



EvoRobogami: Co-designing with Humans in Evolutionary Robotics Experiments

Huang Zonghao

University of Pennsylvania

Philadelphia, Pennsylvania, United States of America

zoh@seas.upenn.edu

David Howard

CSIRO

Brisbane, Queensland, Australia

david.howard@csiro.au

Quinn Wu

University of Pennsylvania

Philadelphia, Pennsylvania, United States of America

quinnwu@seas.upenn.edu

Cynthia Sung

University of Pennsylvania

Philadelphia, Pennsylvania, United States of America

crsung@seas.upenn.edu

ABSTRACT

We study the effects of injecting human-generated designs into the initial population of an evolutionary robotics experiment, where subsequent population of robots are optimised via a Genetic Algorithm and MAP-Elites. First, human participants interact via a graphical front-end to explore a directly-parameterised legged robot design space and attempt to produce robots via a combination of intuition and trial-and-error that perform well in a range of environments. Environments are generated whose corresponding high-performance robot designs range from intuitive to complex and hard to grasp. Once the human designs have been collected, their impact on the evolutionary process is assessed by replacing a varying number of designs in the initial population with human designs and subsequently running the evolutionary algorithm. Our results suggest that a balance of random and hand-designed initial solutions provides the best performance for the problems considered, and that human designs are most valuable when the problem is intuitive. The influence of human design in an evolutionary algorithm is a highly understudied area, and the insights in this paper may be valuable to the area of AI-based design more generally.

CCS CONCEPTS

- Computer systems organization → Evolutionary robotics;
- Mathematics of computing → Evolutionary algorithms;
- Human-centered computing → User studies;
- Computing methodologies → Genetic algorithms.

KEYWORDS

Evolutionary Robotics, interactive evolution, MAP-Elites, quality diversity, design assist.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '22, July 9–13, 2022, Boston, MA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9237-2/22/07...\$15.00

<https://doi.org/10.1145/3512290.3528867>

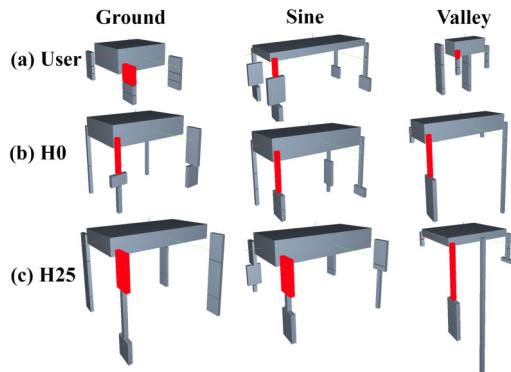


Figure 1: Example robots: (a) Human designs included in the initial population of H25. (b) High-performing designs in final population of H0. (c) High-performing designs in final population of H25. Red marks the left front leg.

1 INTRODUCTION

Evolutionary Robotics (ER) (e.g., [5]) is a powerful tool for robot design, being able to explore interwoven design spaces of coupled body, brain, and environmental interactions. Its fitness-based performance assessment is particularly useful in this role, as unintuitive, surprising designs [20] can be assessed in a bias-free manner and large, complex design spaces can be automatically explored in the pursuit of desired behaviours. ER is overwhelmingly implemented as a fully automated process: a robot-producing black box that relies on computational power, parallelisation, and extensive trial-and-error to tackle high-dimensional design problems over a wide design space. Conversely, human-centered design relies on creativity, intuition, and domain knowledge that digital systems often struggle to replicate, but typically a narrower design space.

Here, we investigate the impact of injecting human designs (and thus human creativity, intuition, and expertise) into the initial population of an ER experiment that produces legged robots (Figure 1). In this way, the human designer can influence the search for solutions that satisfy design requirements, whilst also exploiting the design space coverage afforded by evolutionary search. We implement a quality-diversity [28] evolutionary algorithm based on MAP-Elites

[26], providing an effective performance-design space map. We couple this to the Robogami [31] design tool, which enables interactive human-led design exploration and provides a straightforward way to integrate human designs into the evolutionary process.

We hypothesise that the inclusion of human designs in an ER experiment affects the running progress and the final result of the algorithm, either in terms of the performance of the evolved robots, the coverage of the map, or both. We focus the study mainly on two research questions: **RQ1** What effects would human input have on the intermediate population when algorithm is running? **RQ2** What effects would the human input have on the final population?

First, we collect human designs across a range of terrains using a visual design interface (Figure 5). Humans are tasked to design robot morphologies that generate high-performance locomotion behaviour across a range of increasingly difficult (for the robot) and decreasingly intuitive (for the designer) terrains (Figure 4). In stage 2, we conduct a range of experiments that use varying ratios of human and randomly-initialised designs. We assess the effect of changing ratios of human and random designs in the initial population. Overall, results indicate that human inputs could have either positive or negative impact on the final evolved robot fitness, depending on the quality of the designs, and a negative impact on early stage map coverage. Optimal ratios of human and random designs are shown to depend on environmental complexity.

The main contributions of this work are (i) the first work that uses human designs to evolve robot morphologies, (ii) extensions to the Robogami software that allow for coupling to MAP-Elites and a multi-participant user study, and (iii) a detailed analysis of the effects of adding human inputs to the initial population.

2 BACKGROUND

2.1 Computational Co-design

Human-influenced computational design aims to combine compute-based rapid assessment with human intuition and expert knowledge to collaboratively explore and optimise within a given design space. Typical goals include reducing designer effort and reaching parts of the design space that are otherwise difficult to access. Applications are diverse, including optimising CAD models [32], designing model aeroplanes [37], and assisting in the sketching of levels for computer games [22]. Visualisation plays an important role, particularly when dealing with indirect solution representations which may be unintuitive to explore without visual and performance feedback [14]. Works closely related to ER include interactive design of 3D printable mechanical characters [8] that reproduce desired animatronic motions through the computationally-assisted placement of actuation mechanisms, and robotic creatures [25] including legged robots [9], where computational design is used to generate a plausible gait for a given hand-designed morphology. In all cases, the human predominantly leads and controls the design process, with the computational element used in a supporting role, e.g., to generate suggestions, verify the design, or support the designer by creating working actuation for their designs.

