Quantifying Human Perception with Multi-Armed Bandits

Julien Audiffren [1]

[1]Control and Perception Laboratory, University of Fribourg, Switzerland

Psychophysics

A key objective of Psychophysics is to Quantify Human Perception.

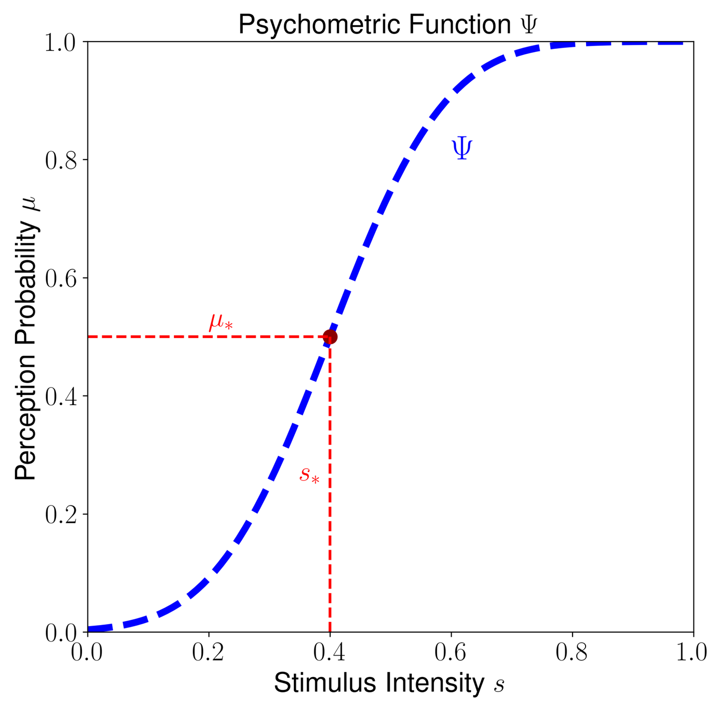

- Human Perception is modeled using a Psychometric Function Ψ
- $\Psi : \mathbb{I} \mapsto [0, 1]$
- $\Psi(s)$ is the probability of the stimulus $s \in \mathbb{I}$ to be perceived
- Here is the space of stimulus is assumed to be $\mathbb{I} = [0, 1]$
- Ψ is known to be non-decreasing, and *somewhat* smooth
- Ψ is **unknown** prior to the experiment

Objective : Quantify human perception.

1. **General Question:** How to estimate Ψ ? [1]
2. **Easier Question:** Given $\mu_* \in [0, 1]$, how to estimate $s_* = \Psi^{-1}(\mu_*)$? [2]

