Transfer Learning for Faster Tensegrity Robot Gaits Discovery

Akshay S. Kashyap John A. Rieffel (Advisor)
Computer Science Department, Union College

Abstract

Tensegrity Robots are a class of soft robots capable of dynamic locomotion. Previous work has shown how optimization techniques can be used to discover optimal gaits in a given environment to move as efficiently as possible. These techniques, however, need to be re-applied to every new environment the robot is deployed in. This research explores using Transfer Learning to speed up this discovery process.

Introduction, Motivation, and Background

Tensegrity Robots are a class of soft robots that are composed of interlocked sets of rigid struts and tensional connections such as springs that give them resilience against damage, the ability to quickly recover from deformations, and adaptability. The modular design of these robots also allows them to be low-cost, reusable, and easy to build. NASA is considering using Tensegrity Robots for planetary exploration [1] and researchers at UCSD are deploying these for indoor air duct maintenance [2].

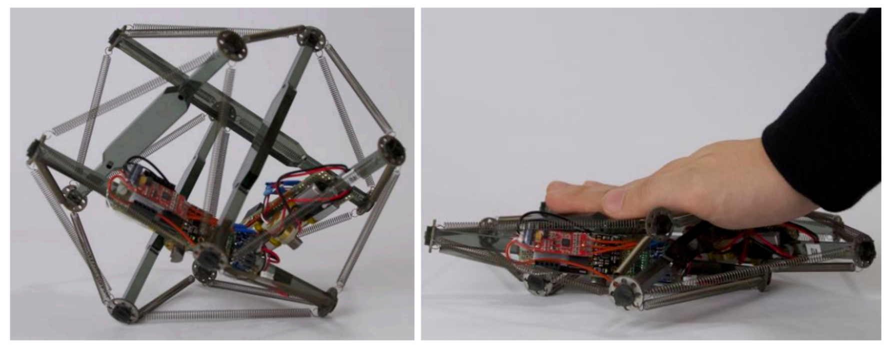

Figure 1: Tensegrity Robot (Source: Union College EvoRobo Group.)

Their physical properties enable them to be nimble and move in ways impossible for rigid robots. Data-driven methods, such as Bayesian Optimization [4], can be used to model the speed of the robot as a function of its gaits to find optimal gaits. An issue is that optimal gaits need to be rediscovered in every new environment the robot is deployed in. This research explores using a transfer learning approach with Bayesian Optimization (BO) to speed up the discovery by using past learning experiences.

Question

Can transferring knowledge from previous learning experiences help speed up the gait discovery process in new environments and damaged states?

Problem Statement and Proposed Solution

For a vibrationally-driven tensegrity robot, a gait can be characterized by one or more motor parameters. We will be optimizing for just frequency in this research. Let $g = (f_{m_1}, f_{m_2}, f_{m_3})$ be a gait of a robot with 3 motors, Let $D = [f_{min}, f_{max}]$ be the domain of motor frequency values. $f_{m_i} \in D$ and $g \in D^3$.

Let S be the source task, the condition the robot has already discovered an optimal gait in using BO. Let T be the target task, the new condition the robot is being deployed in and needs to discover optimal gaits in. Let $p^S(g)$ and $p^T(g)$ be the performance of the robot (measured by speed, in this research) as functions of its gait in S and T respectively.

We will use the transfer learning framework for Bayesian Optimization introduced by Joy TT, et al. [3], called **Efficient-BO**, which models source observations as noisy measurements of target function:

$$y_i^S = p^T(g_i^S) + \epsilon_i^S, \forall i = 1 \ldots n$$

It does this by using the Source task GP to compute noise variance, which is added to the kernel of the Target task GP. (Note: We will refer to BO without transferred knowledge as **Plain-BO**.)

Methods

PyGPGO: An open-source Bayesian Optimization library. This was re-engineered and extended to use the Efficient-BO algorithm.

Tensegrity Simulator: To facilitate the analysis of complex tensegrity dynamics, we built as Tensegrity Simulator platform to model new structures and simulate their physical interactions with the environment, using C++ and the ODE dynamics simulation library.