2.2 Interactive Evolutionary Algorithms

Human input can be integrated into an Evolutionary Algorithm (EA) in several ways. For example, expert knowledge can be used to

set parameter limits and to design fitness functions [17]. Typically, the only feedback the user receives on their choices is at the end of the experiment, when solutions can be analysed to see if these settings produced the intended results. Interactive Evolution [3] is concerned with more in-depth interplay between user input and evolutionary processes. Early examples include Sims' 'Galapagos' exhibit where a human's interest (measured by the time spent looking at a specific screen displaying an evolved art piece) was used to drive the evolutionary process, and indirectly-represented digital art [24], as well as the interactive evolution of dynamical systems [34]. Interactive EAs have applications as diverse as molecular design [19] to evolution of digital images [33] and brochures [29] to game level design [6, 30]. Work on exploring the design space around a provided CAD model [7], e.g., injecting user designs and then evolving them, demonstrates the benefits of mixing user inputs and evolution. We also note the success of interactive evolution in the domain of parametric design [13], which is our target domain.

2.2.1 Interactive Quality-Diversity. Quality-Diversity (QD) algorithms [28] are a family of evolutionary algorithms that aim to produce a wide range of high-quality designs, with two popular variants being NSLC [21] (using Pareto optimisation to maintain diversity) and MAP-Elites [26] (which uses a discretised feature map). QD is particularly suited to design tasks, as a wide variety of performant designs are encouraged to effectively map out an entire design space [12], generating valuable feedback to the designer and identifying regions where interesting solutions might lie. MAP-Elites has also been investigated in the context of interactive evolution, where game level designs can be selected from the map and edited by the user during an evolutionary run [1], with preliminary work showing the ability for human designs to unlock previously undiscovered areas of the design space.

2.2.2 Interactive Evolutionary Robotics. Evolutionary Robotics (ER) (see numerous overviews e.g., [5, 10, 27]) studies the automatic generation of robot morphology and control within a given environment, primarily to generate suitable behaviours, or as a tool to study theories of embodiment, e.g., [15]. Originally simulated [35], ER is now frequently associated with fabrication and physical instantiation of both rigid [23] and soft robots [16] via 3D printing. Robogen is a pertinent UI/simulation enabled evolutionary robotics platform [2]. In our work, we use the Robogami software, which was designed for user interaction and which offers a direct route to eventual fabrication-based studies through 3D printing or folding of the resulting robots through autogenerated fabrication plans. *Interactive* evolutionary robotics is sparsely covered in the literature. ER automation of certain design tasks can reduce barriers to robot design [4], for example, by handling controller optimisation of non-adaptive morphology, by using controller input to add new rules on-the-fly for an evolved classifier-based robot controller [18], and interactive approaches based on cellular representations [11]. To the best of our knowledge, interactive evolutionary morphology generation does not appear in the literature.

2.3 Literature Summary and Motivation

Our work sits at the intersection of Evolutionary Robotics and Computational Co-Design: users interactively design solutions, which

are then harnessed to improve an evolutionary process. Rather than tweaking computed designs, users iteratively tweak their own designs based on fitness and visual behavioural feedback from the simulator. This paper details the first such experiment in an evolutionary robotics context. Our approach combines the benefits of evolved and human designs in a way that does not require the user to constantly interact with the design software to 'tweak' designs, nor stop the evolutionary process to wait for user input. Compared to conventional computational co-design, the emphasis for discovery is shifted more onto the computational (evolutionary) element, which is responsible for ultimately evolving high-performance solutions yet is guided by its given input designs.

3 SYSTEM OVERVIEW

Our study builds a closed-loop system that let users first build robots via a interactive design tool backed by the Robogami software [31], then use a MAP-Elites [26] based searcher to evolve their designs and generate better solutions.

3.1 Robogami Designer

Robogami is a interactive robotic design tool based on a design-by-composition framework that let users design robots by composing independently manufacturable robot parts together. The software also includes algorithm that provides interactive feedback to users, guiding their exploration by checking validity and manufacturability of the design at each step. This provides a novice user a smooth approach to get started and generate a design that reflects their wish. To integrate this front-end to our system loop, we (1) simplify the user interface and add design constraints to the software to limit the design space as described in Section 3.2; (2) link the software to a simulator described in Section 3.3 that evaluates the design's fitness and provides user with visual feedback; (3) develop a compiler that compiles the user design to the design vector that can be loaded by the MAP-Elites searcher described in Section 3.4; (4) wrap the software with a user study guiding system that helps the users go through the user study described in Section 4 with minimum interruption by the researchers.

The resulted UI is shown in Figure 5(a).

3.2 Robot Morphology and Gaits

ER uses a range of robot representations, which can be classified as direct or indirect. Our choice must balance genome complexity and design space expressiveness with the ability for changes in genome parameters to be easily understood by corresponding changes in the robot to aid in human design exploration. As such, we use a direct representation that captures these features.

In this study, we consider legged robots with 2-6 legs, each of which has 2 or 3 links. The genome, allele ranges, and mapping to a phenotype (robot) are presented in Figure 2(a). The genome vector contains 17 to 53 variables depending on the complexity of the robot. Body Shape ID and Link Shape ID correspond to rectangular prisms of different aspect ratios, and the Body Scale and Link Scale are multipliers directly applied to the corresponding dimensions of the parts, providing a flexible and scalable design space suitable for both human and evolutionary exploration.

Legs are evenly distributed on each side of the robot's body and attached to the middle point of the z direction (See Figure 4 for direction definitions). A layout-mirroring flag controls which side of the robot has more legs when an odd number of total legs is inputted. This flag is ignored for robots with even number of legs. Leg links are connected to each other in series with joints located at either end of the link. All joints are revolute joints with axes of rotation parallel to the y axis. The robot parts are assumed all to be of uniform density (2.5 g/cm^3), with dimension details shown in Figure 2.

