OPTIMIZEDDESIGN in Algorithm I).

228 **Genetic algorithm (GA)** GAs [22] are widely used in optimizing black-box functions by relying on
 229 biologically inspired operators such as mutation, crossover and selection, as demonstrated in previous
 230 works on evolving rigid robots [28, 35]. We implement a simple GA using elitism selection and a
 231 simple mutation strategy to evolve the population of robot designs. Specifically, in each generation,
 232 our elitism selection works by keeping the top $x\%$ of the robots from the current population as
 233 survivors and discarding the rest, where x decreases gradually from 60 to 0 over generations. Next,
 234 we iteratively sample and mutate one of those survivors with 10% probability of changing each voxel
 235 of the robot to create more offsprings. Note that by mutating a voxel type from/to empty voxel, we
 236 are able to change the topology of the robot. The crossover operator is not implemented in our genetic
 237 algorithm.

238 **Bayesian optimization (BO)** BO [18, 23] is a commonly used global optimization method for
 239 black-box functions by learning and utilizing a surrogate model, which is usually employed to
 240 optimize expensive-to-evaluate functions, including evolving rigid robots in previous works [26, 19].
 241 Specifically, we choose a batch BO algorithm as described in Kandasamy et al. [16] and implemented
 242 in the GPyOpt package [4] that supports categorical input data. We use Gaussian processes as the
 243 surrogate model, batch Thompson sampling for extracting the acquisition function, and L-BFGS
 244 algorithm to optimize the acquisition function. To ensure a fair comparison with other population-
 245 based evolutionary baseline algorithms, the batch size of this algorithm is set equal to the population
 246 size of other algorithms.

247 **CPPN-NEAT** CPPN-NEAT is the predominant method for evolving soft robot design in previous
 248 literature [6, 7, 8]. In this method, the robot design is parameterized by a Compositional Pattern
 249 Producing Network (CPPN) [30]. The input to a CPPN is the spatial coordinate of a robot voxel
 250 and the output is the type of that voxel. Therefore, by querying the CPPN at all the spatial locations
 251 of a robot, we can obtain the type for each voxel to construct a robot. At the same time the

252 NeuroEvolution of Augmenting Topologies (NEAT) algorithm [31] is used to evolve the structure of
253 CPPNs by working as a genetic algorithm with specific mutation, crossover, and selection operators
254 defined on network structures. Our implementation of CPPN-NEAT is based on the PyTorch-NEAT
255 library [25] and the neat-python library [20].

256 4.2 Control optimization

257 In this section, we introduce the specific control optimization algorithm (OPTIMIZECONTROL in
258 Algorithm 1) that we use in the robot evolution algorithms. In previous works on evolving soft robots,
259 the controller is either encoded as a fixed periodic sequence of actuation [7] or is parameterized
260 as a CPPN that outputs the frequency and phase offset of the periodic actuation for each voxel [8].
261 However, the periodic pattern of the control prevents robots from learning complex non-periodic
262 tasks such as walking on uneven or varying terrains. Therefore, we use reinforcement learning (RL)
263 [32] to train the controller, making it possible for the soft robots to perform arbitrarily complex tasks
264 in our benchmark. Specifically, we apply a state-of-the-art RL algorithm named Proximal Policy
265 Optimization (PPO) [27] for control optimization of robots, with code implementation given by [17].

266 5 Experiments and results

267 In this section we present the evaluation results of baseline robot co-design algorithms on 10 selected
268 benchmark tasks described in Section 3.5. The complete evaluation results on all our benchmark
269 tasks can be found in Appendix E and also on our website.