The Psychometric Experiment

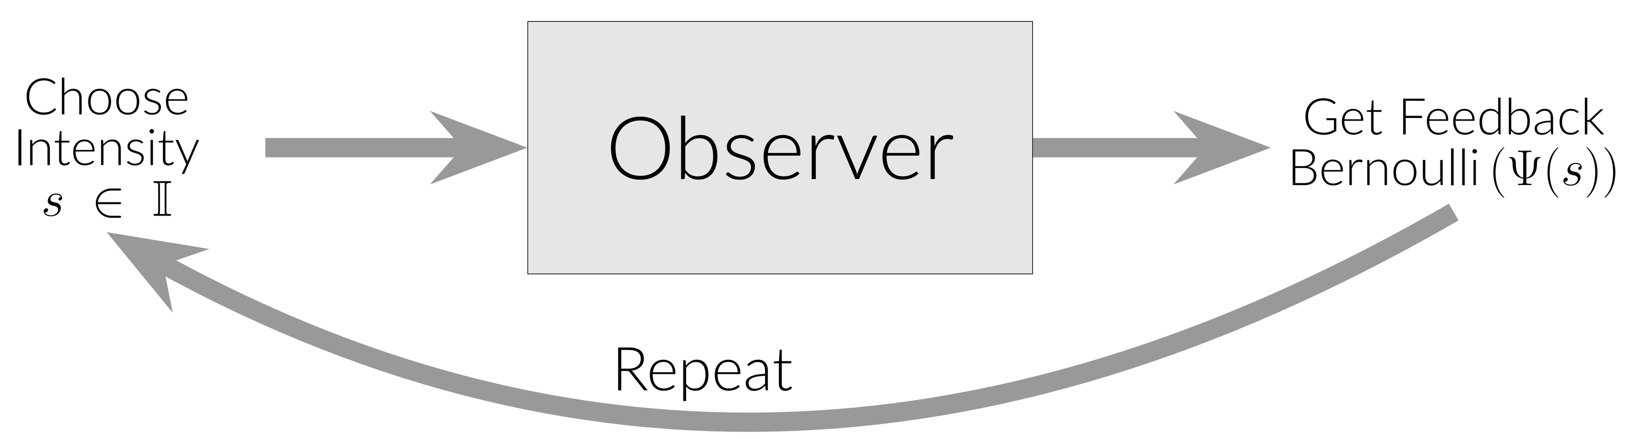

- Each (noisy) evaluation of Ψ Requires the presentation of new stimulus to the observer
- If too many evaluation, problems of *observer fatigue* and *learning*. [1]
- How to solve the easier question as rapidly as possible ?

The Threshold Estimation Problem

Given $\mu_* \in [0, 1]$ (the desired probability) and T (the stimuli budget), find the best possible estimator $\hat{s}$ of $s_* = \Psi^{-1}(\mu_*)$ with at most T evaluation that minimises

$$\mathcal{R}_T(\hat{s}) = |\Psi(\hat{s}) - \mu_*|$$

This is a pure exploration bandit problem !

Dichotomous Optimistic Search (DOS)

Our contribution, DOS, uses a partition tree to estimate s_*, leveraging the non-decreasing property of Ψ.

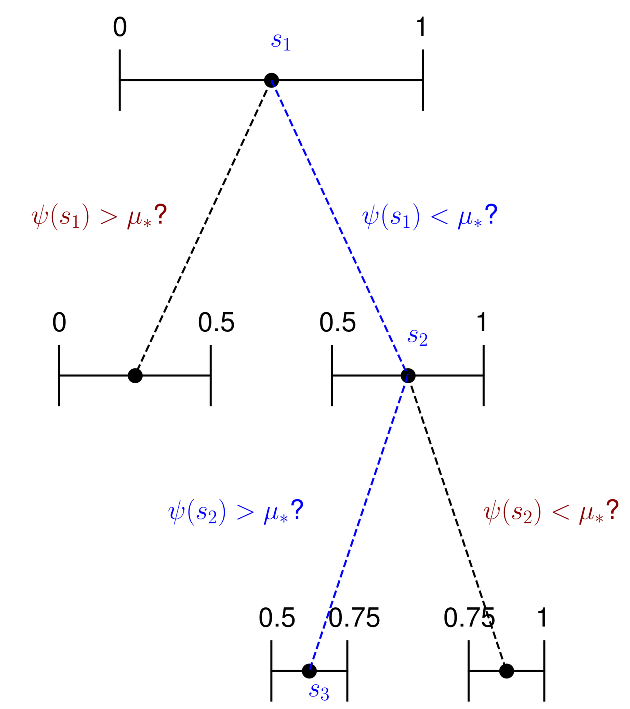

- The Partition tree repeatedly cut the space of stimuli in half
- The agent goes down the tree, choosing the interval that contains x_* according to her estimations.
- Key differences with Binary Search :
  1. No assumptions on the global smoothness of Ψ *(Limited guarantees on the behavior of the partition tree !)*
  2. Only has access to noisy observations *(Presence of uncertainty, Need repeated samples!)*

Challenge: When to move down the partition tree ?

Key Trade-Off: Confidence versus Depth

Confidence. The more s_i is sampled, the more accurate the comparison $s_i > \mu_*$

Depth. The deeper in the tree, the better the estimator $\hat{x}$

DOS Algorithm

<u>Parameters</u> μ_* (objective), T (time horizon)
<u>Initialization</u> $i \leftarrow 1$ (current arm), $s_1 \leftarrow 1/2$ (current stimulus), $N_1 \leftarrow 0$ (number of pulls of s_1), $\hat{\mu}_1 \leftarrow 0$ (empirical average of s_1), $t \leftarrow 0$ (total pulls), $\mathcal{S}$ = null the latest promising arm.
Main Loop
While $t \le T$:
If $|\mu_* - \hat{\mu}_i(t)| > 2\mathcal{B}_T(N_i(t)) \doteq 3\sqrt{\frac{\log(T)}{N_i(t)}}$. or $N_i(t) > N_* \doteq \left\lfloor \frac{T}{(\log T)(\log^2 T)} \right\rfloor$:
If $N_i(t) > N_*$ Then $\mathcal{S} \leftarrow i$ EndIf
Activate new arm: $i \leftarrow i + 1$ and

$$s_i \leftarrow \begin{cases} s_{i-1} + (1/2^i) & \text{if } \mu_* > \hat{\mu}_{i-1} \\ s_{i-1} - (1/2^i) & \text{if } \mu_* \le \hat{\mu}_{i-1} \end{cases}$$