The gait of the robot is directly determined by the morphology, with a set of predefined joint movements based on the design and layout of legs. Each of the legs follows a motion sequence depending on the number of links in the leg (Figure 3) with an offset depending on its placement on the robot body. In particular, the legs are divided into two groups. Starting from the front left leg, alternating group numbers are assigned in clockwise order. Group 1 executes the motion sequence $M_1 M_2 M_3$ simultaneously as Group 2 executes $M_2 M_3 M_1$, and the process repeats over multiple gait cycles. The group switches to the next joint target when all of the joint angles are within 0.01 rad of the target angle, or 3 s after the other group finishes its current motion, whichever happens earlier. The joint angles and motion sequences associated with the joints in the legs are kept constant over all robot designs; that is, they are not included in the evolutionary design. The gait is controlled by a PI controller over both velocity ($k_p = 10, k_i = 0.3$) and position ($k_p = 2, k_i = 0$), with a maximum joint velocity of 2 rad/s . A deterministic controller was used to remove a potential level of indirection, in the form of body-brain optimisation, from the user design problem, allowing a predictable control output per morphology to simplify the design process.

3.3 Task and Fitness

To make sure the participants of the user study could get good intuition on some problems, we define the task for the robot to be simply walking forward in different environments, and define the fitness to be heavily determined by the distance traveled during the given time.

At the start of each simulation, the robot is placed at a fixed point at one end of the environment, with z position determined by having the foot of its longest leg lifted above the terrain surface directly under it by 2 cm . The robot is then simulated forward for 30 s of simulated time with simulation step size of 0.005 s . We use Project Chrono [36] as our simulator.

Since we expect the impact of human-created designs to differ depending on the complexity of the design task and the quality of human intuition, we designed three environments Ground (G), Sine (S), Valley (V) (Figure 4) with increasing degrees of difficulty of walking forward. Ground is expected to be the easiest task for a human as the environment has little impact on robot behaviour. Sine introduces additional difficulty by requiring designers to think how navigate robots through the bumps. Valley is the most difficult due to the introduced asymmetry, which is counter to most human designers' intuition. The fitness objective f is to maximize the distance of the robot moving forward while minimizing divergence

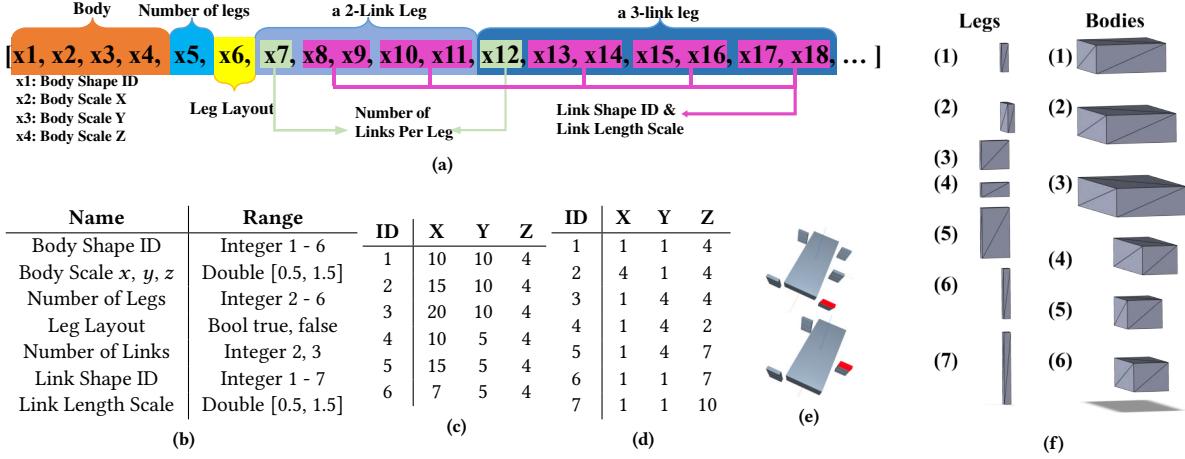


Figure 2: (a) Visual depiction of the robot genome, (b) allele definitions and ranges, (c) body part dimensions (cm), (d) body part dimensions (cm), (e) example robots highlighting the potential for asymmetric morphologies, (f) body and link shapes.

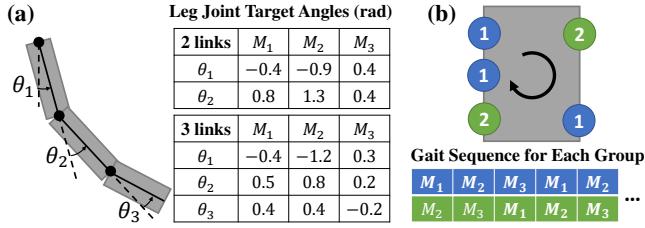


Figure 3: Gait executed by the robot designs. (a) Target joint angles for legs of different numbers of links. (b) Example grouping of legs and corresponding motion sequence. The front of the robot faces up.

from a straight line:

$$f = \Delta x - 0.5 \text{ abs}(\Delta y) \quad (1)$$

where Δx is the net displacement in the forward direction and Δy is the net displacement perpendicular to the forward direction at the time when the simulation stops. If the system detects the robot has fallen off the edge of the environment, the simulation terminates with fitness taken from the robot's last position.

3.4 MAP-Elites Searcher

We build the evolutionary searcher upon MAP-Elites, and link it to the simulator described in Section 3.3 for fitness and feature evaluation. MAP-Elites holds its population in the archive map, which is a discretized low-dimensional projective space of the search space. We use a 2D map of 20×20 cells, with the 2 dimensions (features) being: (1) the length of the body along the x direction (Figure 4 for direction definition), and (2) the standard deviation of the leg lengths. The length of the robot's body directly affects its separation of legs on the forward direction, which then decides robot's maximum feasible leg length and stability given our controller (Figure 3). The standard deviation of the leg lengths quantifies asymmetry of the robot design. These two features were chosen through a small number of test runs using random initial populations. Compared to

other tested features, e.g., average length of legs, maximum length of legs, total volume of robots, length of genome, etc., these two were found to be two of the most significant factors related to the environments and fitness (Eq. 1) we use.

We set the size of initial population to be 30 and run the algorithm for 2000 iterations. For each iteration, 30 parent designs are randomly selected from the archive map to generate a batch of 30 children designs with mutation rate of 0.1 and crossover rate of 0.75. The children designs are evaluated according to a fitness metric and projected back onto and stored in the archive map. If the projected cell is occupied, or if two children designs are projected to the same cell, the individual with the highest fitness is retained.

