Crowdseeding Robot Design

Mark D. Wagy
University of Vermont
Burlington, Vermont, USA
mwagy@uvm.edu

Josh C. Bongard
University of Vermont
Burlington, Vermont, USA
jbongard@uvm.edu

ABSTRACT

Crowdsourcing is a well-known method in which intelligence tasks are completed by an anonymous group of human participants. These are tasks that cannot yet be adequately performed by computers. Rather than performing an intelligence task outright, one crowdsourcing strategy is to use human intelligence to complement machine intelligence. A key point in determining the potential of such a strategy is understanding the ways that human abilities most effectively complement the strengths of machine intelligence. We shed light on this relationship by 'crowdseeding' robot design: we find morphological features common to human-generated robot designs and incorporate them as an additional fitness objective in an evolutionary algorithm that searches over the same space of designs. We demonstrate that this approach outperforms the same evolutionary algorithm that is not crowdseeded in this way.

Categories and Subject Descriptors

I.2.9 [**Computing Methodologies**]: Artificial Intelligence-Robotics

Keywords

Artificial Intelligence, Machine Learning, Evolutionary Computation, Crowdsourcing, Crowdseeding

1. INTRODUCTION

Crowdsourcing is a popular method for distributing intelligence tasks to a group of anonymous human participants over the World Wide Web. Recent research on crowdsourcing has focused on how we can leverage the *wisdom of the crowd* [3] to accomplish more complex tasks than those that are simple and easily separable [2].

Here we present a novel crowdsourcing methodology that we term *crowdseeding*. In this two-stage method, 1) participants are asked to participate in a collective activity over the Web, and then 2) features from the first stage are used to seed a machine learning algorithm with the aim of improving its performance, which in our case is an evolutionary algorithm. We do this by incorporating features that were favored by the crowd in stage 1 as an additional objective in the evolutionary algorithm in stage 2. We demonstrate this methodology using web-based interactive robot design.

2. STAGE ONE: CROWDSOURCING

In the first stage of the experiment, we deployed a web-based interface that enabled participants to design robot bodies and vet the quality of their designs using a hillclimber. Participants were anonymous and unpaid. When they visited the study site, they were presented with a design panel (Figure 1C) in which they were able to design robot bodies within their web browser. This panel consisted of a 5×5 grid of dots. When the user clicked on one dot and dragged their mouse to another dot, a line would be drawn between the dots. Only lines between adjacent dots were allowed. When they were finished 'connecting the dots', they could click a *GO* button, which rendered their drawing in a 3D web-embedded physics simulation (Figure 1B). In this simulation, each line in the design was translated to a rectangular parallelipiped segment and each dot that was adjacent to a line was translated into a cube. Segments were attached to cubes with hinge joints that were actuated with a sinusoidal displacement-controlled signal. The sinusoidal signal was fixed at a 1.5 Hz frequency and at a fixed amplitude that caused the joint to sweep an angle of $[-45°, +45°]$. The phase offset of the signal was assigned to be either $0°$ or $180°$. This offset was determined by a hillclimber algorithm unique to each unique robot design. Thus, each time a user clicked *GO* they would be contributing one run of a hillclimber for that particular design.

Users were exposed to a subsample of all of the designs created by users, which were stored in a central repository. Thirteen randomly selected designs created by other users were displayed at the top of the interface (Figure 1A). Each time a user refreshed their page or clicked *GO*, they would be exposed to a new sample of these past robot designs. They were free to use these to guide their own designs, to repeat them and thus contribute one iteration to the hillclimber, or they could create an altogether new design of their own.

3. STAGE TWO: CROWDSEEDING

In the second stage of the experiment, we used prominent features observed in the first stage to construct an additional objective for a multiobjective genetic algorithm. One characteristic that clearly stood out in designs created by

