Bayesian optimization explains human active search

Ali Borji (1), Laurent Itti (1,2) {borji,itti}@usc.edu; http://ilab.usc.edu/borji

1. Department of Computer Science, University of Southern California (Los Angeles, California)

2. Neuroscience Graduate Program and Department of Psychology University of South California (Los Angeles, California)

Neural Information Processing Systems Foundation 2013

Supported by NSF (CCF-1317433, CMMI-1235539) and ARO (W911NF-11-1-0046, W911NF-12-1-0433)

1. Introduction

Humans utilize sophisticated sequential decision-making strategies for solving real-world problems, but what algorithm do they use? This problem has been addressed in several fields under different names, including active learning [1], Bayesian optimization [2], optimal search [3], optimal experimental design [4], and hyper-parameter optimization.

How do humans compare to algorithms in terms of both performance and behavior? How do humans balance between exploration and exploitation during optimization? Can Gaussian processes [5] offer a unifying theory of human function learning and active search?

2. Exp 1: Optimization I

Subjects click on a blank screen and are shown the ordinate of the function at each clicked abscissa location. Their task is to *find the function's maximum in as few clicks as possible*. Subjects win if they get close enough to the maximum location.

Subjects: 23 non-maths undergrads, Stimuli: 25 1D functions from different families.

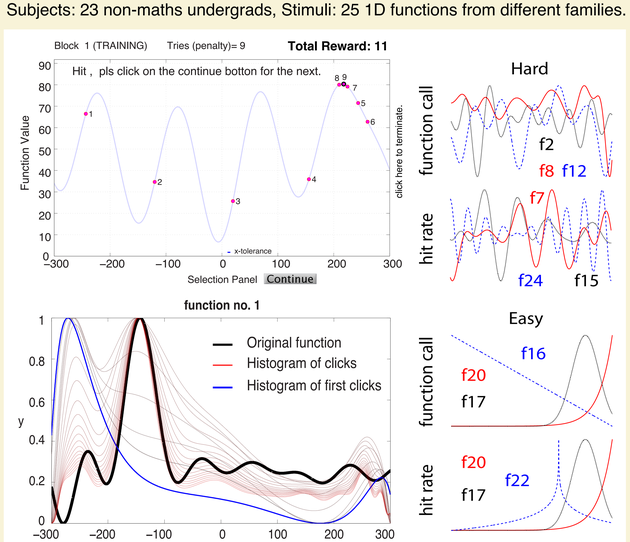

Average number of tries = 12.8 ± 0.4; Average hit rate = 0.74 ± 0.04

3. Humans vs. Optimization Algorithms

In the 1st-level analysis, we perform a grid search over algorithm parameters to sample the hit-rate vs. function-calls plane. Two stopping conditions were considered:

1) we either run an algorithm until a tolerance on x is met (i.e., $|x_{i-1} - x_i| < xTol_m$),
2) we allow it to run up to a variable (maximum) number of function calls.

1st - LEVEL ANALYSIS

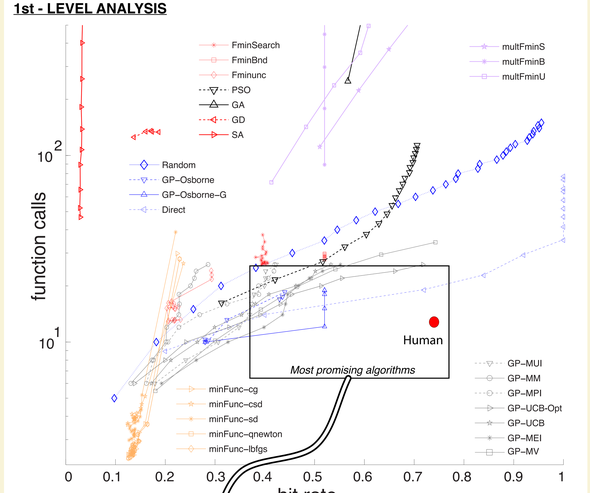

2nd - LEVEL ANALYSIS

Four scores to quantify similarities between algorithms and humans:

1) mean sequence distance between an algorithm's searched locations and human searched locations, in each trial for the first 5 clicks,

2) mean shortest distance between an algorithm's searched locations and all human clicks (i.e., point matching),

3) agreement between probability distributions of searched locations by all humans and an algorithm, and

4) agreement between pdfs of normalized step sizes (to [0 1] on each trial).

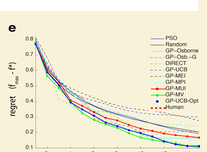

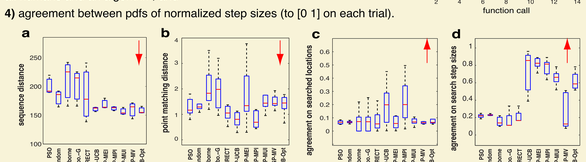

BGO algorithms show higher resemblance to human search behavior

4. Exps 2-6: Interpolation, Extrapolation, and Optimization II

Stimuli: Polynomials of degree 2, 3, and 5. Subjects completed 10 training trials followed by 30 test trials.