4 HUMAN DATA COLLECTION

A user study¹ was conducted to collect human designs with the UI described in Section 3.1. Data was collected from 13 students (5 undergraduate, 8 graduate) who all had some prior experience in robot design. Each participant started with a 10-min. tutorial explanation of the user interface on a sample design, then a 10-min. exploration session to familiarize themselves with the system and ask questions, in which they could design and test robots for an unlimited number of times on a different training environment. After the training, each participant was given the task and the definition of fitness, and asked to design a robot for each of the three environments. At the start of each new environment, all options on the right panel of the UI are set to the middle point of their range, showing the default "neutral robot." The participant would adjust the parameters and design a robot that reflects their intuition and knowledge, and send the robot to simulation for walking visualisation and fitness. Each participant was allowed to update the design 10 times per environment, and each update was considered to include all of design changes made between two consecutive simulation requests. Design genomes were saved at each iteration, resulting in 11 total designs per environment per user.

¹This user study has been approved by the IRB of University of Pennsylvania under the title *Robotic design decision making* (protocol #849907).

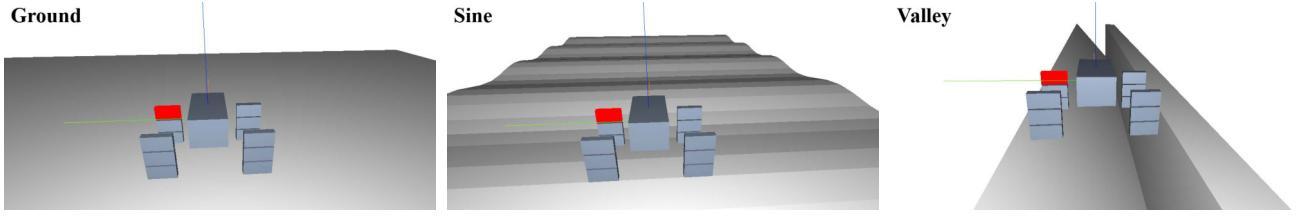


Figure 4: Test environments: Ground (G), Sine (S), and asymmetric Valley (V) with one vertical wall and one wall at 45°. Neutral robots provided for scale. Red marks the left front leg of robot. +x points forward, +y points leftward, +z points upward

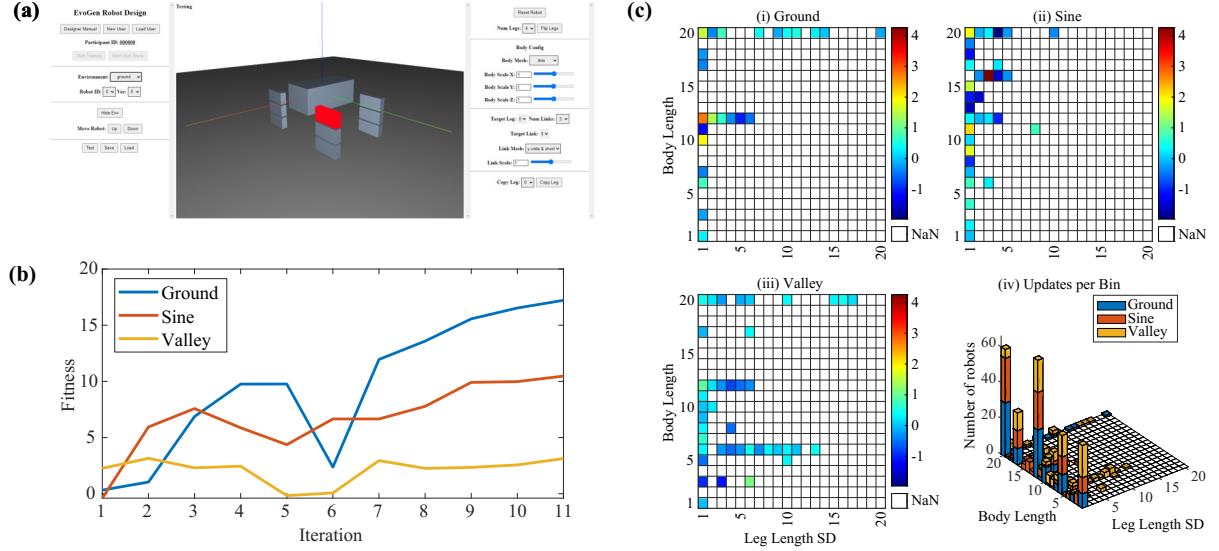


Figure 5: (a) Robot design interface. Options on the right let user tweak the design of robot. The resulting design is displayed in the center. Options on the left guides the user through the user study, change the environment, load a pre-existing design, and simulate the current design. When “Test” button is clicked, a visualisation of the simulation and the final fitness value are shown to the user. (b) Sample fitness progression over 10 design iterations for one user. (c) (i)-(iii) Archive map of all human designs per environment, coloured by fitness. (iv) Number of robots per bin on the map, coloured by environment type.

The order of the environments was randomized for each user to mitigate the bias from one environment to the next, introduced by the increasing familiarity to the system.

We expect the fitness of the design to improve over the iterations up until the user is satisfied with their design quality, at which point they commit to their design by simulating the same design for the remaining iterations or they test minor tweaks.

A total of 143 designs (11 robots from 13 users) were collected for each environment. In a post-processing step, we removed duplicates from when the user requested a simulation multiple times for the same design, resulting in 125 total designs for the Ground, 128 for the Sine, and 130 for the Valley.

Figure 5(b) shows the fitness progression of the design iterations on different environments for one participant. As expected, for the Ground and Sine, designs improved over the course of design iterations for the first few iterations before fitness increases slow down near the final value. The robot designs also undergo larger changes during earlier iterations than later ones. The general fitness increases support the suitability of our approach as a design

tool. In the Valley, however, the fitness achieved by the participant remained low and throughout the iterations and the designs exhibited larger changes between iterations, indicating that users performed more exploration for this task and that it was overall more difficult for them to use their intuition on.