270 We develop three baseline algorithms for robot evolution by combing the three design optimization
271 methods in Section 4.1 and PPO for control optimization in Section 4.2. Since the control optimization
272 method is the same for all baseline algorithms, we simply use GA, BO, CPPN-NEAT to denote
273 these three baseline algorithms with different design optimization methods. The evaluations of our
274 baseline algorithms are performed on machines with Intel Xeon CPU @ 2.80GHz * 80 processors on
275 Google Cloud Platform and GPU is not required. Evaluating one algorithm on a single task usually
276 takes several hours to twenty hours, depending on the number of evaluations, size of population, etc.
277 See Appendix D for more details on hyperparameters of all the experiments.

278 5.1 Comparisons among baseline algorithms

279 We plot the reward curves of the three baseline algorithms on 10 selected benchmark tasks in Figure 3.
280 There is no single optimal algorithm that performs the best on all tasks, but overall, GA outperforms
281 the other two baseline algorithms. This is surprising because our genetic algorithm is implemented
282 with simple and intuitive operators for mutation and selection without sophisticated mechanisms.
283 Therefore, we believe that with more carefully designed operators, GA has the potential to evolve
284 much more intelligent robots. CPPN-NEAT generally performs well on locomotion tasks, as tested
285 by previous works, but performs poorly on more complex manipulation tasks. It is possibly because
286 NEAT favors CPPNs with simpler structures, which encourages CPPNs to generate robots with more
287 regular patterns. However, to succeed in complex manipulation tasks, some agile substructures of
288 the robot must evolve, which might only exist in robots with irregular patterns. Finally, it is not
289 surprising that BO performs poorly on most of the tasks because the high-dimensional categorical
290 input parameter space and the noisy evaluation done by RL together pose a challenge to fitting an
291 accurate surrogate model in BO.

292 5.2 Evolution analysis

293 In Figure 4 we visualize the top four robots in three different generations on training the genetic
294 algorithm for the Carrier, Lifter, and Bridge Walker task. We also show the average reward these
295 designs achieve.

296 In the carrier task, the robot must catch an object that falls from above and then carries that object
297 as far as possible. Therefore, a design for this task requires two key components 1) A structure that
298 will catch and hold an object without letting it fall 2) A mechanism by which to move. These two
299 structures are barely visible in the top survivors of generation 1 (randomly initialized) but become
300 clearly more prominent by generation 10. By generation 30, the algorithm has nearly converged on