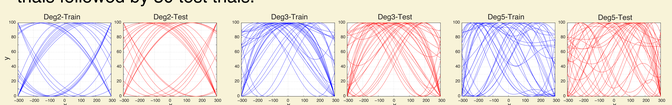

Exp: 2-3. Interpolation: Guess f(0) given 4 points. Active Int.: Choose the best 5th point for guessing f(0).
Exp: 4-5. Extrapolation: Guess f(200) given 3 points. Active Ext.: Choose the best 4th point for guessing f(200).
Exp: 6. Optimization II: Given 4 points, find the maximum location of the function (two chances).

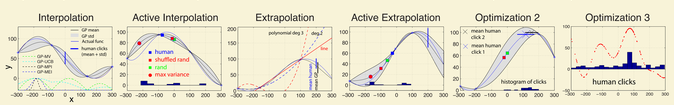

5. Results (Top: Interpolation & Extrapolation, Bottom: Optimization II)

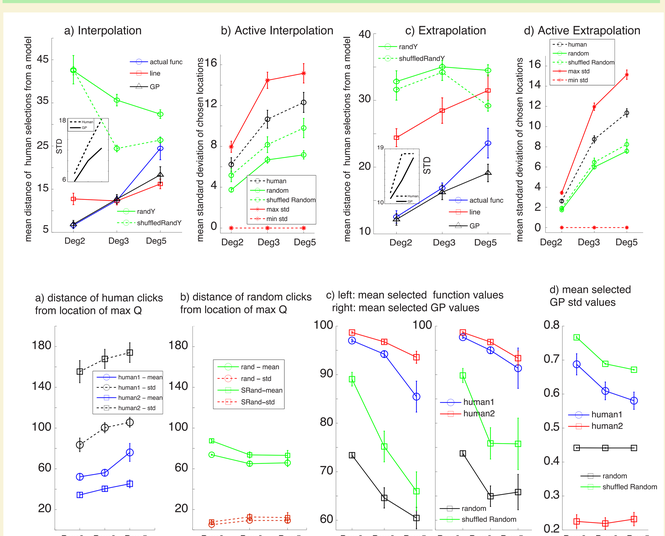

Gaussian processes explain both passive and active function learning in humans

6. Exp 7: Optimization III (Exploration vs. Exploitation)

The task was similar to Optimization I, but here functions were sampled from a Gaussian process with predetermined parameters.

We progressively updated a GP using a subject's clicks on each trial, and exploited this GP to evaluate the next click of the same subject.

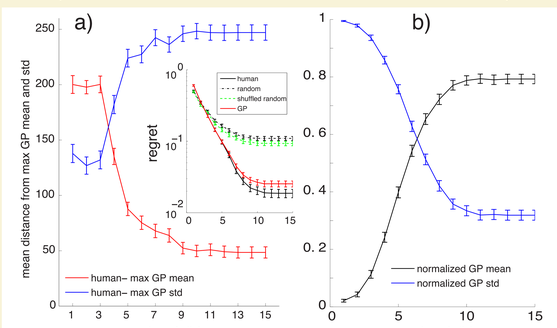

Humans start with exploring the function space and then exploit

7. Discussion and Conclusion

- The high human efficiency in our search tasks does open the challenge that even more efficient optimization algorithms must be possible.
- Gaussian processes provide a reasonable (though not perfect) unifying theoretical account of human function learning, active learning, and search (GP plus a selection strategy).
- Our findings support previous work by Griffiths et al. [6] . Yet, while they showed that Gaussian processes can predict human errors and difficulty in function learning, here we focused on explaining human active behavior with GP, thus extending explanatory power of GP one step ahead.
- Our study shows promising results that may extend to a larger class of natural tasks (e.g., active learning [7], visual search [8]).

References

[1] B. Settles, "Active learning," in Morgan & Claypool, 2012.
[2] E. Brochu, M. Cora, and N. de Freitas, "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," Tech Rep., 2009.
[3] L. D. Stone, "The theory of optimal search," 1975.
[4] V. V. Fedorov, Theory of Optimal Experiments. Academic Press, 1972.
[5] C. E. Rasmussen and C. K. I. Williams, "Gaussian processes for machine learning," in MIT., 2006.
[6] T. L. Griffiths, C. Lucas, J. J. Williams, and M. L. Kalish, "Modeling human function learning with gaussian processes.," in NIPS, 2009.
[7] R. Castro, C. Kalish, R. Nowak, R. Qian, T. Rogers, and X. Zhu, "Human active learning," NIPS, 2008.
[8] J. Najemnik and G. Geisler, "Optimal eye movement strategies in visual search", Nature, 2005.