Avoiding Local Optima with User Demonstrations

Shane Celis[1], Greg Hornby[2], Josh Bongard[3]

University of Vermont

Abstract

Interactive Evolutionary Algorithms (IEA) use human input to help drive a search process. Traditionally, IEAs allow the user to exhibit preferences among some set of individuals. Here we present a system in which the user directly demonstrates what he or she prefers. Demonstration has an advantage over preferences because the user can provide the system with a solution that would never have been presented to a user who can only provide preferences. However, demonstration exacerbates the user fatigue problem because it is more taxing than exhibiting preferences. The system compensates for this by retaining and reusing the user demonstration, similar in spirit to user modeling. The system is exercised on a robot locomotion and obstacle avoidance task that has an obvious local optimum. The system is compared against a general and a specific fitness function designed to remove the local optimum. We show that our proposed system outperforms most variants of these completely automatic methods, providing further evidence that Evolutionary Robotics (ER) can benefit by combining the intuitions of human users with the search capabilities of computers.

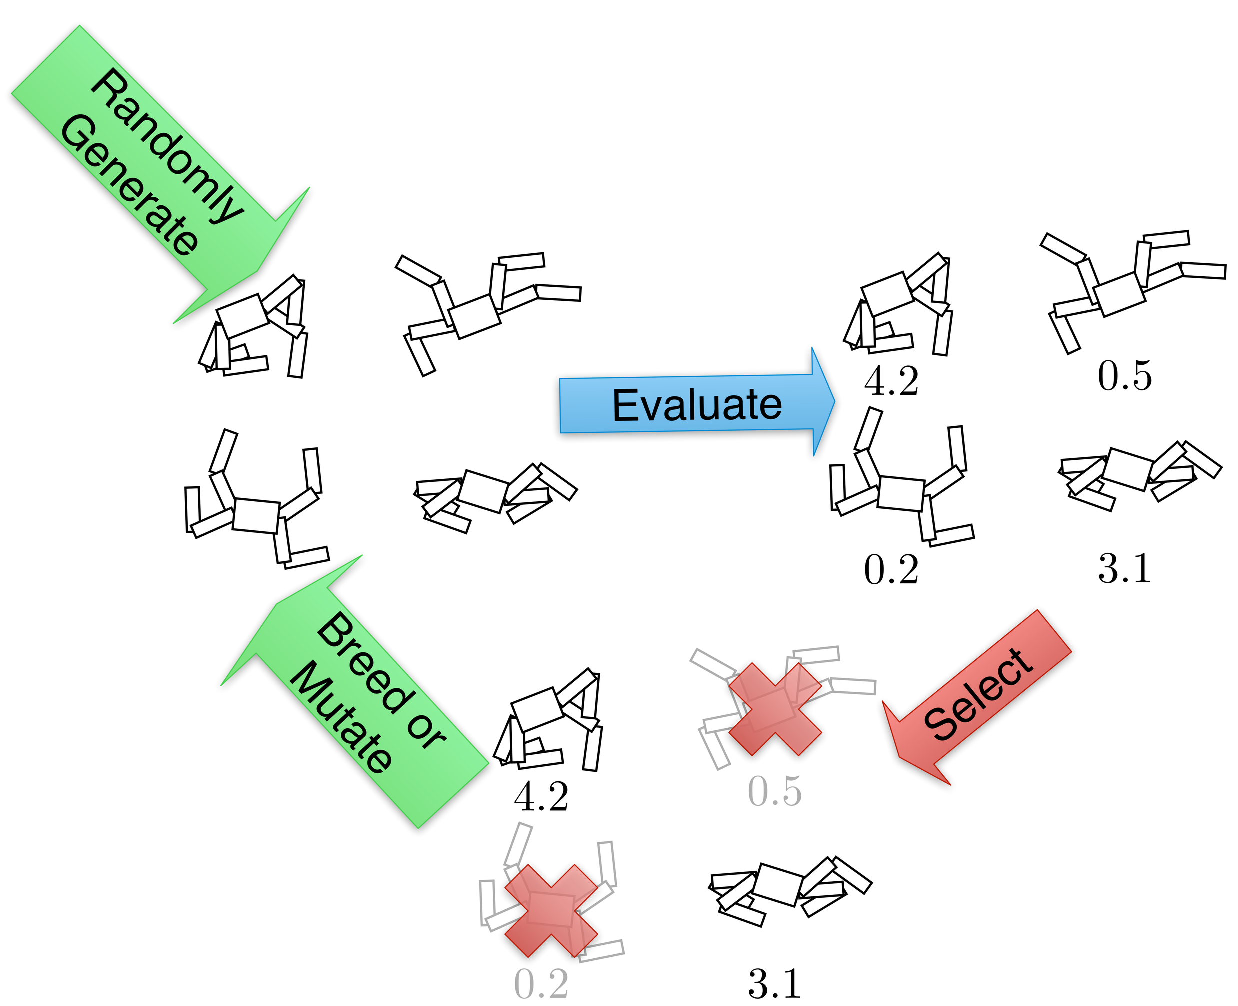

Figure 1: Basic outline of an evolutionary algorithm: Step 1) Randomly generate a population of robots. Step 2) Evaluate performance on a task, e.g., distance traveled. Step 3) Select best performers. Step 4) Breed or mutate selected individuals to create a new population. Repeat steps 2 through 4.

