Modeling visual search in naturalistic virtual reality environments

Angela Radulescu*[1,2], Bas van Opheusden*[1,2], Fred Callaway[2], Thomas Griffiths[2], James Hillis[1]

Contact {angelar@princeton.edu}

[1] Facebook Reality Labs, [2] Princeton University

Main contribution:

Frame visual search as a reinforcement learning problem
- Fixations as information-gathering actions
- Do people employ optimal strategies?

Challenges:

- Representing the state space — **what features does visual system have access to?**
- Finding the optimal policy — **how to balance cost of sampling and performance?**

Naturalistic visual search in VR

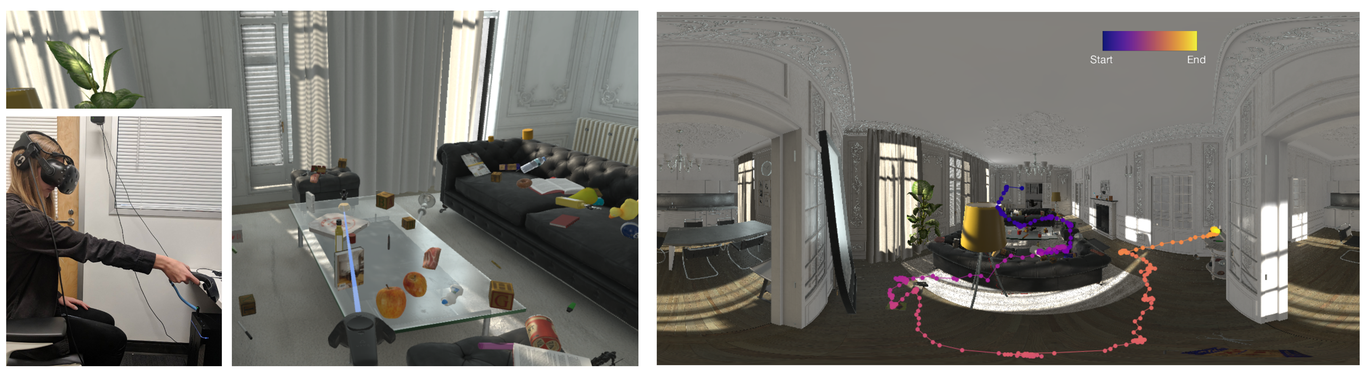

VR + gaze tracking
Fixed camera location; “Find the target within 8 seconds”; 300 unique scenes

Meta-level Markov Decision Process[1]

Latent: $\{F_{true}, i_{true}\}$

- Scene features and target identity unknown to the agent

States: $\{F, J, f_{target}\}$

- Mean and precision of each feature for each object

Actions: $\{o, \perp\}$

- Fixate on object o, or terminate

Transitions: measure $X \sim N(F_{true}, J_{meas})$

- J_{meas} decreases with distance from o
- Integrate X into F and J with Bayesian cue combination

Rewards: if fixating o then $R = -c$; if $\perp$ then $R = 1$ if $argmax(P(target \mid F, J)) = i_{true}$ and 0 otherwise

- Reward agent when most probable target given state matches true target

State representation

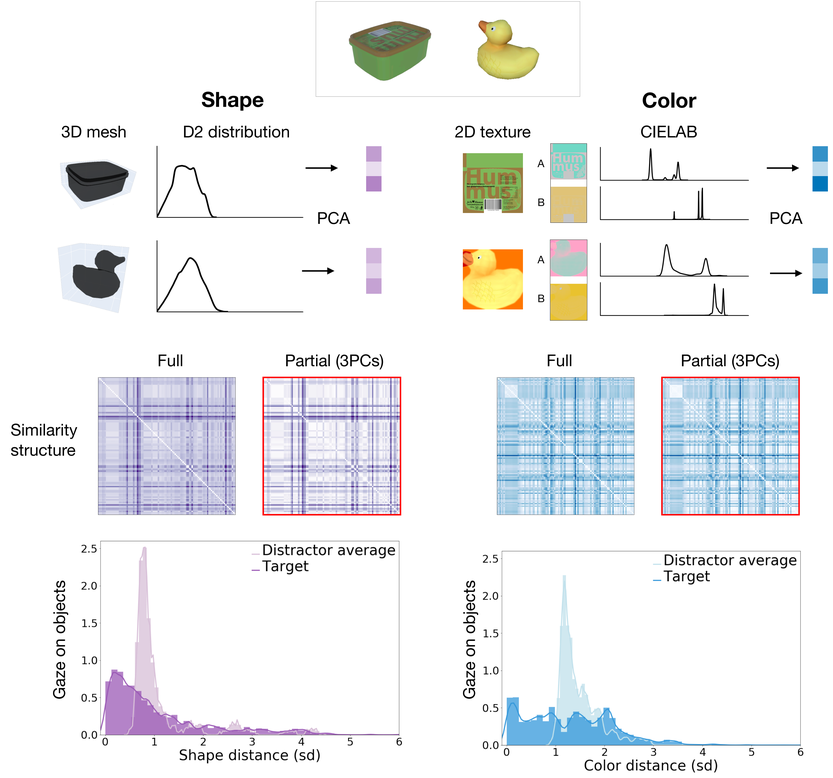

Color and shape features are predictive of gaze.

Is the “ideal searcher”[2] optimal?

- Express “ideal searcher” as policy of Meta-MDP
- Optimize Meta-MDP with deep reinforcement learning

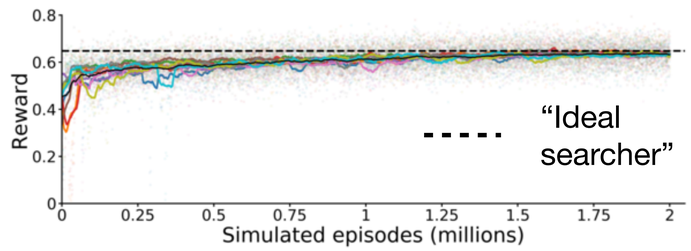

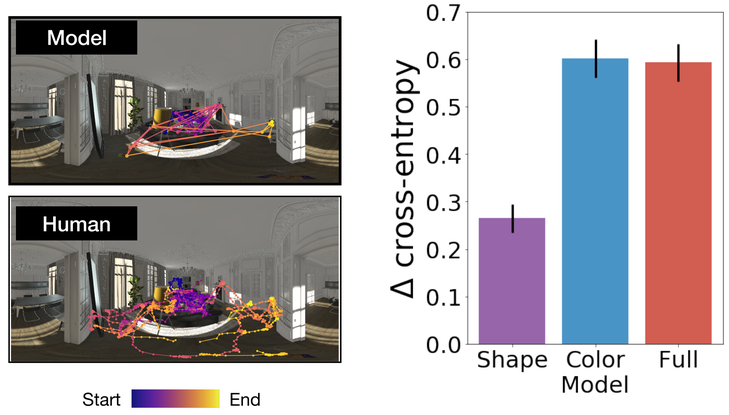

“Ideal searcher” is close to optimal; people are biased searchers.

References [1] Callaway & Griffiths 2018; [2] Najemnik & Geisler, 2005

Open questions: Alternative feature extraction schemes? Is the learned policy truly optimal?