Actively Probing and Modeling Users in Interactive Coevolution

Michael D. Schmidt
Computational Synthesis Laboratory
School of Electrical and Computer Engineering
Cornell University, Ithaca NY 14853, USA
mds47@cornell.edu

Hod Lipson
Computational Synthesis Laboratory
School of Mechanical and Aerospace Engineering
Cornell University, Ithaca NY 14853, USA
hod.lipson@cornell.edu

ABSTRACT
A major challenge in interactive evolution is extracting user preferences with minimal probing. We introduce an interactive multi-objective coevolutionary algorithm that actively selects the most informative probes: We simultaneously coevolve a population of candidate models that explain users' selection so far, and a population of candidate probes that cause the most divergence among model predictions, thereby elucidating model uncertainties (divergence). As progress is made, we begin selecting for probes with the highest expected outcome averaged among different models, thereby exploiting model certainties (consensus). In the evolution of pen stroke drawings, we find this technique to be highly effective at extracting preference models from very limited human interaction. Using only pair-wise preference questions, strategy and preference in pen stroke drawings are extracted in fewer than ten user probes. Our results show that the optimal questions to probe the user need not include drawings similar to the target drawing. Instead, the user models converge on trends in the user responses, thereby extrapolating strong preference for target drawings which the models are never actually trained to prefer.

Categories and Subject Descriptors
I.6.5 [**Computing Methodologies**]: Simulation and Modeling – *model development*

General Terms
Algorithms, Performance

Keywords
Interactive Evolution, Fitness Prediction, Estimation-Exploration Algorithm, Active Learning

1. INTRODUCTION
Interactive evolution is a powerful explorative search technique that utilizes human input to make subjective decisions on potential problem solutions [1]. The fitness landscape in each domain is thereby determined explicitly by the human user. Reliance on human input however, induces two major challenges: First, the time cost to collect human input greatly prohibits the discovery of complex solutions. Second, the quality and accuracy of human input greatly degrades with repeated prompts for input.

In this paper, we introduce a coevolutionary algorithm to maximize the information obtained from the user and minimize the necessary interaction. We coevolve a population of individual solutions with a population of models that predict the user's preference. Coevolved solutions are used both to maximize user preference, and also to probe the user in order to refine uncertainty in the user models, two objectives that are not necessarily aligned.

Our primary hypothesis is that intelligently probing the user for input based on their coevolutionary behavior can generate more accurate user models than conventional modeling from very limited user interaction. New user probes must challenge and refine uncertainty and ambiguity in the model population.. We claim that – like the game of *20 questions* – the coevolved individual solutions provide invaluable information to select these new user probes and find optimal user questions base on answers to previous probes [2, 3].

We then compare our results with the interactive costs of a random search and a perfect local search algorithm. The random search comparison shows how effectively the exponential domain can be narrowed through coevolution. The perfect local search comparison, where the user essentially draws their exact preference explicitly, shows how the interactive coevolution of user models and probes algorithm can extract a specific preferred drawing from the user with fewer user probes and much simpler binary preference questions.

2. ALGORITHM SUMMARY
The interactive coevolution of user models and probes algorithm presented in this paper maintains three essential components: the individual population, a graph of user preference training data, and the comparator model ensemble, motivated by an EEA setup [4, 5]. The algorithm operations on these components in five stages: calculating the best comparison pair, requesting the user's input, generating new relations based on the feedback received, training the comparator model ensemble, and evolving the individual population using the comparator model.

The algorithm predicts the most informative test for the user, refines user models, and then evolves individuals. This process continues until the user is satisfied with the top ranked individual coevolved by the comparator ensemble. As long as the user has a

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Genetic and Evolutionary Computation Conference (GECCO) '06, June 25-29, 2006, Washington DC, USA.
Copyright 2006 ACM 1-58113-000-0/00/0004.$5.00.

consistent preference, further iterations will stabilize and simply fine tune the preferred result.

3. EVOLVING DRAWINGS

3.1 Drawing Encoding

In this experiment we evolve drawings produced by a series of closed pen strokes. Each individual encodes each coordinate drawn to on a 32 by 32 pixel image. In our experiments we predefine the number of strokes for each drawing, but this could easily be evolved as well in future work.

3.2 Discovering Star Shape Preference

3.2.1 Star Drawings

In this experiment, the algorithm evolves drawings with six pen strokes. Here we evaluate the algorithm ability to identify a preference for star-shaped drawings. In this experiment we make two basic algorithm comparisons with random search and local search techniques. These comparisons gauge how effectively our algorithm reduces the interactive cost with the user to discover the target drawing.

