Influence Functions for Adaptive Stimulus Selection

Carnegie Mellon University

Jayanth Koushik Austin Marcus Aarti Singh Michael J. Tarr

Introduction

How do vision scientists decide which images to present to observers?

Visual stimulus selection is often an *ad hoc* process whereby experimenters hand pick, or manually create a restricted set of images designed to address a given theoretical question. In some studies, presentation order is then randomized, but in others presentation order is manually manipulated so as to identify the image or images that maximally activate a single neuron, a collection of neurons, or the BOLD signal in a brain volume or ROI. But, this simple procedure suffers from many drawbacks:

- Since stimuli are picked manually, experimenter bias is inevitable.
- Responses need to be collected for each stimulus in the set.
- As a consequence of the above, the stimulus set is constrained to be small (and discrete) to allow data collection in a reasonable time.
- This procedure can yield misleading results. For instance, Gross's infamous toilet-brush neuron in IT.[2]

A more efficient approach is to use *adaptive* stimulus selection, where the next stimulus to be presented is determined based on previously collected responses. Using suitable algorithms, a maximally activating stimulus can be identified by exploring only a fraction of the potential input space, allowing larger stimulus sets to be experimentally deployed. This adaptive method is usually employed in a closed-loop system as shown in the figure below.

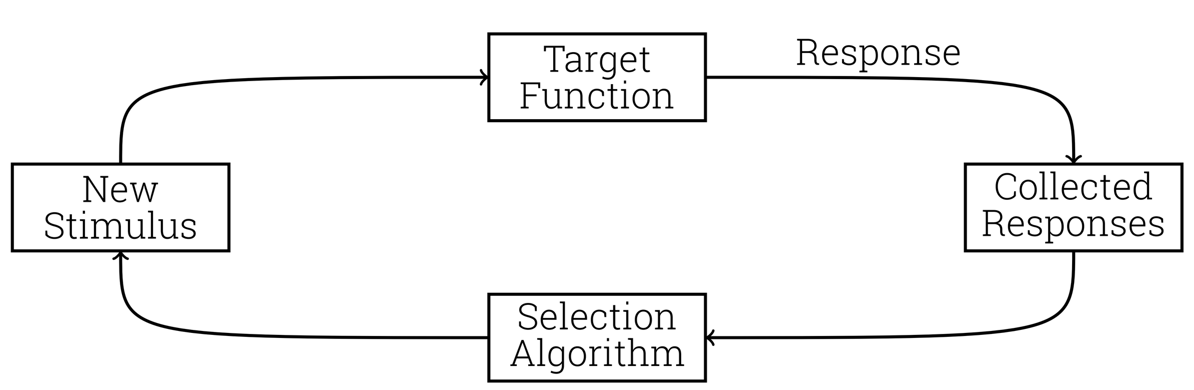

Prior Work

Adaptive stimulus selection has been employed in various contexts. Leeds and Tarr (2016)[3] used fMRI to identify selectivity of neural units and study object representations. Yamane et al. (2008)[5] used an evolutionary strategy to study 3-D object shape in IT. But, a main limitation of previous work is the use of simple optimization methods. To perform more effective search in complex stimulus spaces, we need to use methods that can exploit regularities in complex spaces. In this work, we address this problem, and develop a powerful new optimization algorithm.

Black-Box Optimization

The adaptive stimulus selection problem falls into the general category of black-box optimization. In this context, the function of interest (e.g. brain response) can be interacted with only through evaluation. Mathematically, denoting the function $f : \mathcal{X} \rightarrow \mathbb{R}$, we can get $f(x)$ (with added noise) for any $x \in \mathcal{X}$; but no other information (like gradients), can be obtained. Because of the highly constrained nature, black-box optimization is not amenable to common optimization methods such as gradient descent, and Newton's method. Instead algorithms rely on heuristically searching the input space. The key to efficient search is balancing exploration and exploitation. It is desirable to seek stimuli that are expected to produce a larger response, but being too greedy can lead to insufficient searching. Here we consider a strategy based on confidence bounds to address this trade-off.

Gaussian Process' and Confidence Bounds

We consider a particular black-box optimization algorithm, GP-UCB[4], explained using the figure below, where the solid black line is the function to be optimized. At each iteration, a Gaussian process (GP) is fit to observed data (black dots); this provides a mean fit (dashed blue line), and its variance (blue shaded region). The upper boundary of the variance (blue line) is called the upper confidence bound (UCB), and the next point is selected by maximizing it (blue cross). This leads to a search that selects points with high mean (exploitation), and also those with high variance (exploration).