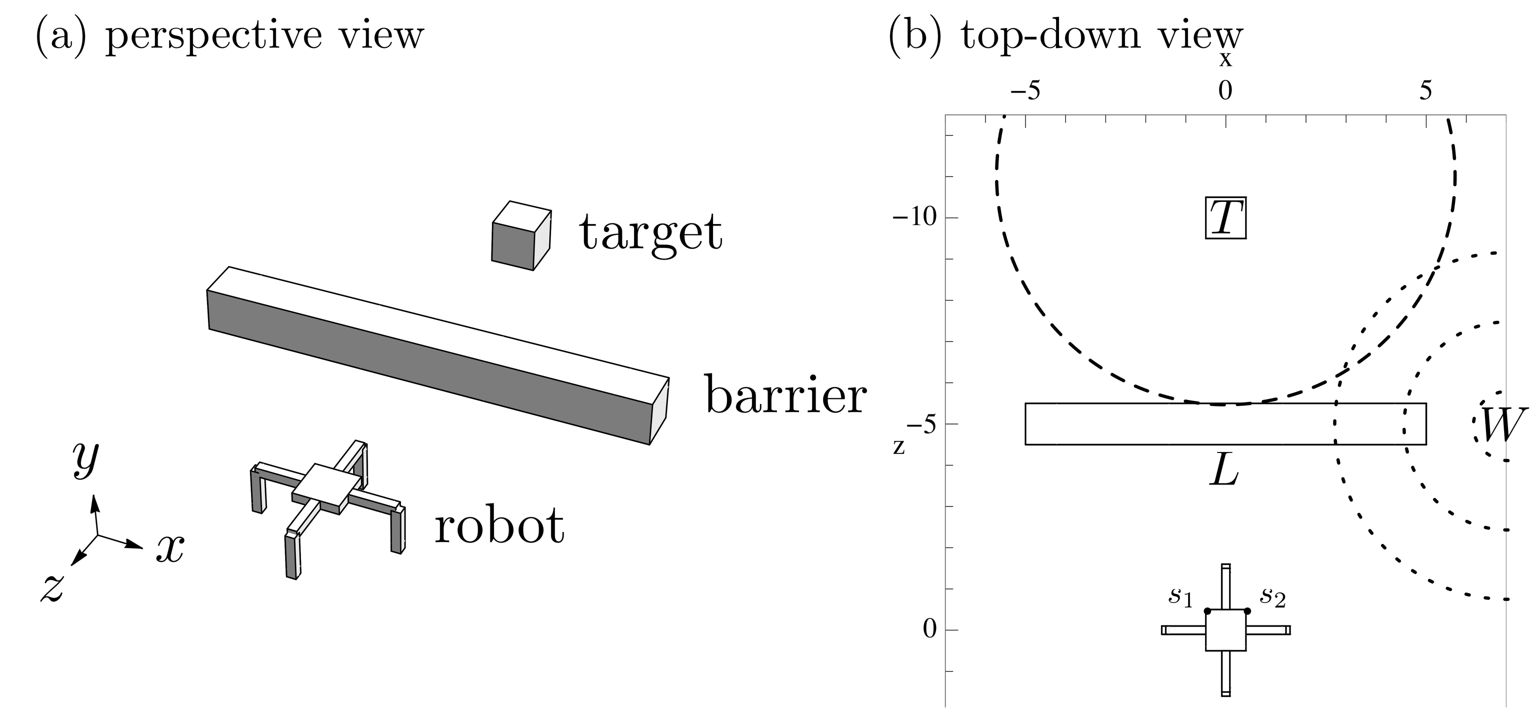

Figure 2: (a) A quadruped robot is tasked with moving to the target. To successfully do this, it must navigate around the barrier. (b) The robot has two line-of-sight target sensors s_1 and s_2. The target T is the global optimum. The robot succeeds if it passes the dashed line bounding T. The local optima for f_{high} is located at L, where it will often become trapped. The mid-level fitness function f_{mid} uses a waypoint W parameterized by α sized radii.

Introduction

The dream of Evolutionary Robotics is to automatically design robot brains and bodies to do any given task. That dream remains unfulfilled, but projects like the protein folding game FoldIt [1] support the idea of not taking humans completely out of the loop. IEAs have traditionally allowed humans to exhibit preferences. In this work, we allow the user to demonstrate how the robot should move to do a task. After all humans have bodies and have intuitions about how to avoid an obstacle, jump over a gap, or ascend a staircase. The proposed hybrid fitness tries to take low-level user demonstrations into account.

Method

The robot used in this investigation is a quadruped walker as shown in Figure 2. The controller is a feed-forward neural network. The robot succeeds if it gets within 5 units of the target. Three fitness functions were evaluated.

1) **High-level fitness** naïvely minimizes the distance between the robot and the target T.
2) **Mid-level fitness** minimizes the distance between the robot and the waypoint W first then minimizes the distance between the robot and the target T. It has an alpha parameter which must be tuned, so three "goldilocks" values are used.
3) **Hybrid fitness** minimizes the distance between the robot and target and concurrently minimizes the user demonstration error. An initial fixed demonstration, a user surrogate, caused the robot to move to the right.

Results and Discussion

Figure 3 shows the results. The high-level fitness function performs the worst, as expected. The mid-level fitness functions perform better than the high-level fitness function for some parameter values. The hybrid fitness does statistically the same or better than all the other fitness functions.