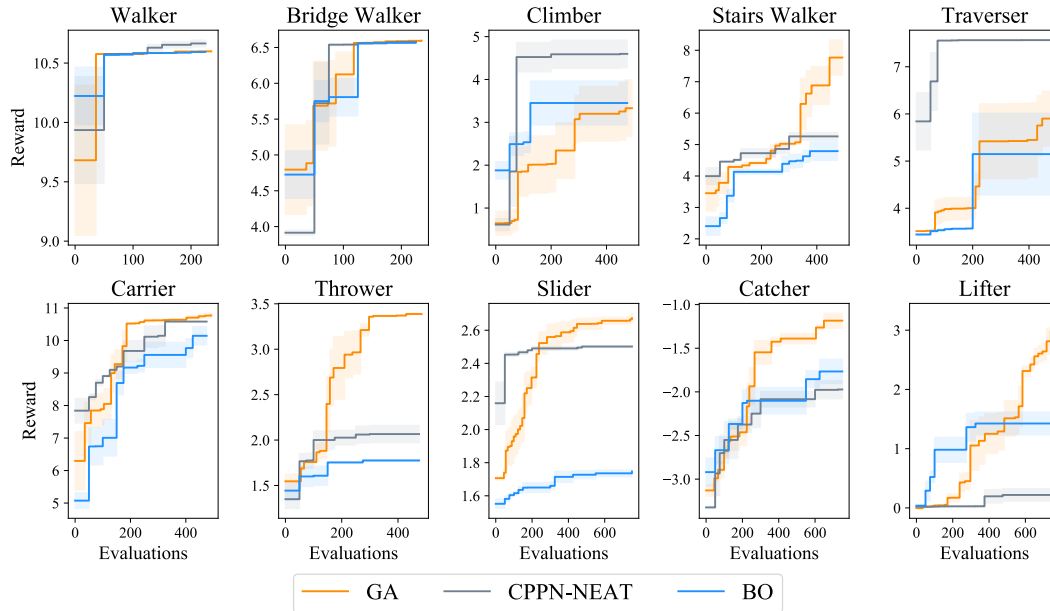


Figure 3: **Performance comparison among baseline algorithms.** We plot the best performance of robots that each algorithm has evolved w.r.t. the number of evaluations on each task. All the curves are averaged over 3 different random seeds, and the variance is shown as a shaded region.

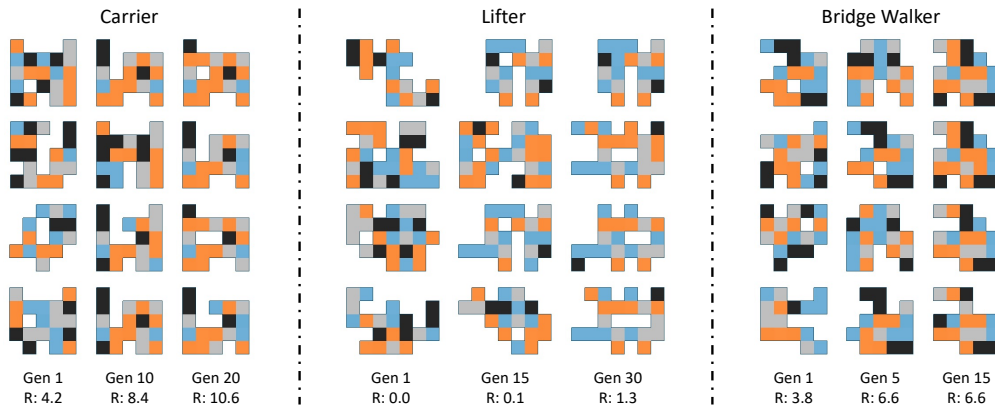


Figure 4: **Evolution of robot designs.** For each of the three selected tasks, we visualize the population in three different generations. Each column corresponds to one generation for which we show the four top performing robots along with their average reward.

301 a design with two legs and a structure up top that prevents the block from sliding off as the robot
 302 moves. A similar comparison pattern can be seen in the Lifter task, where the algorithm learns a
 303 parallel gripper-like shape at the bottom in order to manipulate an object. Unlike in the carrier task,
 304 the design structures that the algorithm generates are not prominently found in the initial generation.
 305 Finally, these patterns are echoed in the Bridge Walker task. Here the robot learns to evolve a large
 306 front foot to maximize their surface area and friction force to best walk across the soft rope bridge.

307 5.3 Comparison against hand-designed robots

308 We compare the performances of robots optimized by algorithm and the hand designed robots on
 309 several tasks to show the necessity of a co-design algorithm (Figure 5). The structure of the hand
 310 designed robots are bio-inspired and manually constructed according to our best intuition, and their
 311 control are optimized by PPO.

312 For every task, the hand designed robots are outperformed by at least one algorithm (usually more).
 313 For instance, for the Climber task we tested numerous natural robot design. However, none of them