Figure 5(c) (i) - (iii) show the archive maps with all designs from all users for the corresponding environment are inputted. A large number of designs are concentrated on the left side of the map corresponding to symmetric designs for all environments. Comparing to the sample archive maps of the final population of H25 shown in Figure 7, we see the users successfully placed their designs in the high-performing region for Ground and Sine, but not for the Valley, indicating the less productive intuition for the latter. And we see more cells in Figure 5(c)(iii) are filled, as the users were forced to explore more areas in the design space. Figure 5(c)(iv) shows the number of robots that goes into each cell in the feature map, with all robots included. The preference of cells is another indicator that humans' exploration of the design space is relatively condensed.

Env	Mean Fitness					Human / Final H0
	Human	H25 I	H25 F	H0 I	H0 F	
G	7.4	5.76	8.09	-0.64	6.5	113.85%
S	5.99	4.61	6.25	-0.56	5.79	103.45%
V	1.24	1.02	5.83	-0.06	6.52	19.02%
G(B)	11.96	11.96	17.29	3.08	12.19	89.87%
S(B)	9.2	9.2	13.89	3.05	10.68	104.31%
V(B)	3.14	9.2	10.84	2.16	13.45	26.87%

Table 1: Mean fitness for the human designs, Initial (I) and final (F) population of H0 and H25. In the first column, the letters indicate the environment: G(round), S(ine), or V(alley). We also show the best fitness of human design, and the mean of best single fitness of different test runs (B).

5 RESULTS

To explore **RQ1** and **RQ2**, we set up 5 test conditions with varying numbers of human and random designs.

- (H0) No human designs and 30 random robots.
- (H5) 5 human designs and 25 random robots. Human inputs were the top 5 designs, with at most 1 design per user.
- (H15) 15 human designs and 15 random robots. Human inputs were the top 15 designs, with at most 2 designs per user.
- (H25) 25 human designs and 5 random robots. Human inputs were the top 25 designs, with at most 3 designs per user.
- (H30) 30 human designs and 0 random robots. Human inputs were the top 30 designs, with at most 3 designs per user.

Each test condition Hx is repeated 10 times with random designs uniformly re-sampled for each repeat. The experiments are conducted on a desktop computer with a quad core 10th Gen Intel i7 @ 3.4GHz CPU, taking approximately 6 hours per run.

The curves in Figure 6(a) show the transitional changes of key metrics during the running of algorithm. Box plots in Figure 6(b) show key MAP-Elites metrics of the final population. [26] Figure 7 shows sample archive maps from H0 and H25 for three environments, and Figure 1 shows initial and high-performing robots sampled from those archive maps. Figure 8 shows the Mann-Whitney U-test result for mean and best fitness of the final population of different test conditions.²

5.1 (RQ1) Effects on intermediate population

Mean fitness marks the average performance of the population, a good indicator of the progress of the algorithm run. The random H0 test condition follows a generally increasing trend. However, test conditions with human input shows sharp dips in fitness in the first few iterations: the human designs are relatively high fitness and early additions to the map decrease the mean. For Ground and Sine, higher amounts of human input maintain higher mean and best fitness. For the Valley, the difficulty of the design task reduces effectiveness of the human input and the fitness curves largely overlap.

Best fitness is the fitness of the highest-performing robot of the population, which directly shows the quality of solution that the

algorithm has found at any point of the run. The test conditions with human inputs all start at the same high fitness value on Ground and Sine, as the best human design is guaranteed to be included in the initial population. The curve for H0 starts at a significantly lower point than the rest, and also converges to a lower value. In contrast, for the Valley, the curve of H0 is indistinguishable from the others. The differences between H0 and other conditions at the start of the curve in different environments also indicates that the Ground and Sine are easy environments for the users.

QD-score is a performance measurement that considers both quality and diversity of the population [28]. For the Ground, all test conditions have a QD-score curve starting roughly at the same point, but test conditions with human designs increase at a faster rate, which the largest final QD-score being achieved by H25 on average. The results of the Sine show a similar trend, but with narrower gap between H25 and H0. In contrast, for the Valley, H0 achieves the best QD-score shortly after a few iterations and holds the position for the rest of the run. This implies human designs with bad quality bring negative effect to QD-score.

Coverage curves show a rapid increase in early iterations, followed by a slower convergence to 100% coverage. Lower amounts of human input achieve higher coverage faster, however 100% coverage is attained within 400 iterations in all cases, indicating that in our experiment setup, human input affects only the rising speed of coverage.

5.2 (RQ2) Effects on final population

For the Ground and Sine, human designs help to achieve better QD-score, global performance, reliability and precision in all test conditions. And in most cases, H25 outperforms other test conditions in all metrics. For the Valley, H0 records the best QD-score, precision and reliability. For the Ground and Sine, the human designs had a positive effect on the mean fitness of the final archive map, and clearly beats H0 for the best fitness. However, for the Valley, the H0 initial population produced an archive map with a mean fitness that was statistically significantly higher than almost all other test conditions. The low mean fitness of the user designs (1.24) compared to the final mean fitness of the H0 runs (6.52) indicates that users found the Valley task difficult and biased the search to an unpromising area of the design space.

Results indicate that using only human designs in the initial population can hinder the search. For the Ground, H30 condition resulted in a statistically significantly lower mean fitness in the archive map compared to H25. While for the Sine, the mean fitness of the two archive maps were approximately the same, indicating no great benefit from the additional human designs. This implies that although human intuition was valuable in directing the search to particular portion of the archive map, there is a benefit to adding random robots to diversify the population and encourage design exploration. Importantly, the evolutionary process does not dilute the benefits of using user designs; stark fitness benefits are observed even after 2000 iterations of evolution with no additional interaction. Human designs are also frequently replaced in their niche by evolved designs; evolution uses the provided genetic material to improve solutions throughout the map.

²Supplementary data and resources are available at: <https://sung.seas.upenn.edu/publications/evorobogami-gecco-2022>.