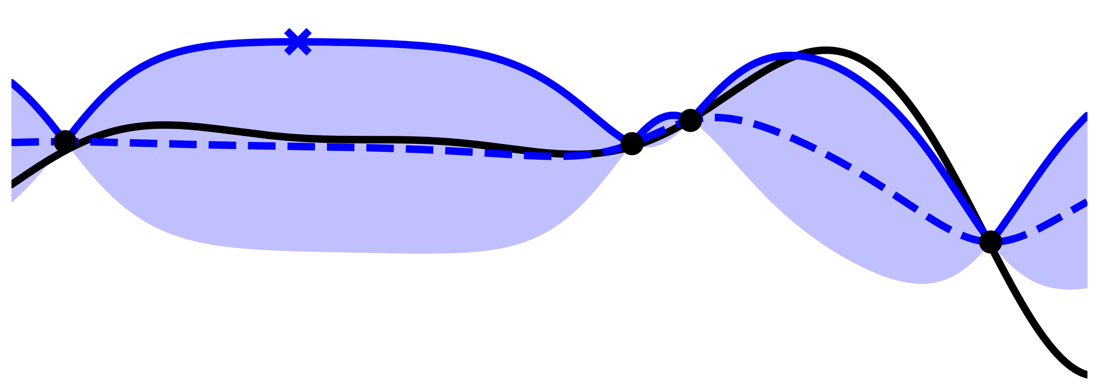

Influence Functions

A limitation of the above method is that it is restricted to use GPs to model data. This can be inefficient in high-dimensional spaces, like visual spaces. Our algorithm uses the statistical concept of influence functions[1] to estimate variance, and obtain confidence bounds for arbitrary models. Then, input selection follows the same procedure as for GP-UCB. The basic idea behind influence functions is shown in the figure below. The black line is the least squares fit obtained using all the data. On removing the green cross, the fit (green line) does not change by much. But on removing the blue cross, the fit (blue line) changes a lot. So there's greater uncertainty about points close to the blue cross, since the prediction is susceptible to larger change. This basic principle can be made rigorous to obtain confidence bands for arbitrary models.

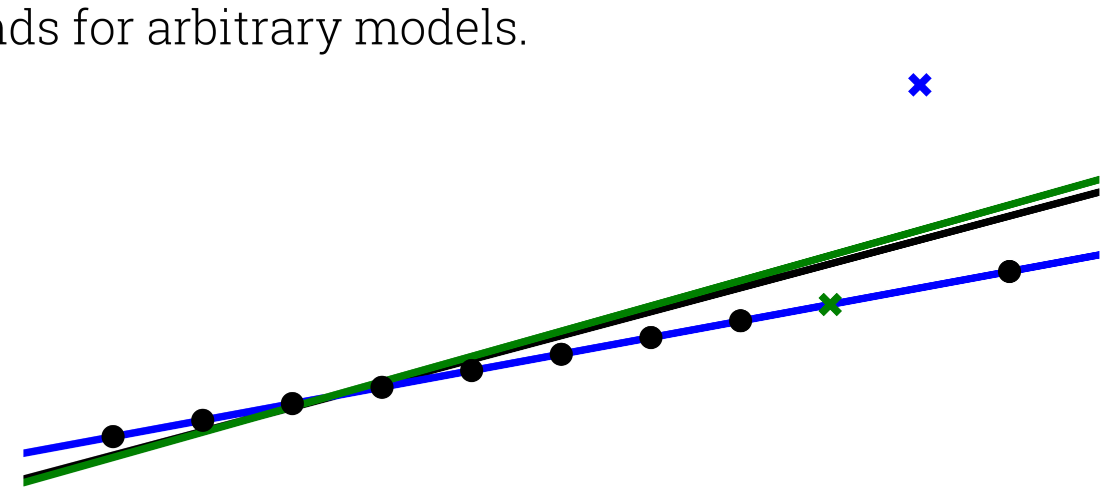

Conclusion

We have developed a new algorithm for black-box optimization, which can be deployed in real-time closed loop systems using a wide class of models. Using a model suitable to the domain can make optimization more efficient by exploiting regularities in the input space. For instance, convolutional neural networks (CNNs) can be used to perform stimulus search in arbitrary image spaces (discrete or continuous).

From a vision science perspective, improved real-time stimulus selection methods will enable more time-efficient and theoretically relevant experimental paradigms. For example, visual neuroscience has often involved searching for the (near-)optimal stimuli driving a given neural unit (a neuron or voxel) - more efficient and reliable search methods will enable a better understanding of the neural representation of visual information.

Acknowledgements

This work is supported by NIH grant MH R90DA023426-12, AFRL grant FA8750-17-2-0212, and CMU's BrainHub ProSEED grant.

References

[1] Cook, R. D. and S. Weisberg
1980. Characterizations of an empirical influence function for detecting influential cases in regression. Technometrics, 22(4):495--508.

[2] Gross, C. G.
2008. Single neuron studies of inferior temporal cortex. Neuropsychologia, 46(3):841--852.

[3] Leeds, D. D. and M. J. Tarr
2016. A method for real-time visual stimulus selection in the study of cortical object perception. Neuroimage, 133:529--548.

[4] Srinivas, N., A. Krause, S. M. Kakade, and M. Seeger
2009. Gaussian process optimization in the bandit setting: No regret and experimental design. arXiv preprint arXiv:0912.3995.

[5] Yamane, Y., E. T. Carlson, K. C. Bowman, Z. Wang, and C. E. Connor
2008. A neural code for three-dimensional object shape in macaque inferotemporal cortex. Nature neuroscience, 11(11):1352.

Contact e-mail: `jayanthkoushik@cmu.edu`