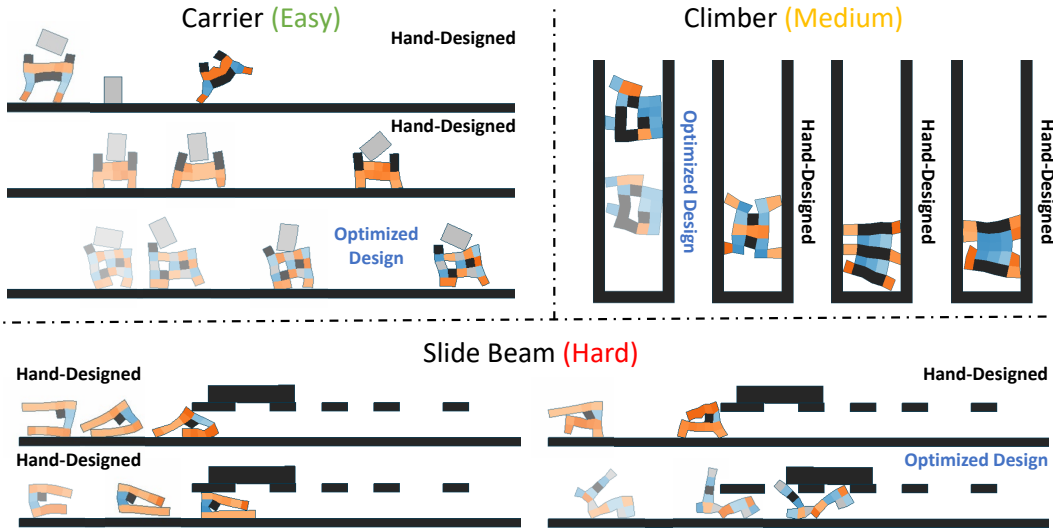


Figure 5: **Comparison between algorithm-optimized robots and hand designed robots on three tasks.** In each task, we visualize one robot optimized by the algorithm and several hand-designed robots.

314 make it very far up. The issue with our designs is that we could not find the right trade off between
 315 getting traction on the wall, and accelerating upwards. The genetic algorithm, however, is able to
 316 find this balance. It develops leg-like structures that help the robot make forward progress, as well
 317 as a long flat back that maximizes contact/frictional forces with the wall. Additionally, the genetic
 318 algorithm selects for having a hole in the center of its body, which helps it achieve a certain optimized
 319 walking motion.

320 For other tasks, the performance between the hand designed robots and the robots produced by
 321 the algorithms is quite comparable. This is the case with the Carrier robots, as a very natural
 322 hand-designed Carrier robot performs almost as well as the optimal robots produced by the design-
 323 optimization algorithms.

324 In the final case, there are tasks where neither a hand designed nor robot produced by the algorithm
 325 could achieve satisfying performance. One such environment is the Slider environment. For this task,
 326 many of the hand design robots fail to even achieve the first part of the goal and position themselves
 327 underneath the beam. While there is one robot produced by the genetic algorithm that does slide
 328 the beam across several pegs, it still comes nowhere close to achieving the optimal reward in the
 329 environment. This suggests that further work is needed in designing co-optimization algorithms that
 330 can complete these hard tasks.

331 **6 Conclusion and future work**

332 In this paper we proposed Evolution Gym, the first large-scale benchmark for evolving the structure
 333 and control of soft robots. Through the wide spectrum of tasks in Evolution Gym, we systematically
 334 studied the performance of current state-of-the-art co-design algorithms. As a result, we observed how
 335 intelligent robots could be evolved autonomously from scratch yet still be capable of accomplishing
 336 some surprisingly complex tasks. We also discovered the limitations of existing techniques for
 337 evolving more intelligent embodied systems.

338 There are several potential directions to be explored in the future. First, with the help of our proposed
 339 benchmark, it is desirable to develop more advanced co-design algorithms to solve the difficult tasks
 340 which the existing methods cannot address. Second, a robot will be considered more successful if it
 341 can perform multiple different tasks. Our benchmark suite naturally provides a comprehensive set of
 342 tasks and can potentially promote more exciting research work about multi-task or multi-objective
 343 robot co-design algorithms. Finally, since tasks in Evolution Gym are currently limited to either
 344 locomotion or manipulation, we plan to further extend Evolution Gym to additional task categories
 345 such as flying or swimming by incorporating new simulation capabilities.