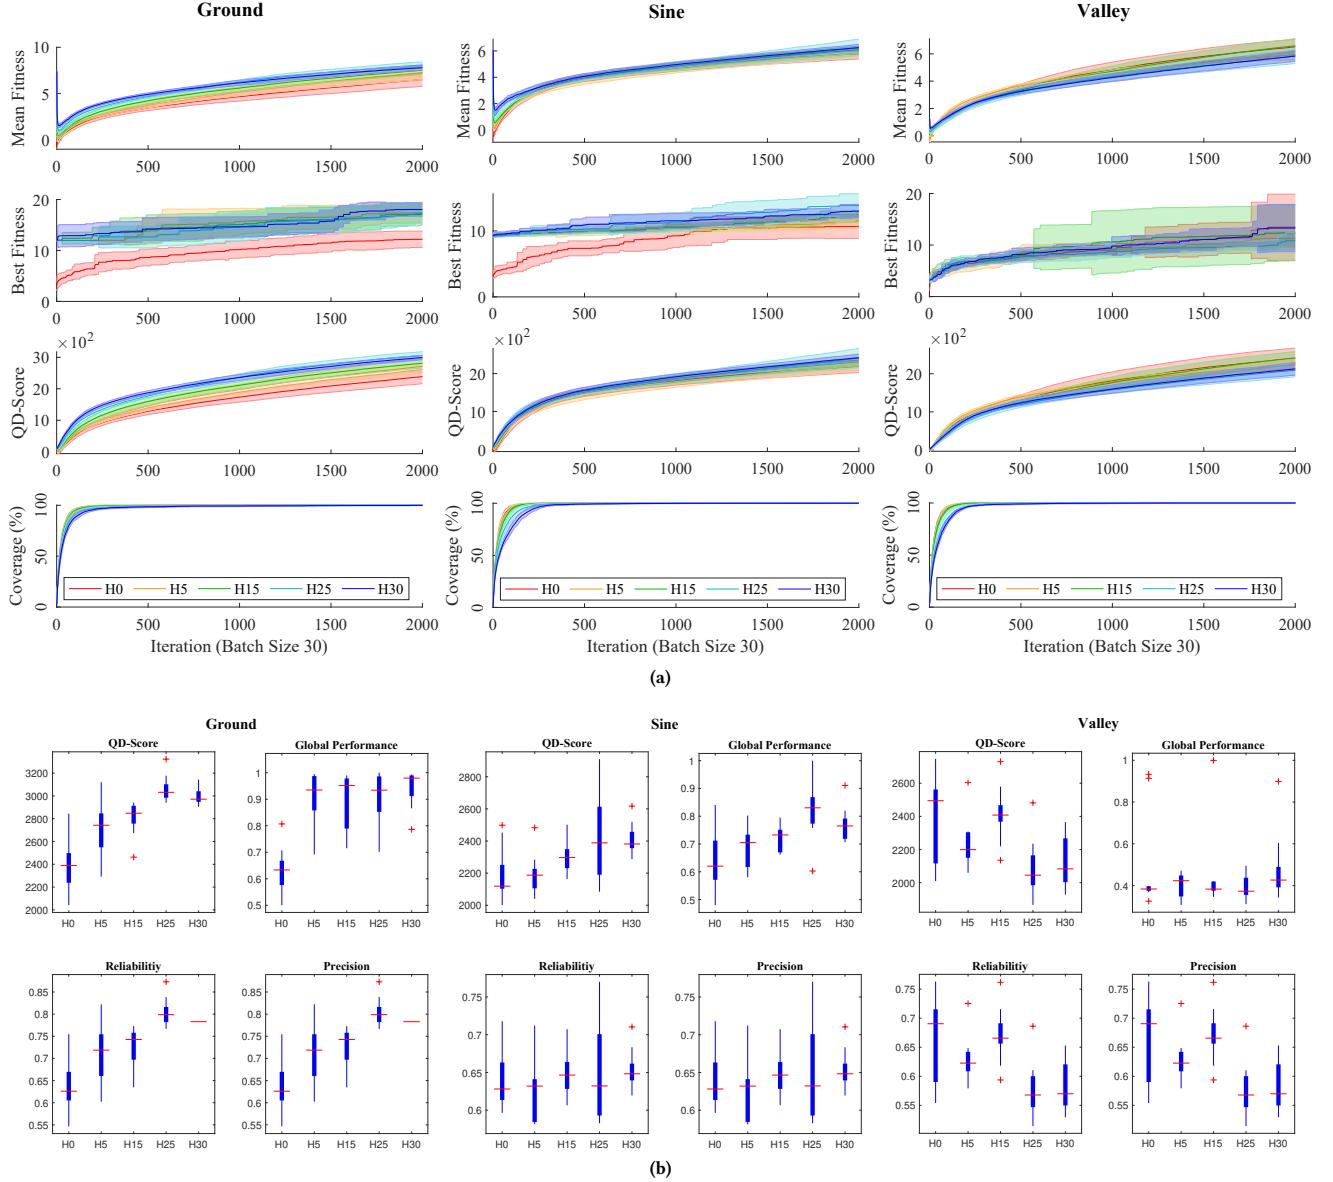


Figure 6: Results for 10 runs of each test condition. (a) Mean archive fitness, best single fitness of archive, QD-score, and coverage at each iteration. Shaded area marks ± 1 standard deviation. **(b)** Box plots of final QD-score, global performance, reliability, and precision of each test condition. Red line marks median, + indicates outliers.

Given best fitness performance, we conclude that human inputs are beneficial to evolution in all cases, but more so when the user has a sound grasp of the problem to be solved. When this is the case, the human inputs bias the search process towards high-performing regions of the design space, as well as providing useful genetic material for further tinkering by evolution.

6 DISCUSSION

In summary, this paper shows the effects of adding human designs to the initial population of an evolutionary robotics experiment. Our

approach is ‘set and forget’: it does not require constant/repeated user interaction to achieve its stated benefits, and is an accessible way to bring intuition into evolutionary design problems without slowing down evolutionary run time. It can also be viewed as the inverse of related approaches wherein computational designs are suggested to the user for tweaking; here the human does the suggesting and evolution does the tweaking.

