Modular Robot Design Optimization with Generative Adversarial Networks

Jiaheng Hu, Julian Whitman, Matthew Travers and Howie Choset
The Robotics Institute, Carnegie Mellon University

Introduction

Modular robots inspire a concept where instead of bringing a family of robots into the field, one can bring a toolbox of modules which are configured, and perhaps re-configured, to provide a customized system design for a particular task. This unique benefit poses many challenges, including the need to determine a particular system design for each given task, often with limited computation and time.

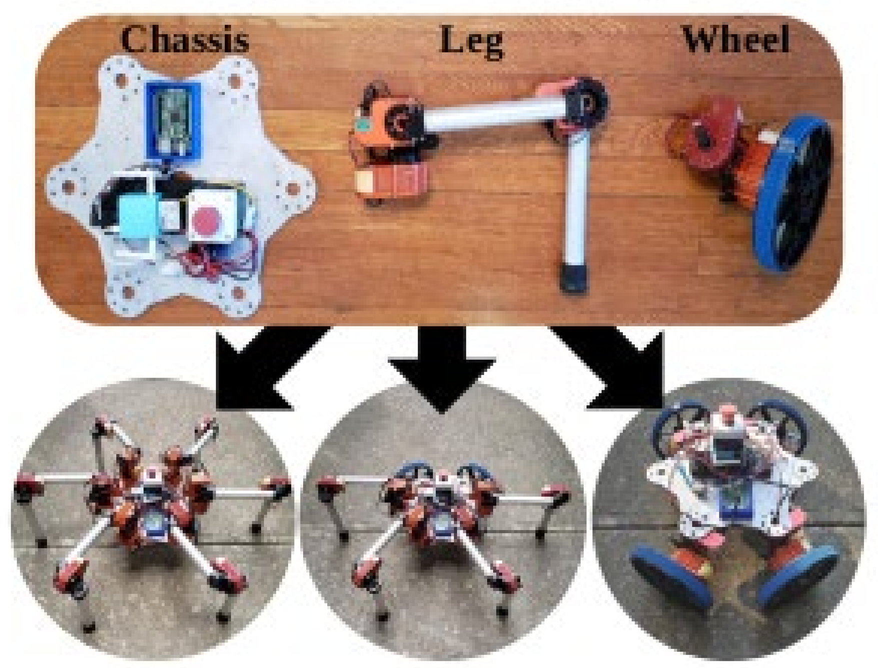

A set of modules - body, legs, and wheels – can generate many different robots system configurations.

Previous work in design automation [1, 2, 3] achieve efficient run-times by learning a one-to-one mapping from task to design. However, robot design problems are often multimodal, where multiple distinct designs can equivalently achieve a task. This paper presents a novel method based on generative adversarial networks (GANs) that learns a one-to-many mapping from task to a distribution of designs, and can be used to generate diverse and high quality designs in a computationally-efficient manner.