3.2.2 Results

Figure 1 shows one randomly chosen run to evolve an approximate star shaped drawing. The user has a general strategy when answering comparison prompts to prefer shapes with multiple sharp pointed corners such as a star shape has. The user's preferred drawing for each comparison is shown in green, non-preferred drawings are shown in red, and drawings deemed to be equivalent are shown in dark yellow.

Notice in Figure 1 that a very well formed star is derived as the top predicted shape after five user prompts. No prior prompts required a star shape. Instead the comparator model inferred the star shape as an optimum solution given the user responses favoring shapes with multiple sharp edges.

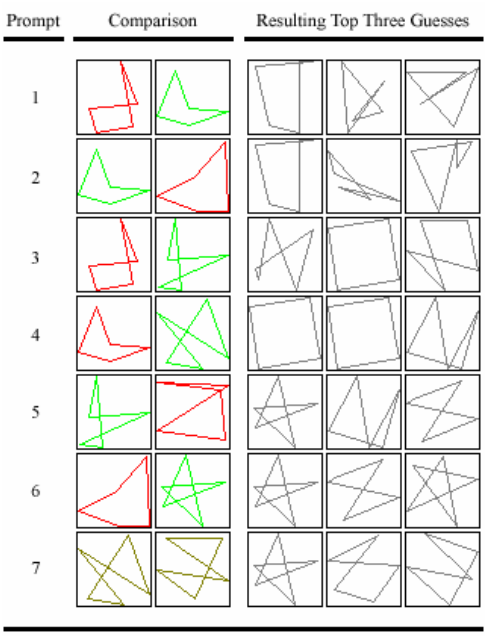


Figure 1. The prompts given to the user and the resulting top three guesses over seven iterations.

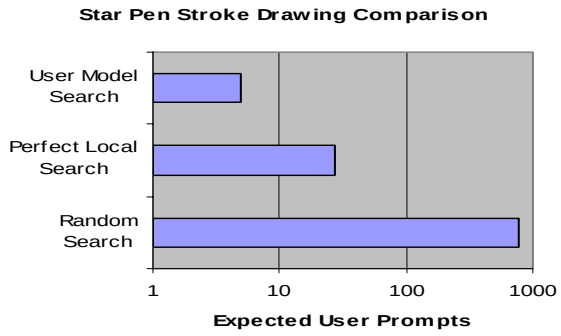


Figure 2. The number of user prompts expected between the compared algorithms to find the target star shape.

Figure 2 shows the comparison with the user comparator model algorithm with the perfect local search and random search algorithms. The logarithmic scale shows that the user comparator model makes significant improvement over the perfectly performing local search and random search.

4. CONCLUSION

Empirical results in this paper show the interactive coevolution of user models and probes algorithm to be very effective at extracting preference models and resulting solutions from very limited human interaction. Using only pair-wise preference questions, strategy and preference in pen stroke drawings are extracted in fewer than ten user probes.

In comparison to the perfect local search algorithm, where the user essentially draws their exact preference explicitly, the interactive coevolution of user models and probes algorithm requires roughly ten times fewer total user input. Furthermore, the prompts to the user are presented in much simpler binary preference questions.

Finally, our results show that the optimal questions to probe the user need not include drawings similar to the target drawing. Instead, the user models converge on trends in the user responses, thereby extrapolating strong preference for target drawings which the models are never actually trained to prefer.

5. REFERENCES

[1] Ecemis, I, E. Bonabeau, and T. Ashburn. "Interactive estimation of agent-based financial markets models: modularity and learning". p.1897-1904, Proc. GECCO 2005.

[2] Hillis, W. "Co-evolving improves simulated evolution as an optimization procedure." In Langton, C. et al. (Eds.), Artificial Life II. Addison Wesley, 1992.

[3] Koza, J. "Genetic Programming: On the Programming of Computers by Means of Natural Selection." Cambridge, MA: The MIT Press, 1992.

[4] Larrañaga, P., and J. Lozano. "Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation." Norwell: Kluwer Academic Publishers, 200.

[5] Zykov, V., J. Bongard, and H. Lipson "Co-evolutionary variance guides physical experimentation in evolutionary system identification." In The 2005 NASA/DoD Conference on Evolvable Hardware, 2005.