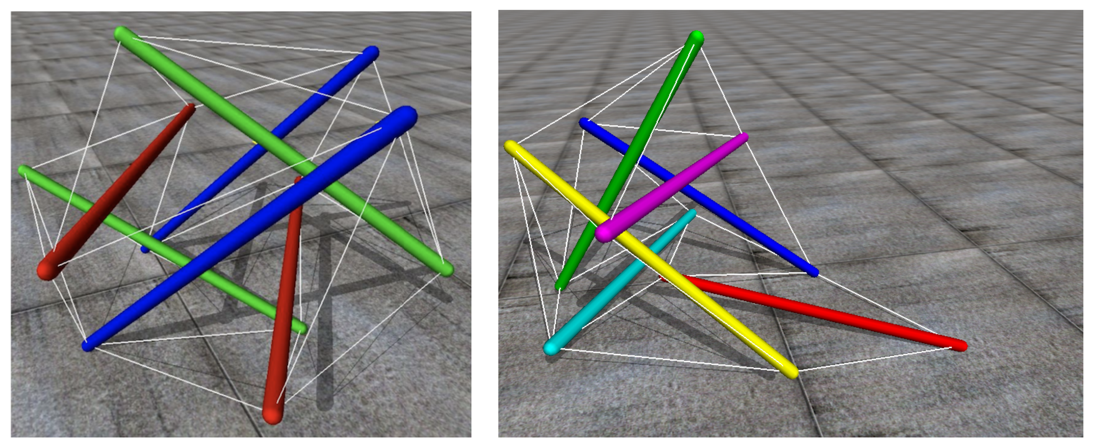

Figure 2: Simulations. Six Strut Tensegrity (L); Broken Robot (R)

Experiment

An experiment consists of two setups S and T, each with different environment variables or a damaged state of the robot such as a broken spring or a damaged strut.

1. Step 1: Perform gait discovery using Plain-BO in S.
2. Step 2: Perform gait discovery in T with Plain-BO and with Efficient-BO.

The expectation is that Efficient-BO will perform better. The experimental setups for evaluation are as follows:

Experiment	Source Task S	Target Task T
1	Gravity: -0.1 units Friction: 0.5 units	Gravity: -0.5 units Friction: 0.75 units
2	Plain Ground Terrain	Hilly Terrain
3	Undamaged Tensegrity	Tensegrity with Broken Spring

Table 1: Experimental Setups

Results

Each graph below shows the speed of the robot vs. optimization trials and a box plot of max speed achieved. (Blue = Plain-BO in S, Orange = Plain-BO in T, Green = Efficient-BO in T)

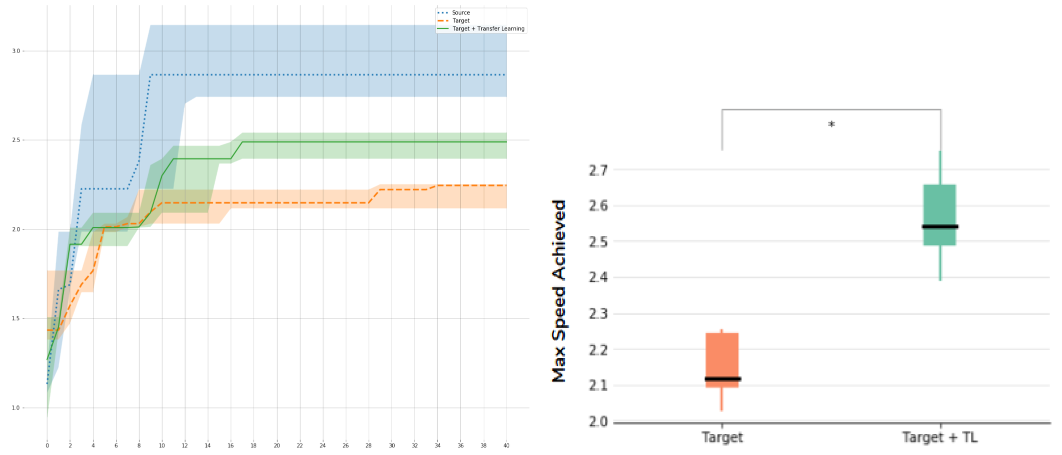

Figure 3: Experiment 1

Efficient-BO achieves a **max speed 11.3% higher** in **50% fewer trials**. **p-value** < 0.05.