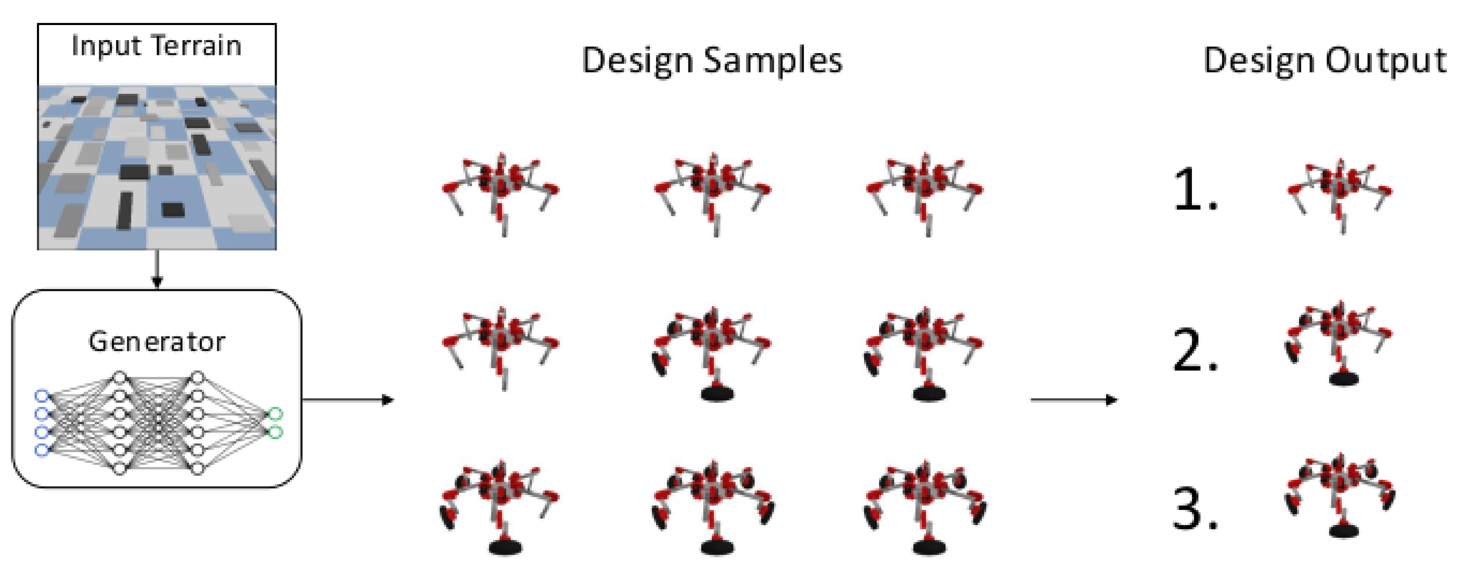

We present a generative approach to create task-conditioned modular robot designs. The generator takes in a task description (in this example, a terrain height map, left) and outputs a distribution of designs, where high-performing designs have high probability mass.

RoboGAN training pipeline

At each training iteration, a task (here, a terrain to traverse) is generated at random and abstracted into a 2-D height map. The generator implicitly maps this height map into a population of designs, depicted in the red box. The self-guided data creation step explores around the generated designs by evolving them using a procedure inspired by Evolutionary Algorithms [4], and creates a population of evolved designs, depicted in the green box. The discriminator takes as input the terrain and a robot design, which may be either from the generated designs or the evolved designs. The discriminator learns to distinguish which population the design comes from. The output of the discriminator feeds into the loss L_G, guiding the generator towards generating high-performing designs.