Results demonstrate that when humans have a reasonable understanding of the problem to be solved, their inputs have beneficial effects on the subsequent evolutionary process, yielding higher

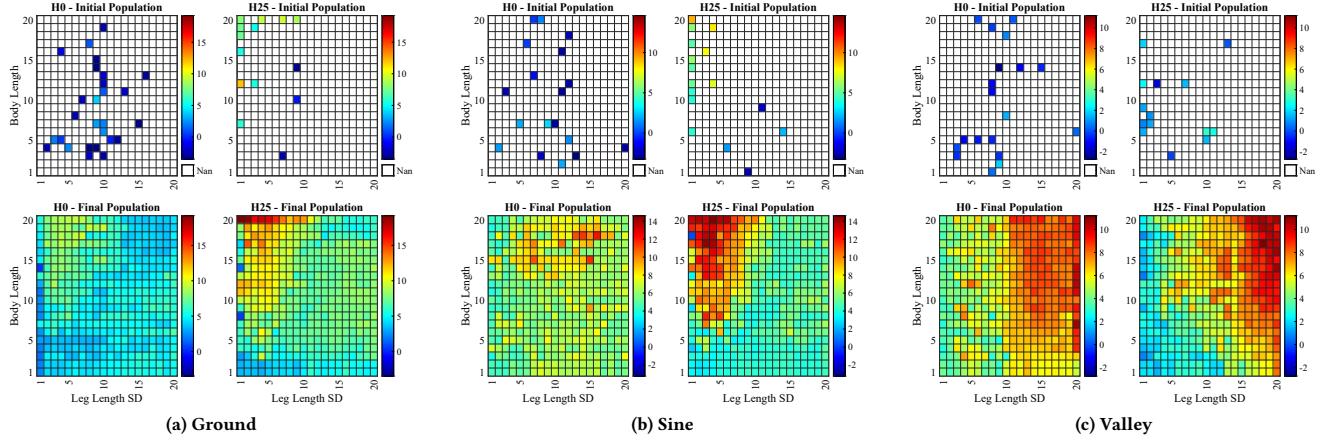


Figure 7: Representative maps from a single run with H0 and H25 test cases. Colour represents fitness.

A		H5	H15	H25	H30	B		H5	H15	H25	H30
H0	+	+	+	+	+	H0	+	+	+	+	+
H5	+		+	+		H5		~	~	+	
H15		+		+		H15			+	+	
H25				-		H25				+	

(a) Ground

A		H5	H15	H25	H30	B		H5	H15	H25	H30
H0	~	+	+	+	+	H0	+	+	+	+	+
H5	+		+	+		H5		+	+	+	
H15		+		+		H15			+	+	
H25			~	H25						-	

(b) Sine

A		H5	H15	H25	H30	B		H5	H15	H25	H30
H0	-	+	+	-	-	H0	-	-	-	-	-
H5	+		-	-		H5	+	-	+		
H15		-		H15		H15		-	+		
H25			~	H25					+		

(c) Valley

Figure 8: Statistical significance test of final archive maps for three environments. "+" means the column result is larger than the row result. "-" indicates the opposite. "~" means difference is less than 0.5% of the smaller value. Shaded cells indicate statistically significant differences confirmed by Mann-Whitney U-test at $P < 0.05$. "A" indicates comparison on mean fitness of archive map; "B" indicates comparison on best single fitness of archive map.

best fitness and higher mean archive fitness. Evolution biases the initial population into high-performing regions of the search space (a form of expert knowledge injection), and subsequently exploits the high-quality genetic material to generate performant designs. We note that evolution plays the key designer role in this process; initial human inputs are not seen in the final populations and are frequently overtaken in their niches by evolved designs based on

their genome parameters. We also see that the fitness benefits are not simply wiped out by the ongoing genetic process; instead fitness benefits are lasting and readily observable in the final evolved populations. Additionally, archive coverage is not unduly affected by the use of user designs. Maximum map coverage is readily achieved in all cases within 400 iterations, so would only be detrimental in very short runs that are commonly not used with MAP-Elites. A mixture of human and random designs (H25) achieves good results in all metrics for Ground and Sine, where the human intuition is reasonable, suggesting an optimal balance of the two sources of designs.

We note some limitations with this study. First, it is dependent on human input and a cohort of experienced users, which may not be available. However this limitation is true of any user-based study. Furthermore, many CAD packages and robot simulators can readily support this functionality, and many potential users exist in academia (with access to students) or industry (with access to peer employees). The range of applications can therefore be considered to be relatively large. Finally, results suggest that our approach is highly beneficial, but is particularly powerful when the human designers have a reasonable grasp of the problem. This effect is still to be quantified, but nevertheless this observation may guide the application of this technique into other domains.

Future work will focus on assessing the impact of representation on a user's ability to design, and the interaction of different representations and evolutionary methods with the process. We note that Robogami is capable of much more design diversity than used in this study, so expanding the design space is a key goal. Additional studies will explore questions such as: Are there regions that humans can reach that evolution struggles with, or *vice versa*? Similarly, we wish to exploit the ability to easily fabricate the robots and move some of the experimentation to reality.

ACKNOWLEDGMENTS

The work was supported in part by the National Science Foundation Grant #1845339.

REFERENCES

[1] Alberto Alvarez, Jose Font, Steve Dahlskog, and Julian Togelius. 2021. Assessing the Effects of Interacting with MAP-Elites. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 17, 1 (Oct. 2021), 124–131. <https://ojs.aaai.org/index.php/AIIDE/article/view/18899>

[2] Joshua Auerbach, Deniz Aydin, Andrea Maesani, Przemyslaw Kornatowski, Titus Cieslewski, Grégoire Heitz, Pradeep Fernando, Ilya Loshchilov, Ludovic Daler, and Dario Floreano. 2014. Robogen: Robot generation through artificial evolution. In *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*. MIT Press, 136–137.

[3] Wolfgang Banzhaf. 2000. Interactive evolution. *Evolutionary Computation* 1 (2000), 228–234.

[4] Josh Bongard, Paul Beliveau, and Gregory Hornby. 2012. Avoiding Local Optima with Interactive Evolutionary Robotics. In *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation* (Philadelphia, Pennsylvania, USA) (GECCO '12). Association for Computing Machinery, New York, NY, USA, 1405–1406. <https://doi.org/10.1145/2330784.2330955>

[5] Josh C Bongard. 2013. Evolutionary robotics. *Commun. ACM* 56, 8 (2013), 74–83.

[6] Luigi Cardamone, Daniele Loiacono, and Pier Luca Lanzi. 2011. Interactive evolution for the procedural generation of tracks in a high-end racing game. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. 395–402.