References

- [1] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- [2] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- [3] Jacob Austin, Rafael Corrales-Fatou, Sofia Wyetzner, and Hod Lipson. Titan: A parallel asynchronous library for multi-agent and soft-body robotics using nvidia cuda. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7754–7760, 2020.
- [4] The GPyOpt authors. GPyOpt: A bayesian optimization framework in python. <http://github.com/SheffieldML/GPyOpt>, 2016.
- [5] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- [6] Nicholas Cheney, Jeff Clune, and Hod Lipson. Evolved electrophysiological soft robots. In *Artificial Life Conference Proceedings 14*, pages 222–229. MIT Press, 2014.
- [7] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. *SIGEVolution*, 7(1):11–23, August 2014.
- [8] Francesco Corucci, Nick Cheney, Francesco Giorgio-Serchi, Josh Bongard, and Cecilia Laschi. Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions, 2017.
- [9] Francesco Corucci, Nick Cheney, Francesco Giorgio-Serchi, Josh Bongard, and Cecilia Laschi. Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions. *Soft robotics*, 5(4):475–495, 2018.
- [10] Francesco Corucci, Nick Cheney, Hod Lipson, Cecilia Laschi, and Josh Bongard. Evolving swimming soft-bodied creatures. In *ALIFE XV, The Fifteenth International Conference on the Synthesis and Simulation of Living Systems, Late Breaking Proceedings*, volume 6, 2016.
- [11] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, pages 1329–1338. PMLR, 2016.
- [12] Christer Ericson. *Real-time collision detection*. CRC Press, 2004.
- [13] Jonathan Hiller and Hod Lipson. Dynamic simulation of soft multimaterial 3d-printed objects. *Soft robotics*, 1(1):88–101, 2014.
- [14] Zhiao Huang, Yuanming Hu, Tao Du, Siyuan Zhou, Hao Su, Joshua B. Tenenbaum, and Chuang Gan. Plasticinelab: A soft-body manipulation benchmark with differentiable physics, 2021.
- [15] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.
- [16] Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabás Póczos. Parallelised bayesian optimisation via thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142. PMLR, 2018.
- [17] Ilya Kostrikov. Pytorch implementations of reinforcement learning algorithms. <https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail>, 2018.
- [18] Harold J Kushner. A new method of locating the maximum point of an arbitrary multippeak curve in the presence of noise. 1964.