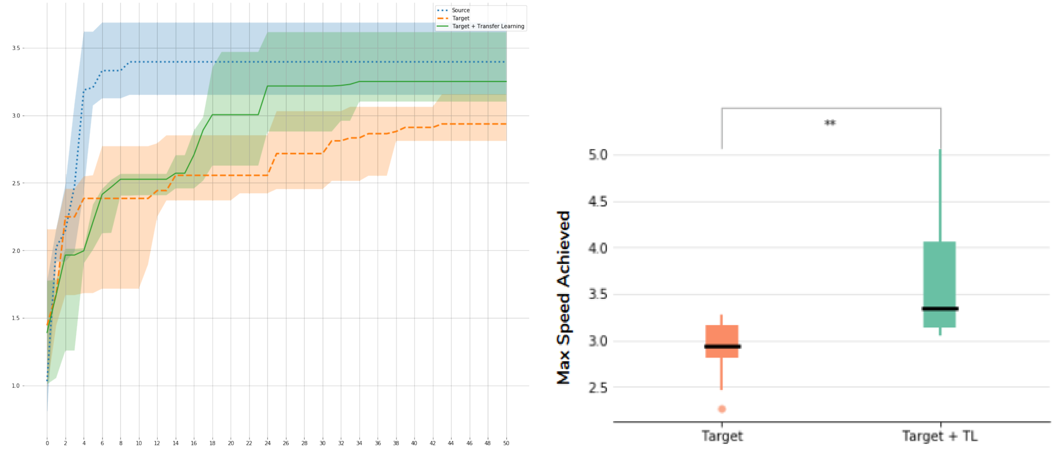

Figure 4: Experiment 2

Efficient-BO achieves a **max speed 12.1% higher** in **50% fewer trials**. **p-value** < 0.01.

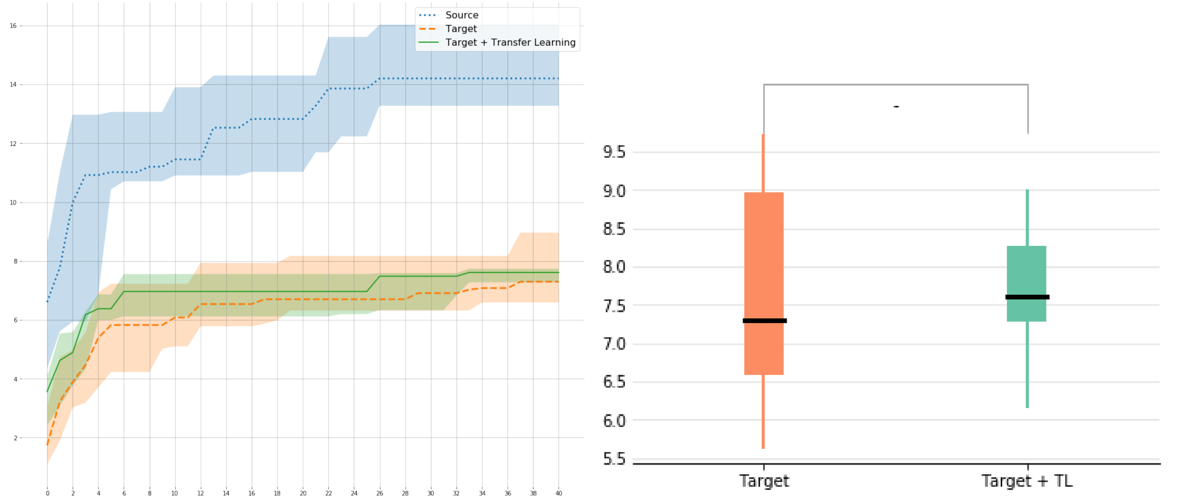

Figure 5: Experiment 3

Efficient-BO achieves a **max speed 4.31% higher** in **50% fewer trials**. **p-value** > 0.05.

Analysis

The results show that using transferred knowledge does improve gait discovery in new, unseen environments. The improvement in the damaged state was not statistically significant, and this could be because the damaged tensegrity has dynamics different from the one intact and there was therefore minimal transfer of useful knowledge.

Conclusion

In conclusion, this research presents a solution for making tensegrity robots more adaptable and resilient to change, and hopes to help make Tensegrity Robots usable for solving practical problems in the real world, (or as NASA hopes, on the Moon.)

References

[1] Bruce, J., Caluwaerts, K., Iscen, A., Sabelhaus, A. P., and SunSpiral, V. Design and evolution of a modular tensegrity robot platform. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on* (2014), IEEE, pp. 3483–3489.

[2] Friesen, J., Pogue, A., Bewley, T., de Oliveira, M., Skelton, R., and Sunspiral, V. Ductt: A tensegrity robot for exploring duct systems. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on* (2014), IEEE, pp. 4222–4228.

[3] Joy, T. T., Rana, S., Gupta, S. K., and Venkatesh, S. Flexible transfer learning framework for bayesian optimisation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining* (2016), Springer, pp. 102–114.

[4] Rieffel, J., and Mouret, J.-B. Soft tensegrity robots. *arXiv preprint arXiv:1702.03258* (2017).