[7] J. Clune, A. Chen, and H. Lipson. 2013. Upload any object and evolve it: Injecting complex geometric patterns into CPPNs for further evolution. In *2013 IEEE Congress on Evolutionary Computation*. 3395–3402. <https://doi.org/10.1109/CEC.2013.6557986>

[8] Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. 2013. Computational Design of Mechanical Characters. *ACM Trans. Graph.* 32, 4, Article 83 (jul 2013), 12 pages. <https://doi.org/10.1145/2461912.2461953>

[9] Ruta Desai, Beichen Li, Ye Yuan, and Stelian Coros. 2018. Interactive co-design of form and function for legged robots using the adjoint method. *arXiv preprint arXiv:1801.00385* (2018).

[10] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E Gusz Eiben. 2015. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI* 2 (2015), 4.

[11] Fr ed eric Grua and Kameel Quatramaran. 1996. *Cellular encoding for interactive evolutionary robotics*. Technical Report.

[12] Adam Gaier, Alexander Aстроth, and Jean-Baptiste Mouret. 2018. Data-efficient design exploration through surrogate-assisted illumination. *Evolutionary computation* 26, 3 (2018), 381–410.

[13] John Harding and Cecile Brandt-Olsen. 2018. Biomorpher: Interactive evolution for parametric design. *International Journal of Architectural Computing* 16, 2 (2018), 144–163. <https://doi.org/10.1177/1478077118778579> arXiv:<https://doi.org/10.1177/1478077118778579>

[14] Frank Hoisl. 2012. *Visual, interactive 3D spatial grammars in CAD for computational design synthesis*. Ph.D. Dissertation. Technische Universität München.

[15] David Howard, Agoston E Eiben, Danielle Frances Kennedy, Jean-Baptiste Mouret, Philip Valencia, and Dave Winkler. 2019. Evolving embodied intelligence from materials to machines. *Nature Machine Intelligence* 1, 1 (2019), 12–19.

[16] David Howard, Jack O'Connor, Jordan Letchford, James Brett, Therese Joseph, Sophia Lin, Daniel Furby, and Gary W. Delaney. 2021. Getting a Grip: in-Material Evolution of Membrane Morphology for Soft Robotic Jamming Graspers. [arXiv:2111.01952 \[cs.RO\]](https://arxiv.org/abs/2111.01952)

[17] Gerard Howard, Thomas Lowe, and Wade Geles. 2020. Diversity-based design assist for large legged robots. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*. 81–82.

[18] D. Katagami and S. Yamada. 2002. Interactive evolutionary robotics from different viewpoints of observation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 2. 1108–1113 vol.2. <https://doi.org/10.1109/IRDS.2002.1043879>

[19] Eric-Wubbo Lameijer, Joost N Kok, Thomas Bäck, and Ad P IJzerman. 2006. The molecule evaluator. An interactive evolutionary algorithm for the design of drug-like molecules. *Journal of chemical information and modeling* 46, 2 (2006), 545–552.

[20] Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Lee Altenberg, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Beslon, David M Bryson, et al. 2020. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life* 26, 2 (2020), 274–306.

[21] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. 211–218.

[22] Antonios Liapis, Georgios N Yannakakis, and Julian Togelius. 2013. Sentient sketchbook: computer-assisted game level authoring. (2013).

[23] Hod Lipson and Jordan B Pollack. 2000. Automatic design and manufacture of robotic lifeforms. *Nature* 406, 6799 (2000), 974–978.

[24] Jon McCormack et al. 1993. Interactive evolution of L-system grammars for computer graphics modelling. *Complex Systems: from biology to computation* (1993), 118–130.

[25] Vittorio Megaro, Bernhard Thomaszewski, Maurizio Nitti, Otmar Hilliges, Markus Gross, and Stelian Coros. 2015. Interactive Design of 3D-Printable Robotic Creatures. *ACM Trans. Graph.* 34, 6, Article 216 (oct 2015), 9 pages. <https://doi.org/10.1145/2816795.2818137>

[26] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).

[27] Stefano Nolfi and Dario Floreano. 2000. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press.

[28] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. 2016. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3 (2016), 40.

[29] Juan C. Quiroz, Amit Banerjee, Sushil J. Louis, and Sergiu M. Dascalu. 2009. *Document design with interactive evolution*. 309–319. https://doi.org/10.1007/978-3-642-02937-0_28 Copyright: Copyright 2009 Elsevier B.V., All rights reserved.

[30] Jacob Schrum, Jake Gutierrez, Vanessa Volz, Jialin Liu, Simon Lucas, and Sebastian Risi. 2020. Interactive Evolution and Exploration within Latent Level-Design Space of Generative Adversarial Networks. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (Cancún, Mexico) (GECCO '20). Association for Computing Machinery, New York, NY, USA, 148–156. <https://doi.org/10.1145/3377930.3389821>

[31] Adriana Schulz, Cynthia Sung, Andrew Spielberg, Wei Zhao, Robin Cheng, Eitan Grinspun, Daniela Rus, and Wojciech Matusik. 2017. Interactive robogami: An end-to-end system for design of robots with ground locomotion. *The International Journal of Robotics Research* 36, 10 (2017), 1131–1147. <https://doi.org/10.1177/0278364917723465> arXiv:<https://doi.org/10.1177/0278364917723465>

[32] Adriana Schulz, Jie Xu, Bo Zhu, Changxi Zheng, Eitan Grinspun, and Wojciech Matusik. 2017. Interactive design space exploration and optimization for cad models. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14.

[33] Jimmy Secretan, Nicholas Beato, David B D Ambrosio, Adelein Rodriguez, Adam Campbell, and Kenneth O Stanley. 2008. Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1759–1768.

[34] Karl Sims. 1992. Interactive evolution of dynamical systems. In *Toward a practice of autonomous systems: Proceedings of the first European conference on artificial life*. 171–178.

[35] Karl Sims. 1994. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. 15–22.

[36] Alessandro Tasora, Radu Serban, Hammad Mazhar, Arman Pazouki, Daniel Melanz, Jonathan Fleischmann, Michael Taylor, Hiroyuki Sugiyama, and Dan Negrut. 2015. Chrono: An open source multi-physics dynamics engine. In *International Conference on High Performance Computing in Science and Engineering*. Springer, 19–49.

[37] Nobuyuki Umetani, Yuki Koyama, Ryan Schmidt, and Takeo Igarashi. 2014. Pteromys: Interactive design and optimization of free-formed free-flight model airplanes. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.