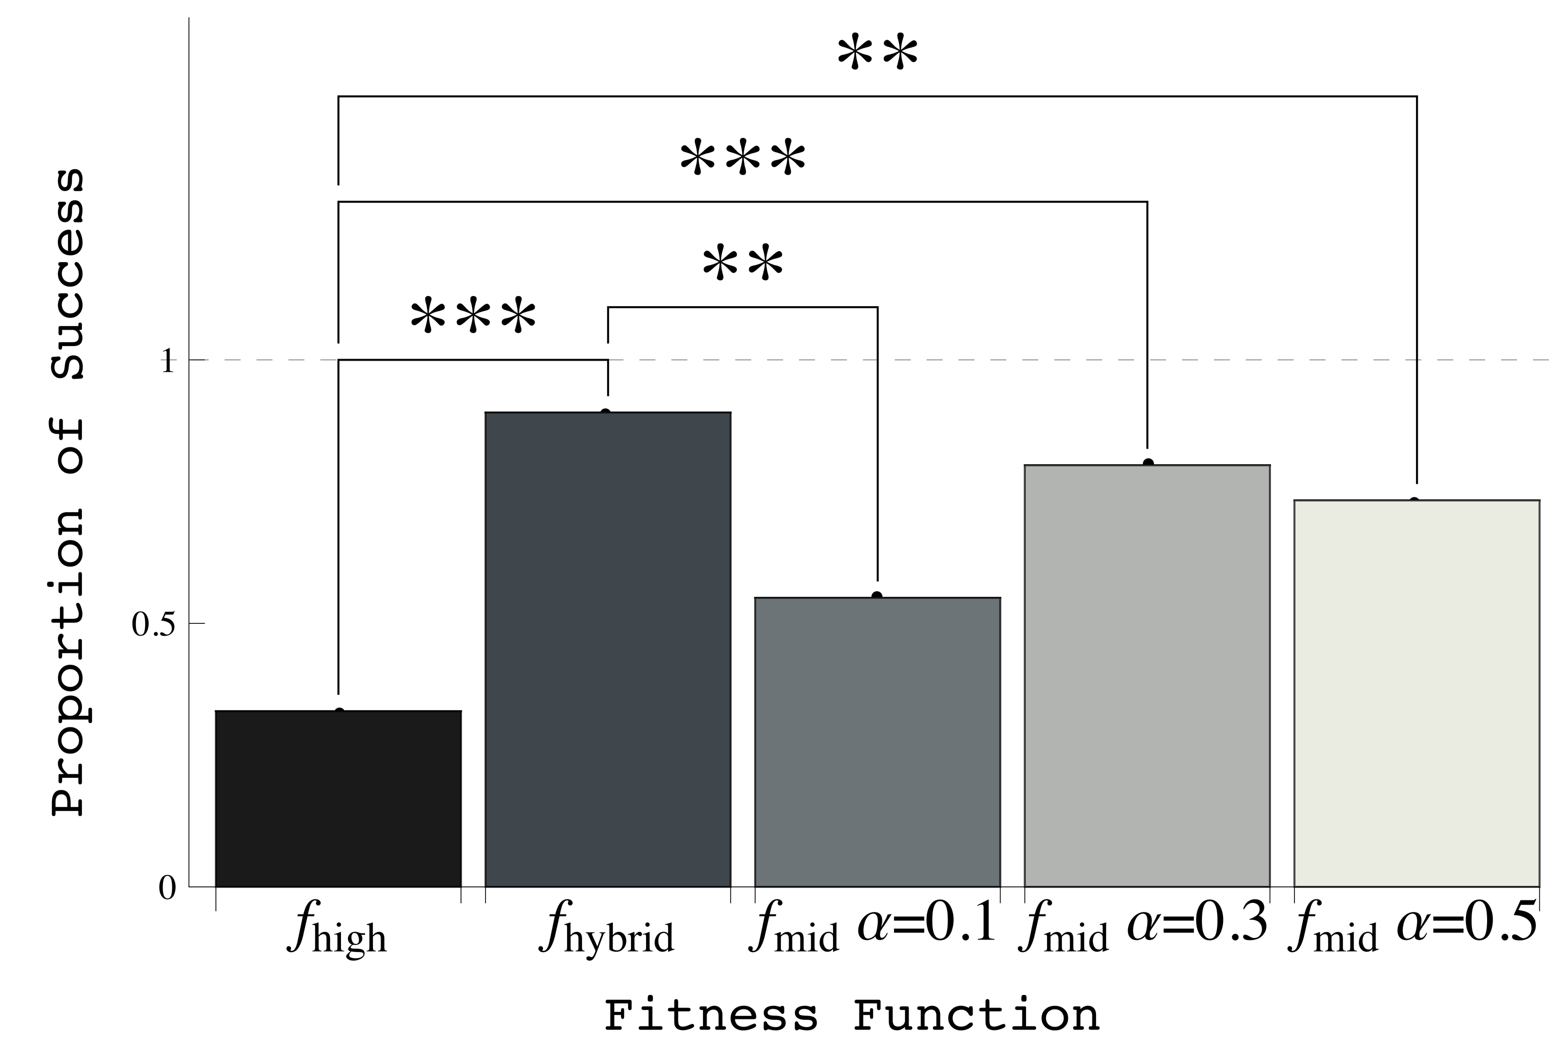

Figure 3: The results show the proportion of successful evolutionary trials that produced a robot that avoided the obstacle. Thirty trials were performed for each fitness function with a population of 20 for 100 generations using NSGA-II. The stars represent significantly different results determined by the Fischer Exact Test.

Future Work

This paper used a user surrogate. However, harnessing human intuition remains the object of this work. A study using this system with human participants on this task and other tasks that do not decompose as easily, like gap jumping, remains to be done.

Selected References

[1] C. B. Eiben, J. B. Siegel, J. B. Bale, S. Cooper, F. Khatib, B. W. Shen, F. Players, B. L. Stoddard, Z. Popovic, and D. Baker. Increased Diels-Alderase activity through backbone remodeling guided by Foldit players. *Nature Biotechnology*, 30(2):190–192, Jan. 2012.

[2] M. Schmidt. Actively probing and modeling users in interactive coevolution. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2006.

[3] J. Bongard. Multiobjective Preference-Based Policy Learning for Evolutionary Robotics. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2013.

[1] shane.celis@uvm.edu, [2] gregory.s.hornby@nasa.gov, [3] josh.bongard@uvm.edu