- 393 [19] Thomas Liao, Grant Wang, Brian Yang, Rene Lee, Kristofer Pister, Sergey Levine, and Roberto
394 Calandra. Data-efficient learning of morphology and controller for a microrobot. In *2019*
395 *International Conference on Robotics and Automation (ICRA)*, pages 2488–2494. IEEE, 2019.
- 396 [20] Alan McIntyre, Matt Kallada, Cesar G. Miguel, and Carolina Feher da Silva. neat-python.
397 <https://github.com/CodeReclaimers/neat-python>
- 398 [21] Eric Medvet, Alberto Bartoli, Andrea De Lorenzo, and Stefano Seriani. 2d-vsr-sim: A simulation
399 tool for the optimization of 2-d voxel-based soft robots. *SoftwareX*, 12:100573, 2020.
- 400 [22] Zbigniew Michalewicz. *Genetic algorithms+ data structures= evolution programs*. Springer
401 Science & Business Media, 2013.
- 402 [23] J. Močkus. On bayesian methods for seeking the extremum. In G. I. Marchuk, editor, *Optimiza-*
403 *tion Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, pages 400–404, Berlin,
404 Heidelberg, 1975. Springer Berlin Heidelberg.
- 405 [24] Andrew Nealen, Matthias Müller, Richard Keiser, Eddy Boxerman, and Mark Carlson. Physi-
406 cally based deformable models in computer graphics. In *Computer graphics forum*, volume 25,
407 pages 809–836. Wiley Online Library, 2006.
- 408 [25] Uber Research. pytorch-neat. <https://github.com/uber-research/PyTorch-NEAT>.
- 409 [26] Charles Schaff, David Yunis, Ayan Chakrabarti, and Matthew R Walter. Jointly learning
410 to construct and control agents using deep reinforcement learning. In *2019 International*
411 *Conference on Robotics and Automation (ICRA)*, pages 9798–9805. IEEE, 2019.
- 412 [27] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
413 policy optimization algorithms, 2017.
- 414 [28] Karl Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on*
415 *Computer graphics and interactive techniques*, pages 15–22, 1994.
- 416 [29] Andrew Spielberg, Allan Zhao, Yuanming Hu, Tao Du, Wojciech Matusik, and Daniela Rus.
417 Learning-in-the-loop optimization: End-to-end control and co-design of soft robots through
418 learned deep latent representations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-
419 Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*,
420 volume 32. Curran Associates, Inc., 2019.
- 421 [30] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of develop-
422 ment. *Genetic programming and evolvable machines*, 8(2):131–162, 2007.
- 423 [31] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting
424 topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- 425 [32] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press,
426 2018.
- 427 [33] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David
428 Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin
429 Riedmiller. Deepmind control suite, 2018.
- 430 [34] Merel Van Diepen and Kristina Shea. A spatial grammar method for the computational design
431 synthesis of virtual soft locomotion robots. *Journal of Mechanical Design*, 141(10), 2019.
- 432 [35] Tingwu Wang, Yuhao Zhou, Sanja Fidler, and Jimmy Ba. Neural graph evolution: Towards
433 efficient automatic robot design, 2019.
- 434 [36] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson
435 env: Real-world perception for embodied agents. In *Proceedings of the IEEE Conference on*
436 *Computer Vision and Pattern Recognition*, pages 9068–9079, 2018.
- 437 [37] Allan Zhao, Jie Xu, Mina Konaković-Luković, Josephine Hughes, Andrew Spielberg, Daniela
438 Rus, and Wojciech Matusik. Robogrammar: Graph grammar for terrain-optimized robot design.
439 *ACM Trans. Graph.*, 39(6), November 2020.

440 **Checklist**

- 441 1. For all authors...
- 442 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
443 contributions and scope? [Yes]
- 444 (b) Did you describe the limitations of your work? [Yes] See Section 6
- 445 (c) Did you discuss any potential negative societal impacts of your work? [N/A]
- 446 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
447 them? [Yes]
- 448 2. If you are including theoretical results...
- 449 (a) Did you state the full set of assumptions of all theoretical results? [N/A]
- 450 (b) Did you include complete proofs of all theoretical results? [N/A]
- 451 3. If you ran experiments...
- 452 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
453 mental results (either in the supplemental material or as a URL)? [Yes] The URL is
454 presented in the abstract.
- 455 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
456 were chosen)? [Yes] See Appendix D
- 457 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
458 ments multiple times)? [Yes] We ran experiments with multiple random seeds and
459 reported error bars. See Section 5
- 460 (d) Did you include the total amount of compute and the type of resources used (e.g., type
461 of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5
- 462 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 463 (a) If your work uses existing assets, did you cite the creators? [Yes] The existing code
464 implementation for our baseline algorithms are cited in Section 4
- 465 (b) Did you mention the license of the assets? [Yes] This benchmark platform will be
466 released under the MIT license. See Section 1
- 467 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
468 The URL is presented in the abstract.
- 469 (d) Did you discuss whether and how consent was obtained from people whose data you’re
470 using/curating? [N/A]
- 471 (e) Did you discuss whether the data you are using/curating contains personally identifiable
472 information or offensive content? [N/A]
- 473 5. If you used crowdsourcing or conducted research with human subjects...
- 474 (a) Did you include the full text of instructions given to participants and screenshots, if
475 applicable? [N/A]
- 476 (b) Did you describe any potential participant risks, with links to Institutional Review
477 Board (IRB) approvals, if applicable? [N/A]
- 478 (c) Did you include the estimated hourly wage paid to participants and the total amount
479 spent on participant compensation? [N/A]