GECCO '15 July 11-15, 2015, Madrid, Spain

ACM ISBN 978-1-4503-3488-4/15/07.
DOI: http://dx.doi.org/10.1145/2739482.2764648

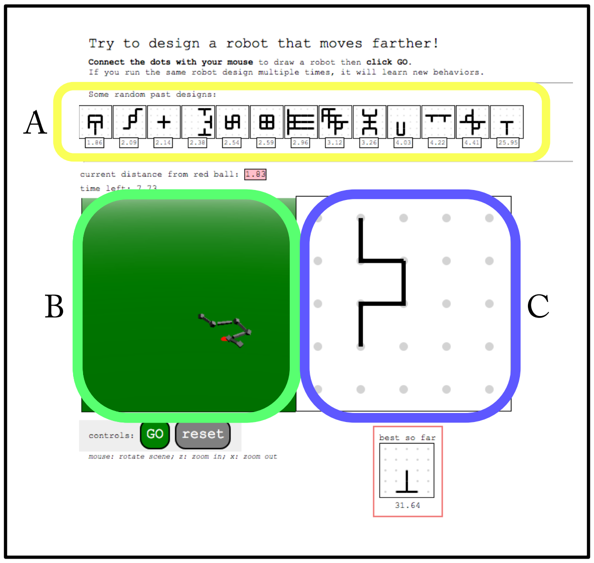

Figure 1: Screenshot of user interface. Users could see a sample of designs produced by other participants at the top (A, yellow box). They could "connect the dots" to draw robot morphologies on the right (C, blue box) and when they clicked on *GO*, they would see a simulation of the robot (B, green box).

the crowd in stage 1 of the experiment was a high degree of reflective symmetry. Approximately seven out of every ten unique designs were perfectly symmetric with respect to reflection over a vertical, horizontal or diagonal axis (and many more were highly symmetric), despite the paucity of symmetric designs in the total number of designs possible using this interface (approximately two per million).

In addition to the primary objective – the distance that a particular robot was able to move from a fixed point – we introduced a secondary objective to maximize symmetry of the robot design to seed a genetic algorithm. The search space of this genetic algorithm was the same as in the first stage. Genotypes in the evolutionary population consisted of a bitstring, composed of bit triplets for each possible line: the first and third bits denoted whether a hinge joint was actuated with a $0°$ (0) or $180°$ (1) phase-offset sinusoid and the middle bit denoted the presence (1) or absence (0) of a segment between each pair of neighboring dots.

4. RESULTS AND DISCUSSION

In the first stage of the experiment, 947 users participated in designing 2292 unique robot bodies. We compared 100 independent trials of a uniobjective control treatment to 100 independent trials of the seeded, bi-objective experimental treatment. The control treatment selected only on the basis of our primary objective: the distance that a particular robot was able to move from an initial, fixed point. The experimental treatment selected for both the primary distance objective as well as the additional, seeded objective, which was to maximize symmetry.

We found that robots created using the experimental treatment methodology were able to move significantly farther than those in the control treatment (Figure 2, $p < 0.05$; Mann-Whitney U-Test) despite the decreased selection pressure resulting from the introduction of an additional objective.

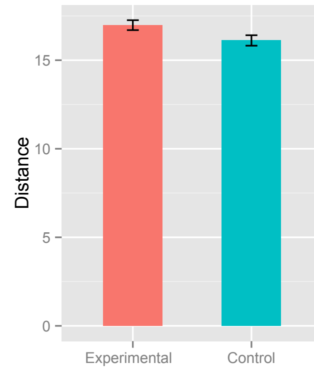

Figure 2: Best distance of control treatment versus experimental treatment at the end of the evolutionary run (Mann-Whitney U-test; p=0.038).

We hypothesize that the crowd was influenced by the predominance of symmetry in locomoting organisms found in Nature; that the situatedness of the human participants in the crowdsourced portion of the experiment contributed to the ability of the evolutionary algorithm in the second stage to find superior robot designs. Thus a group of loosely coordinated non-experts were able to successfully contribute an objective to a machine learning algorithm known through scientific literature to be beneficial to forward locomotion [1] but by no means obvious to that same group in the crowdsourced portion of the experiment. As such, this methodology shows promise as a means for extracting intuitions from the crowd to complement a machine learning algorithm. Future work will focus on developing methods to automate the feature detection stage of this process and to investigate whether there exist less obvious features in the crowd-generated designs which might be complementary to machine intelligence.

5. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) under projects DGE-1144388 and PECASE-0953837, and by the Defense Advanced Research Projects Agency (DARPA) under grant FA8650-11-1-7155.

6. REFERENCES

[1] Josh C Bongard and Chandana Paul. Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. In *From Animals to Animats: The Sixth International Conference on the Simulation of Adaptive Behaviour*. Citeseer, 2000.

[2] Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E Kraut. Crowdforge: Crowdsourcing complex work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 43–52. ACM, 2011.

[3] James Surowiecki, Mark P Silverman, et al. The wisdom of crowds. *American Journal of Physics*, 75(2):190–192, 2007.