Actively Probing and Modeling Users in Interactive Coevolution

Michael Schmidt, Hod Lipson

Computational Synthesis Lab, Cornell University, Ithaca NY, USA

http://ccsl.mae.cornell.edu

Computational Synthesis Lab

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).

Interactive evolution is an explorative search technique that utilizes human input to make subjective decisions on potential problem solutions [1]. User feedback selection can be both time consuming and laborious.

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 to dynamically adapt solution fitness models can generate more accurate preference 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].

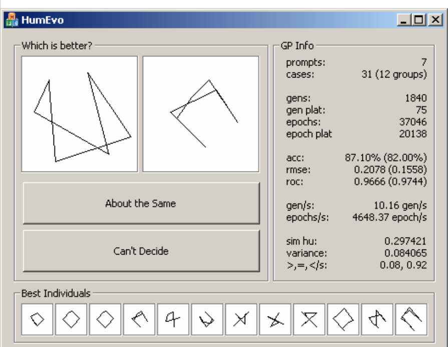

Figure 1. *The user interface and a binary preference prompt.*

Algorithm Overview

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

Empirical Results

Figure 2 who shows the median of 10 experimental runs to discover a user preference for a star-shaped drawing. A star is found after 5 user inputs before a star is even shown to the user.

Figure 3 shows a fitness landscape for a clock drawing. Two maxima are found at the user's preferred times 3:00 and 6:00.

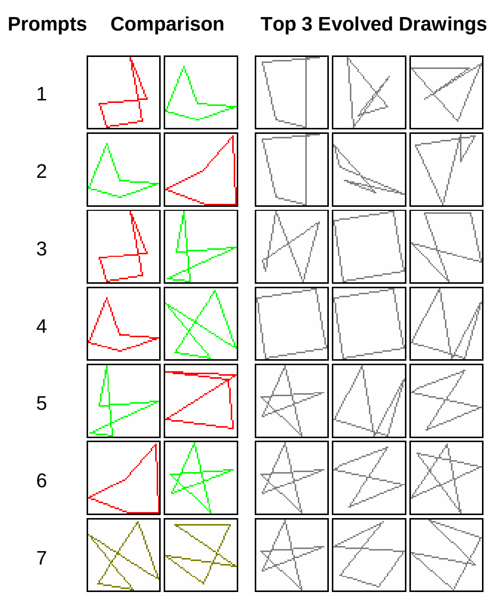

Figure 2. *The prompts and resulting evolved drawings for a user preferring star shaped drawings.*

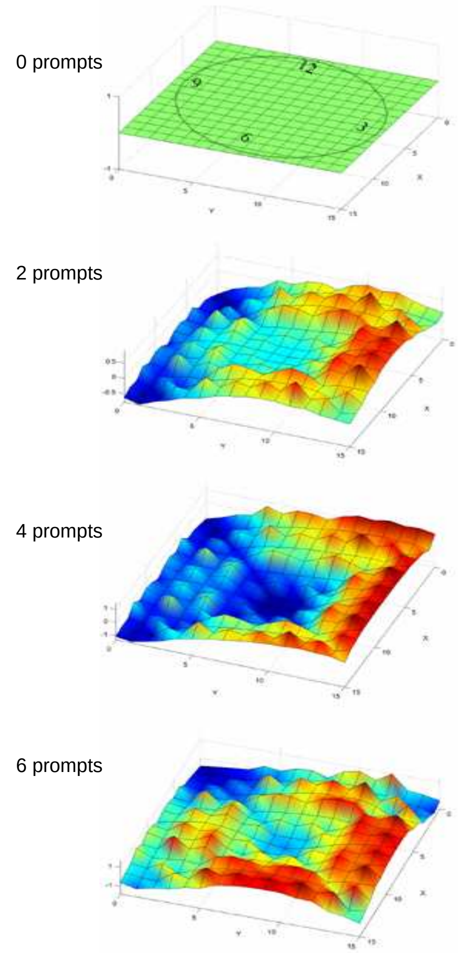

Figure 3. *The effective fitness landscape coevolved after 0, 2, 4, and 6 user prompts for a clock drawing and two preferred times.*

Conclusions

Empirical results show the coevolution of user models to be effective at extracting preference 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.

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.

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.

For more information visit http://ccsl.mae.cornell.edu