EndIf
Sample arm s_i, update $t, N_i, \hat{\mu}_i$
EndWhile
<u>Output:</u> s_{i_*}, where $i_* = \begin{cases} \mathcal{S} & \text{if } \mathcal{S} \neq \text{null}, \\ \kappa & \text{otherwise.} \end{cases}$

Regret Bounds

If Ψ is Hölder Continuous in a neighbourhood of s_*, then the regret of DOS is upper bounded by

$$\mathbb{E}(\mathcal{R}_T) \le \mathcal{O}\left(\sqrt{\frac{(\log T)^2(\log\log T)}{T}}\right).$$

Experiments

- We empirically evaluated the performance of DOS, and another hierarchical bandit based algorithm POO (Parallel Optimistic Optimization [3]).
- We also used Staircase [4], and PsiMethod [2], two methods from Psychophysics
- We used three psychometric functions : A Gaussian c.d.f. (right), a Beta c.d.f. (center), and an arbitrary Hölder function (left)
- We set $\mu_* = 0.5$, and $T = 200$, and did 100 runs for each experiment.

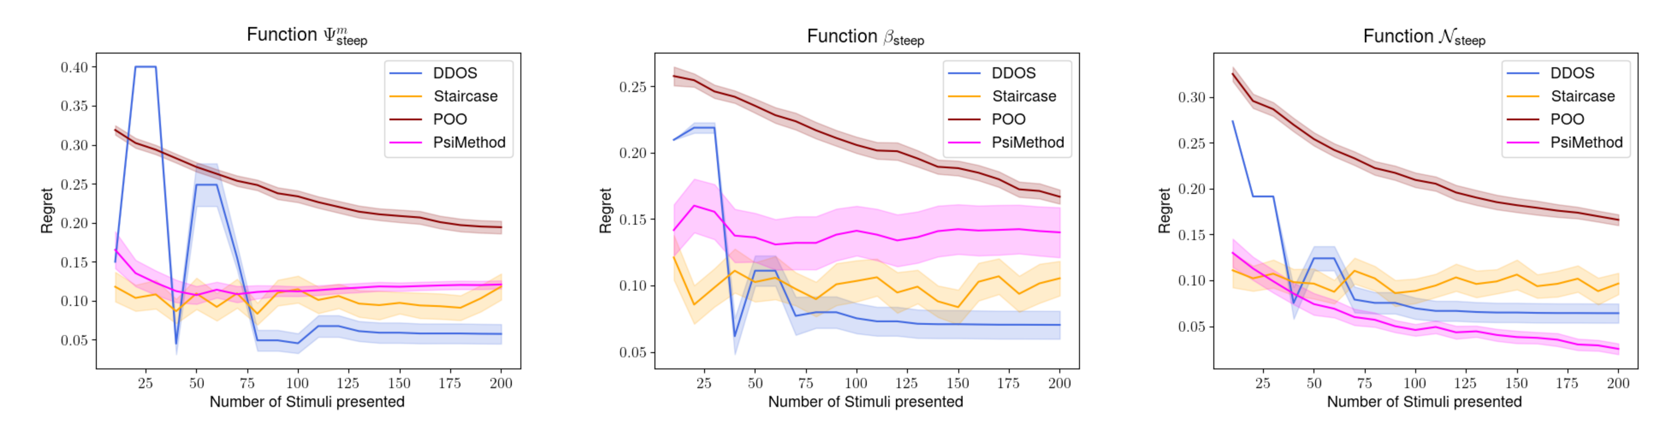

Results

- PsiMethod outperforms other algorithms for Gaussian c.d.f. -- as it is able to leverage its Bayesian assumptions, but performs poorly in other settings
- Second, POO seems to converge, it achieves the worst performance, as its rate of convergence is slow.
- DOS provides one of the best estimation in all these settings.

References

[1] Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. Perception & psychophysics

[2]. Kontsevich, L. L., & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. Vision research

[3] Levitt, H. C. C. H. (1971). Transformed up-down methods in psychoacoustics. The Journal of the Acoustical society of America

[4]. Valko, M., Carpentier, A., & Munos, R. (2013, May). Stochastic simultaneous optimistic optimization. In International Conference on Machine Learning (pp. 19-27). PMLR.

20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), 3-7 May 2021, Virtual