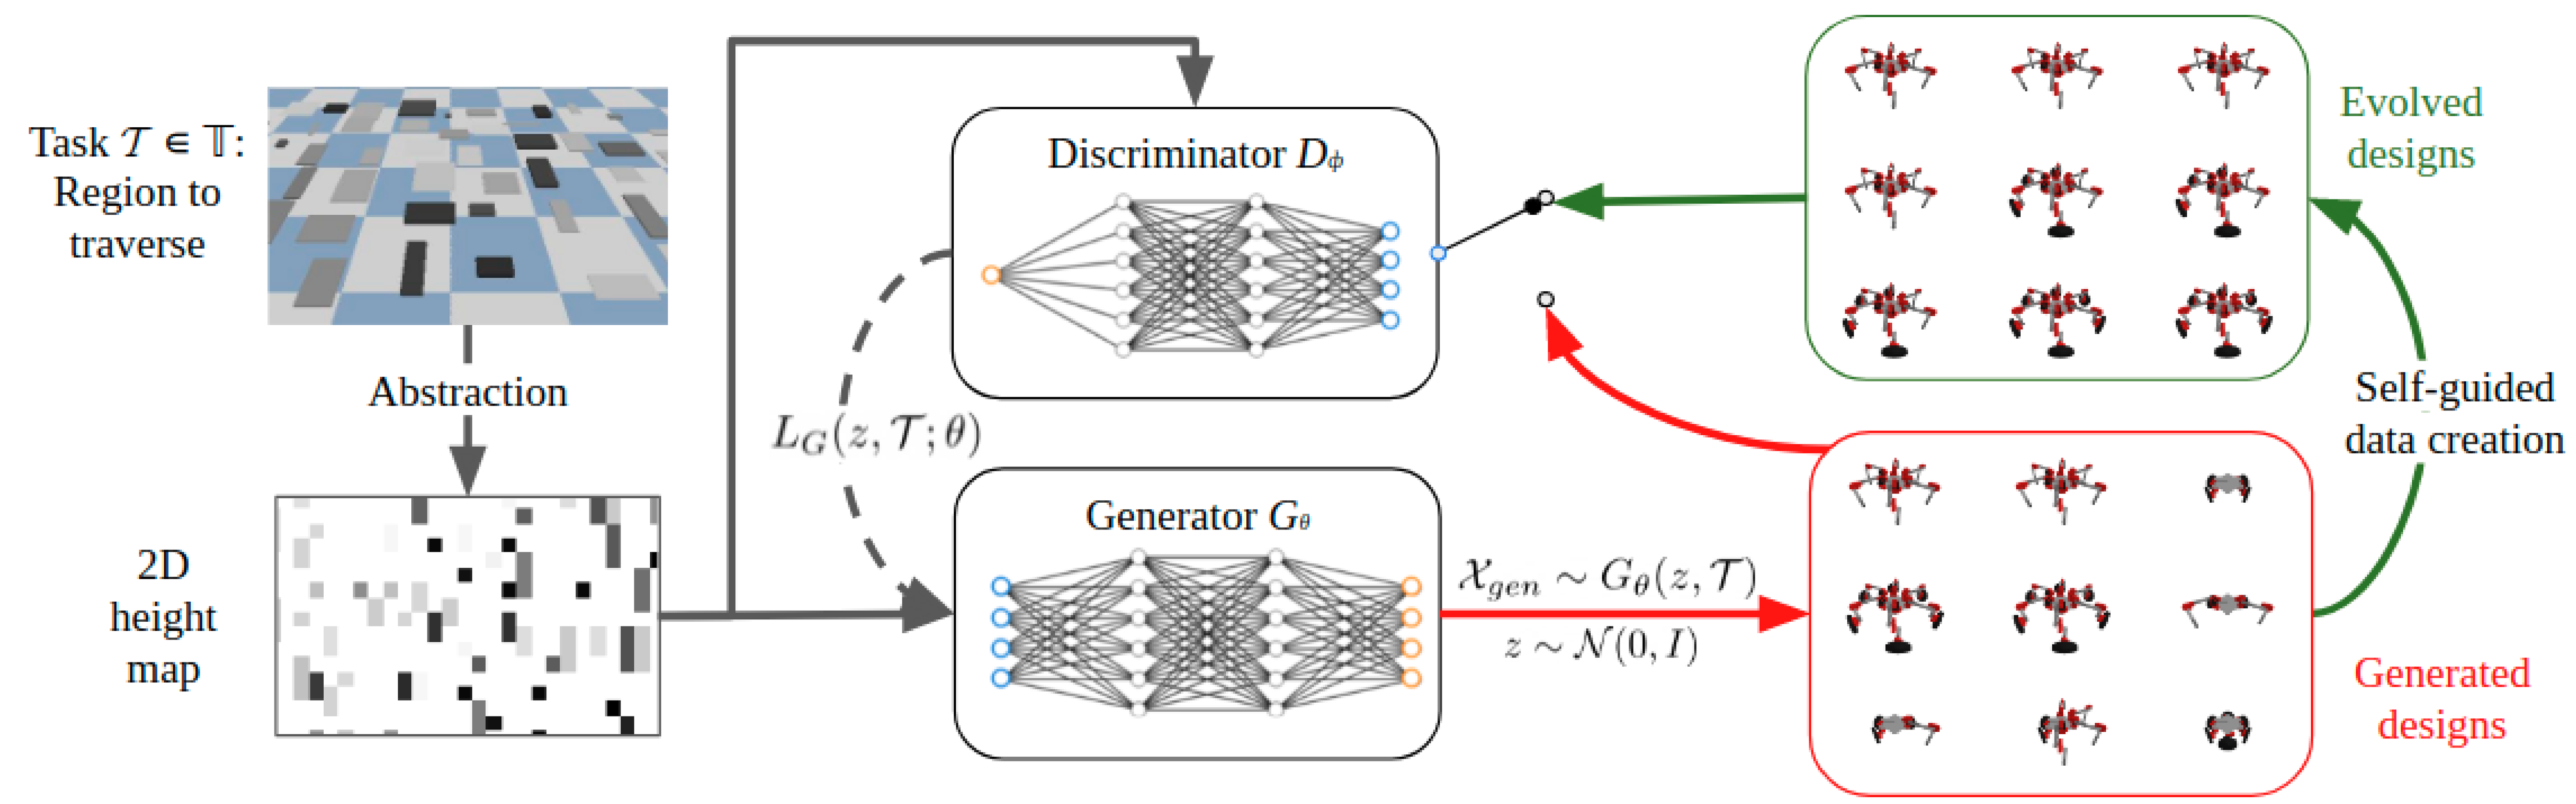

Analysis of the latent space

The trained generator learns a mapping from a latent space to designs conditioned on the task. The latent space mapping serves as a distribution over designs that we can use to decide which design to use for the given task. This figure shows the output of the generator for different points in the latent space given a single terrain map input.

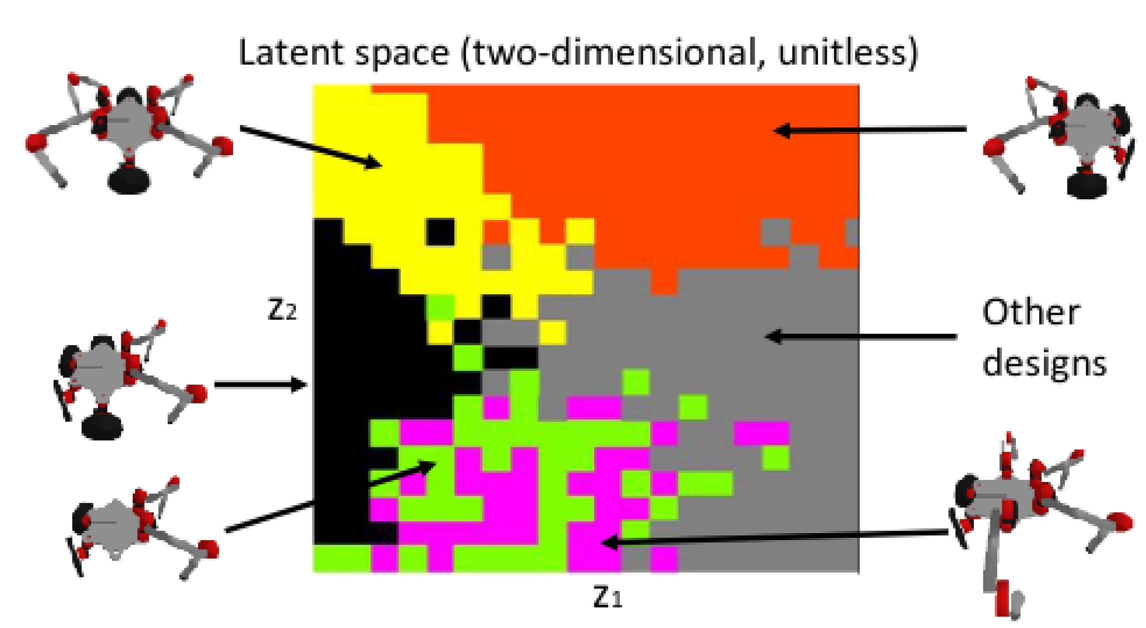

Comparisons to related methods

Compared to previous learning-based approaches, our algorithm produces designs that are superior both in performance and solution diversity.

Method	Max. Performance (m/s)	Number of unique designs	Deployment Performance (m/s)
Fit2form [1]	4.03 ± 0.91	1.2±0.2	2.55±1.31
Q-learning [2]	4.31±0.26	2.7±0.9	3.10±0.92
MolGAN [3]	3.82±0.55	1.8±0.3	2.93±1.00
RoboGAN	**5.73±0.22**	**5.8±1.5**	**5.36±0.52**

References

[1] H. Ha, S. Agrawal, and S. Song, "Fit2Form: 3D generative model for robot gripper form design, in Conf. on Robotic Learning, 2020.

[2] J. Whitman, R. Bhirangi, M. Travers, and H. Choset, "Modular robot design synthesis with deep reinforcement learning," in Proc. of the AAAI Conf. on Artificial Intelligence, 2020

[3] N. De Cao and T. Kipf, "MolGAN: An implicit generative model for small molecular graphs," ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models, 2018

[4] G. S. Hornby, H. Lipson and J. B. Pollack, "Evolution of generative design systems for modular physical robots," Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation