From heuristic to optimal models in naturalistic visual search

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

facebook Reality Labs [1]

PRINCETON UNIVERSITY [2]

Bridging AI and Cognitive Science workshop, *ICLR*
April 24th, 2020

An everyday problem…

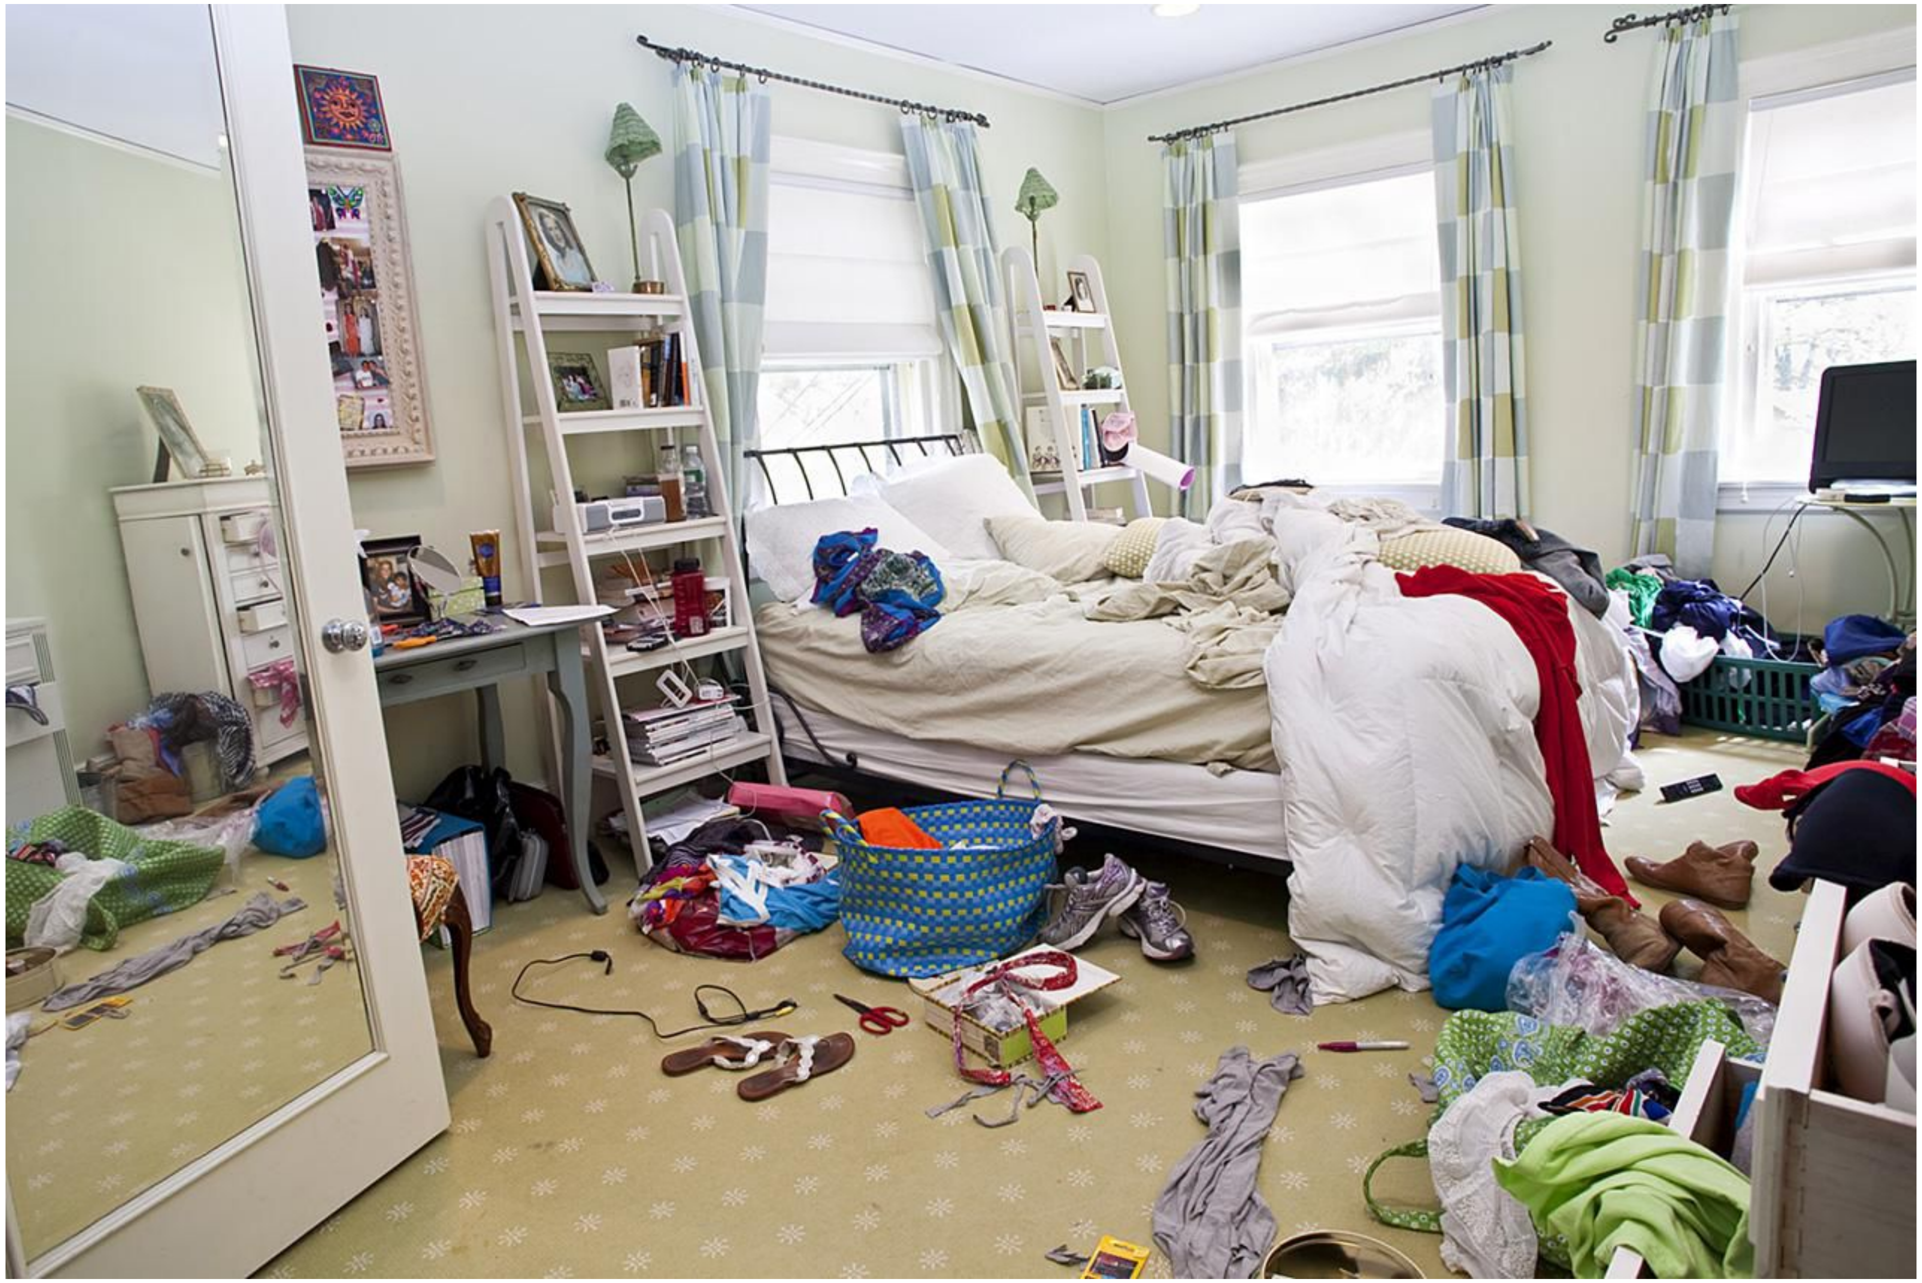

…where are the keys?

Resource allocation in visual search

- **Main contribution**: frame visual search as a **reinforcement learning** problem
    - Fixations as information-gathering actions
    - Do people employ optimal strategies?

Resource allocation in visual search

- **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 — world is high-dimensional; **what features does visual system have access to**?

    - Finding the optimal policy — reward function is sparse; **how to balance cost of sampling and performance**?

Naturalistic visual search in VR

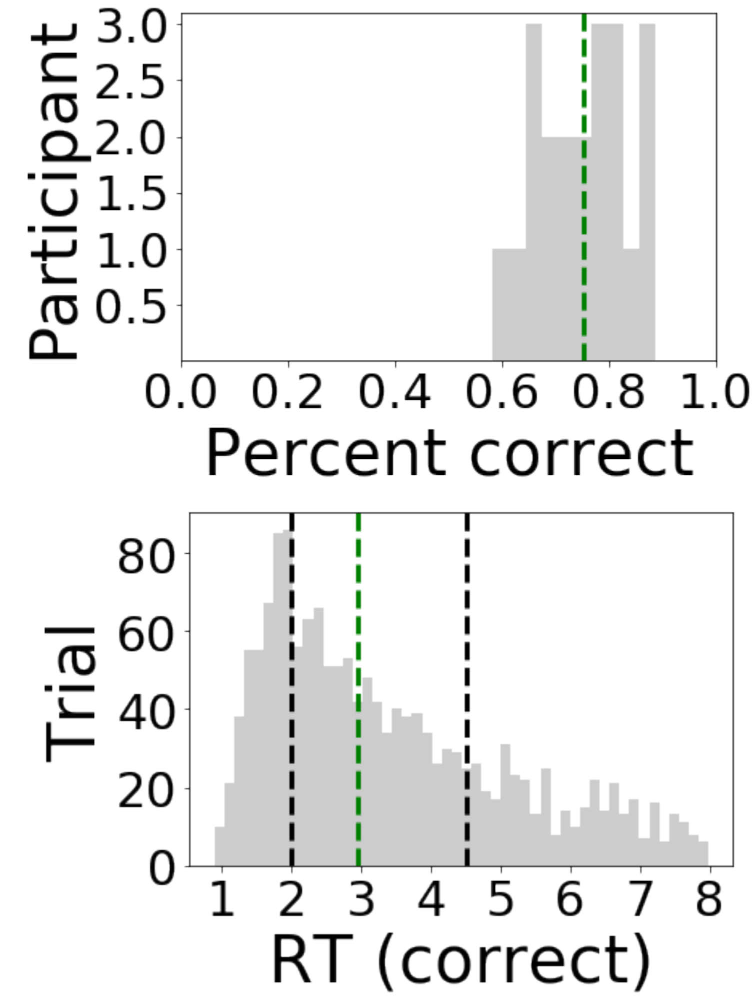

- VR + gaze tracking, fixed camera location
- Cluttered room, 1 target among many distractors
- “Find the target within 8 seconds”
- 6 different rooms x 5 locations per room x 10 trials per location = 300 unique scenes
- Some trials assisted

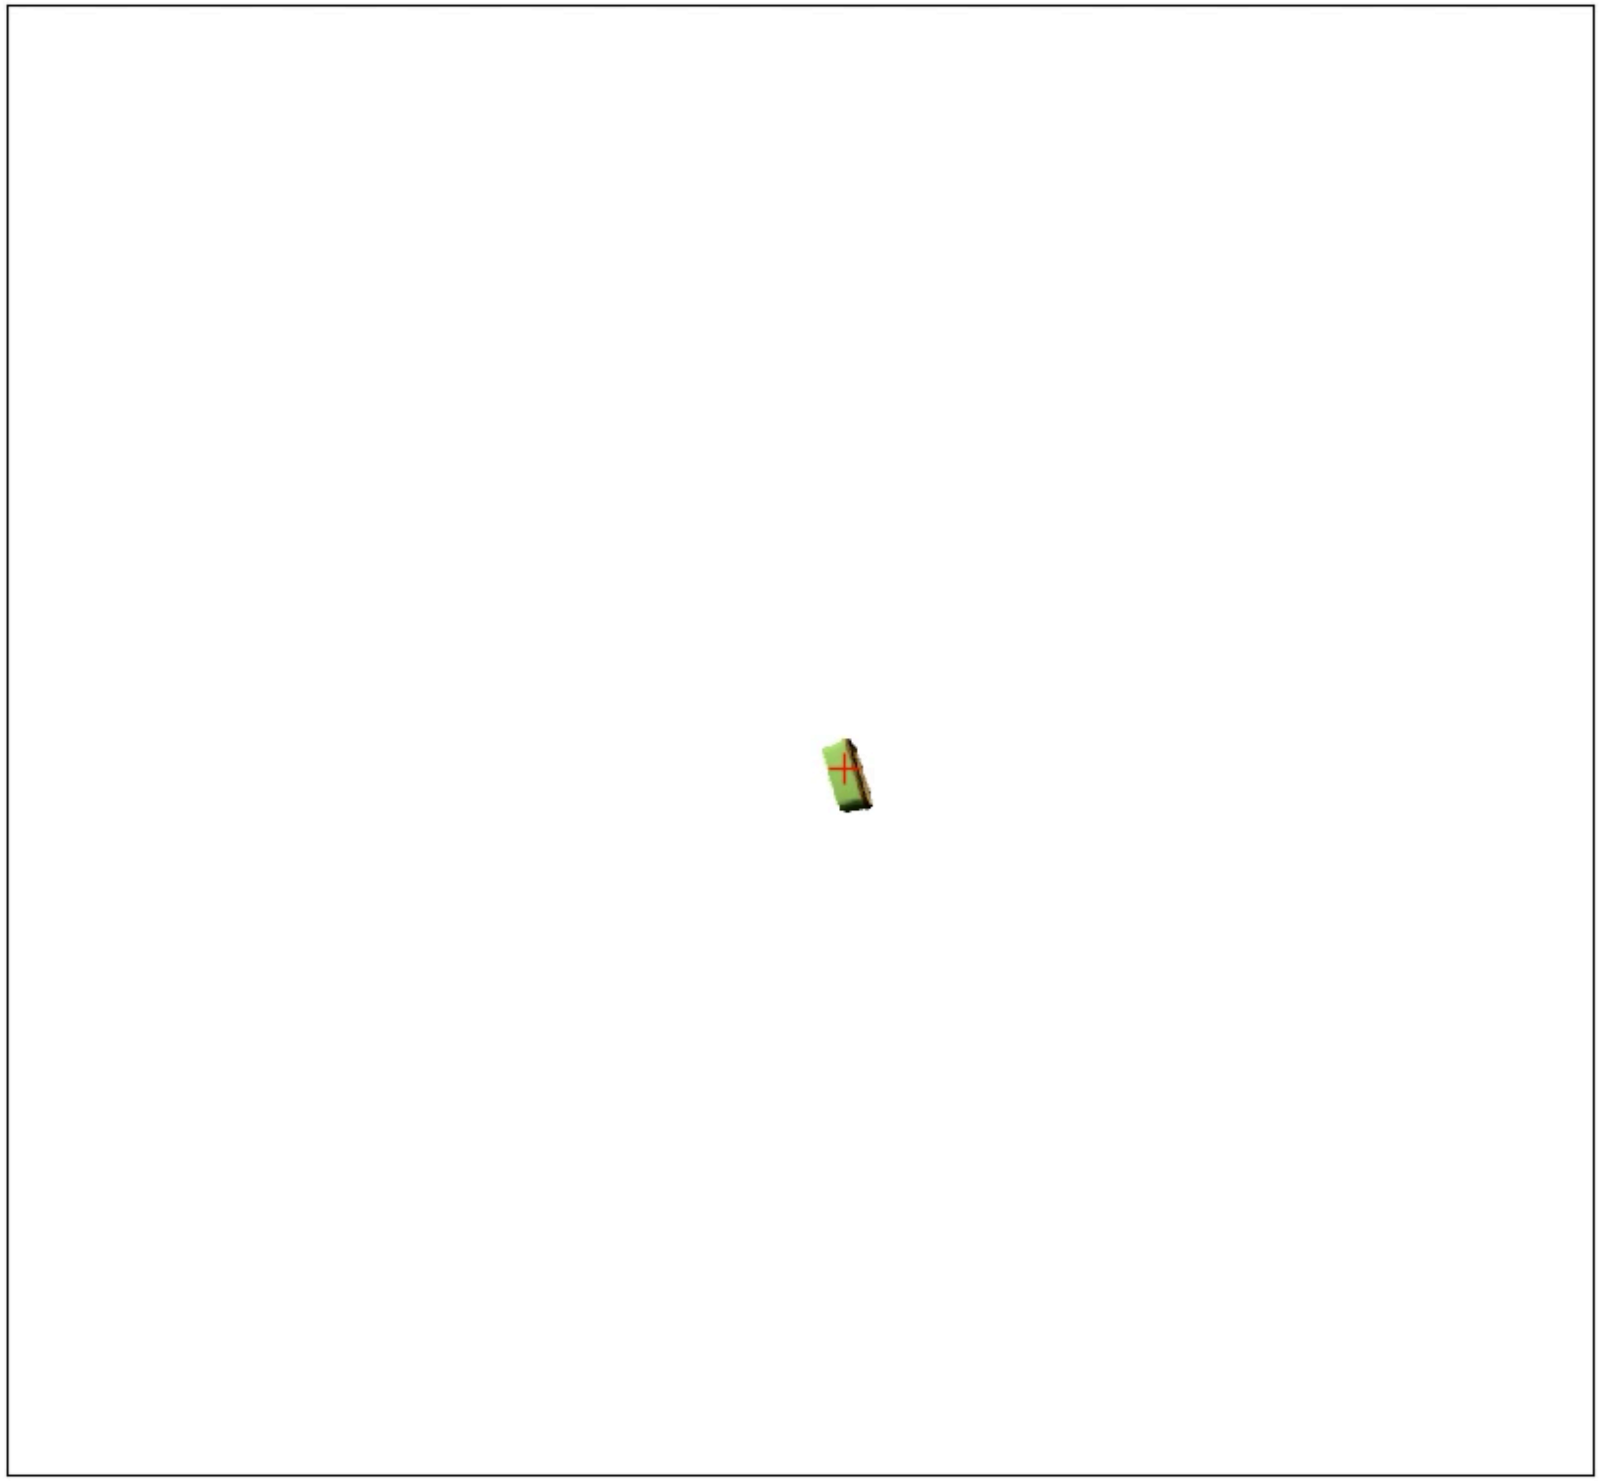

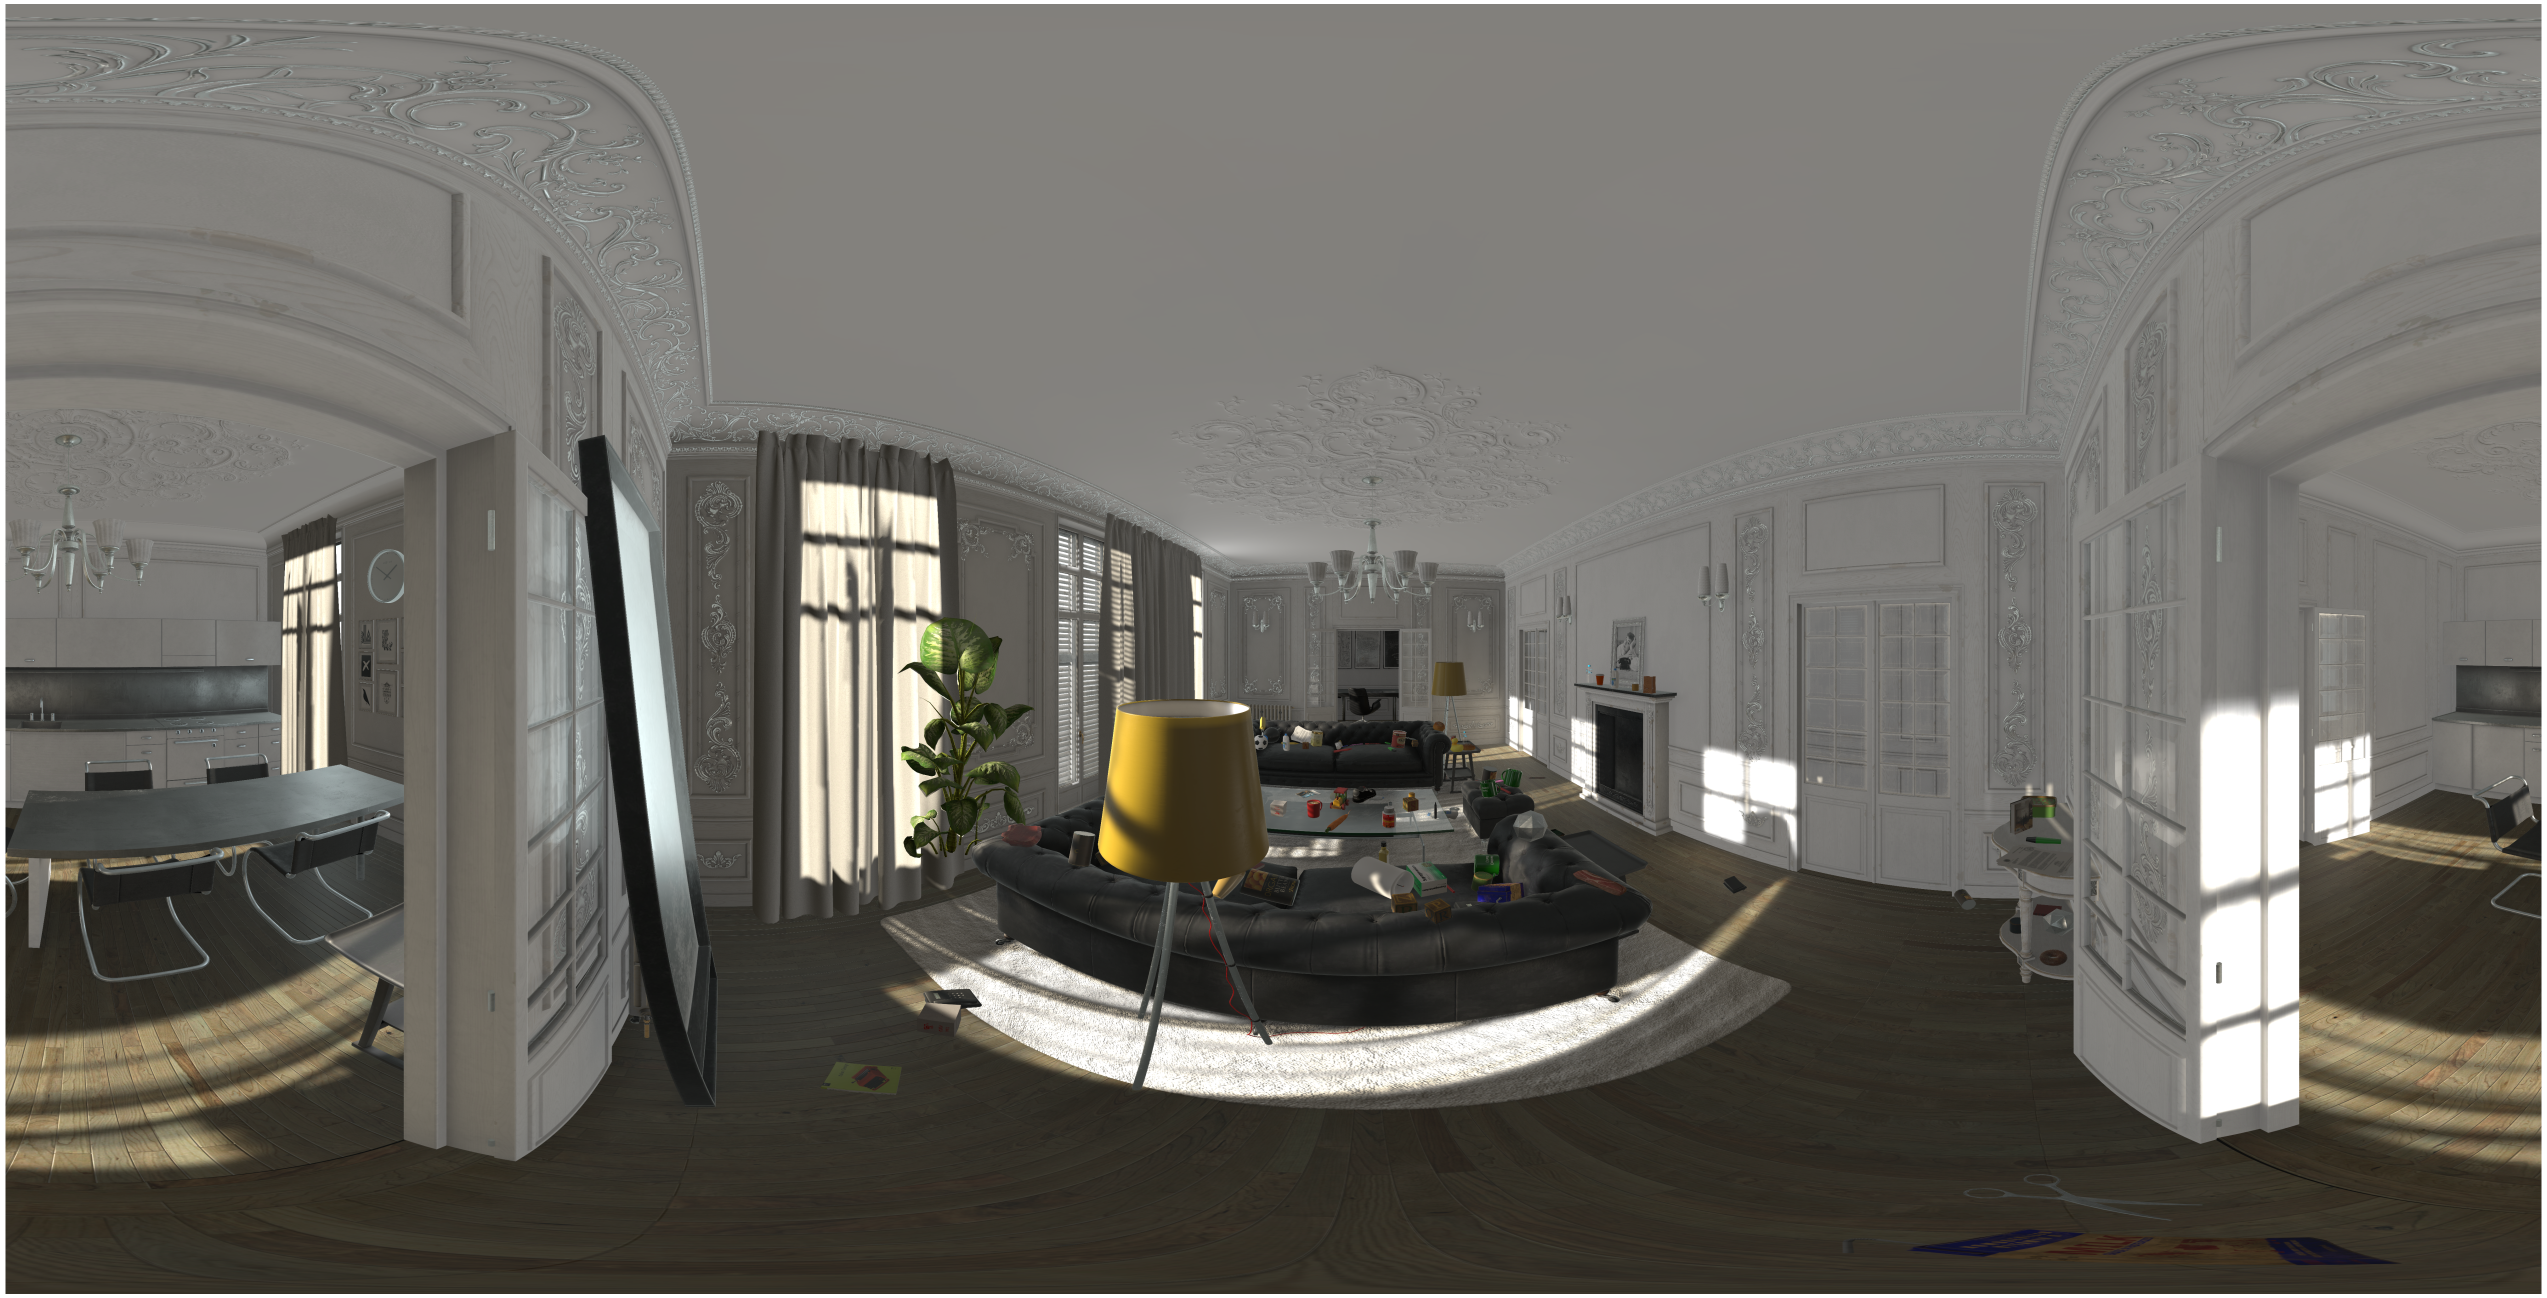


Start
End

Meta-level Markov Decision Process

Callaway & Griffiths 2018

Meta-level Markov Decision Process

- **Latent**: $\{F_{true}, i_{true}\}$
  - Scene features and target identity unknown to the agent

Callaway & Griffiths 2018

Meta-level Markov Decision Process

- **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

Callaway & Griffiths 2018

Meta-level Markov Decision Process

- **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

Callaway & Griffiths 2018

Meta-level Markov Decision Process

- **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

Callaway & Griffiths 2018

Meta-level Markov Decision Process

- **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

Callaway & Griffiths 2018

Challenge I: representing the belief space

Challenge II: finding the optimal policy

Which features to include?

Which features to include?

Objects

Target

Shape

Color

Object attributes

Treisman & Gelade, 1980
Horowitz & Wolfe, 2017

Which features to include?

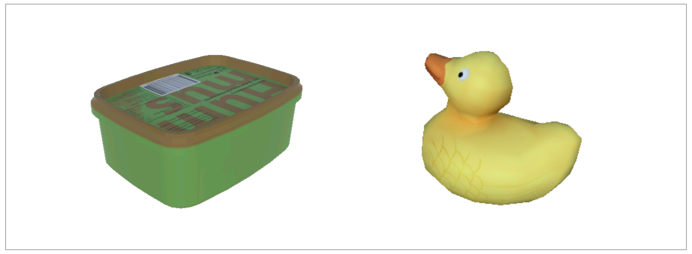

Which features to include?

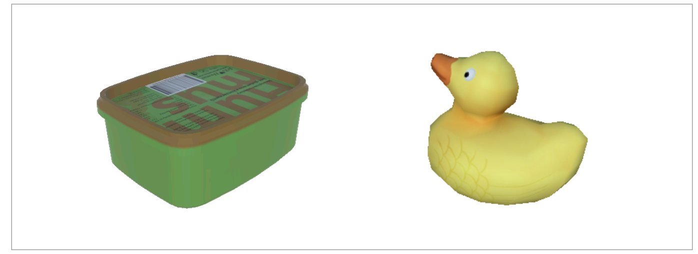

Shape

3D mesh | D2 distribution

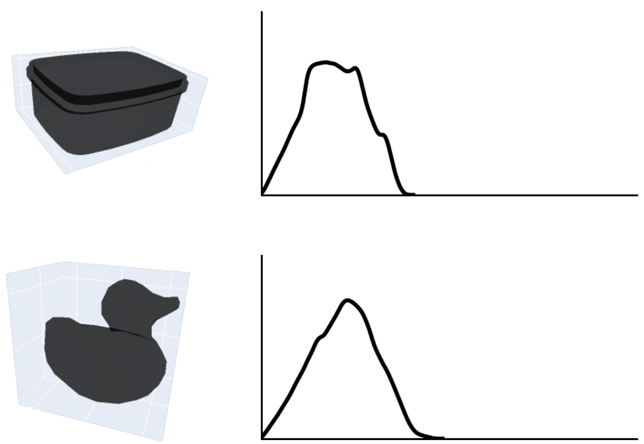

Which features to include?

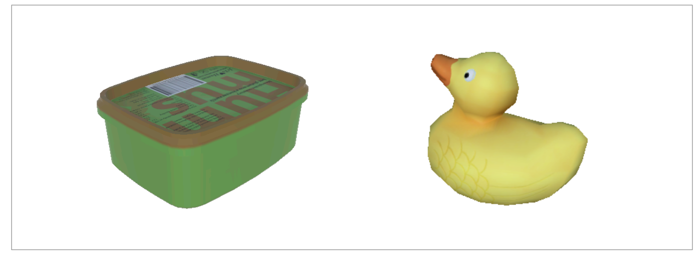

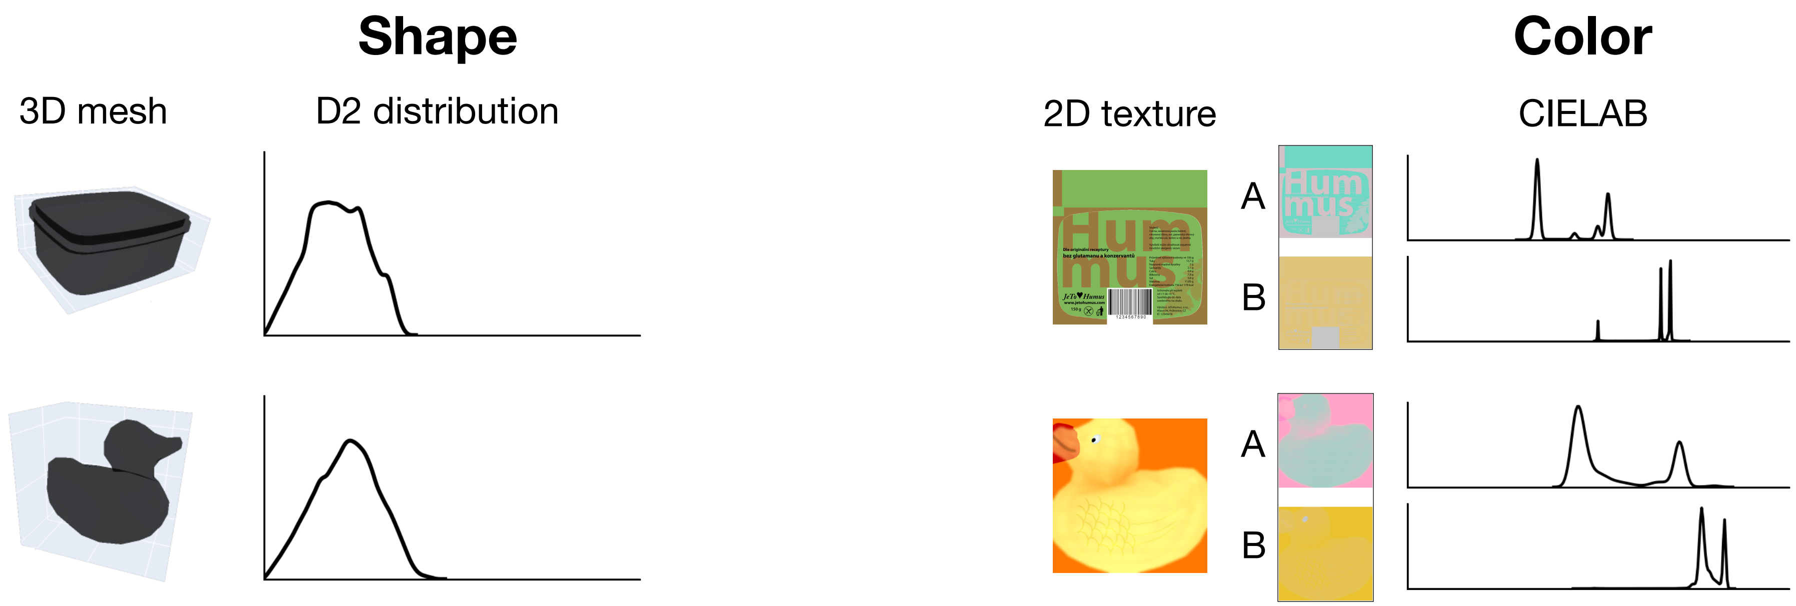

Which features to include?

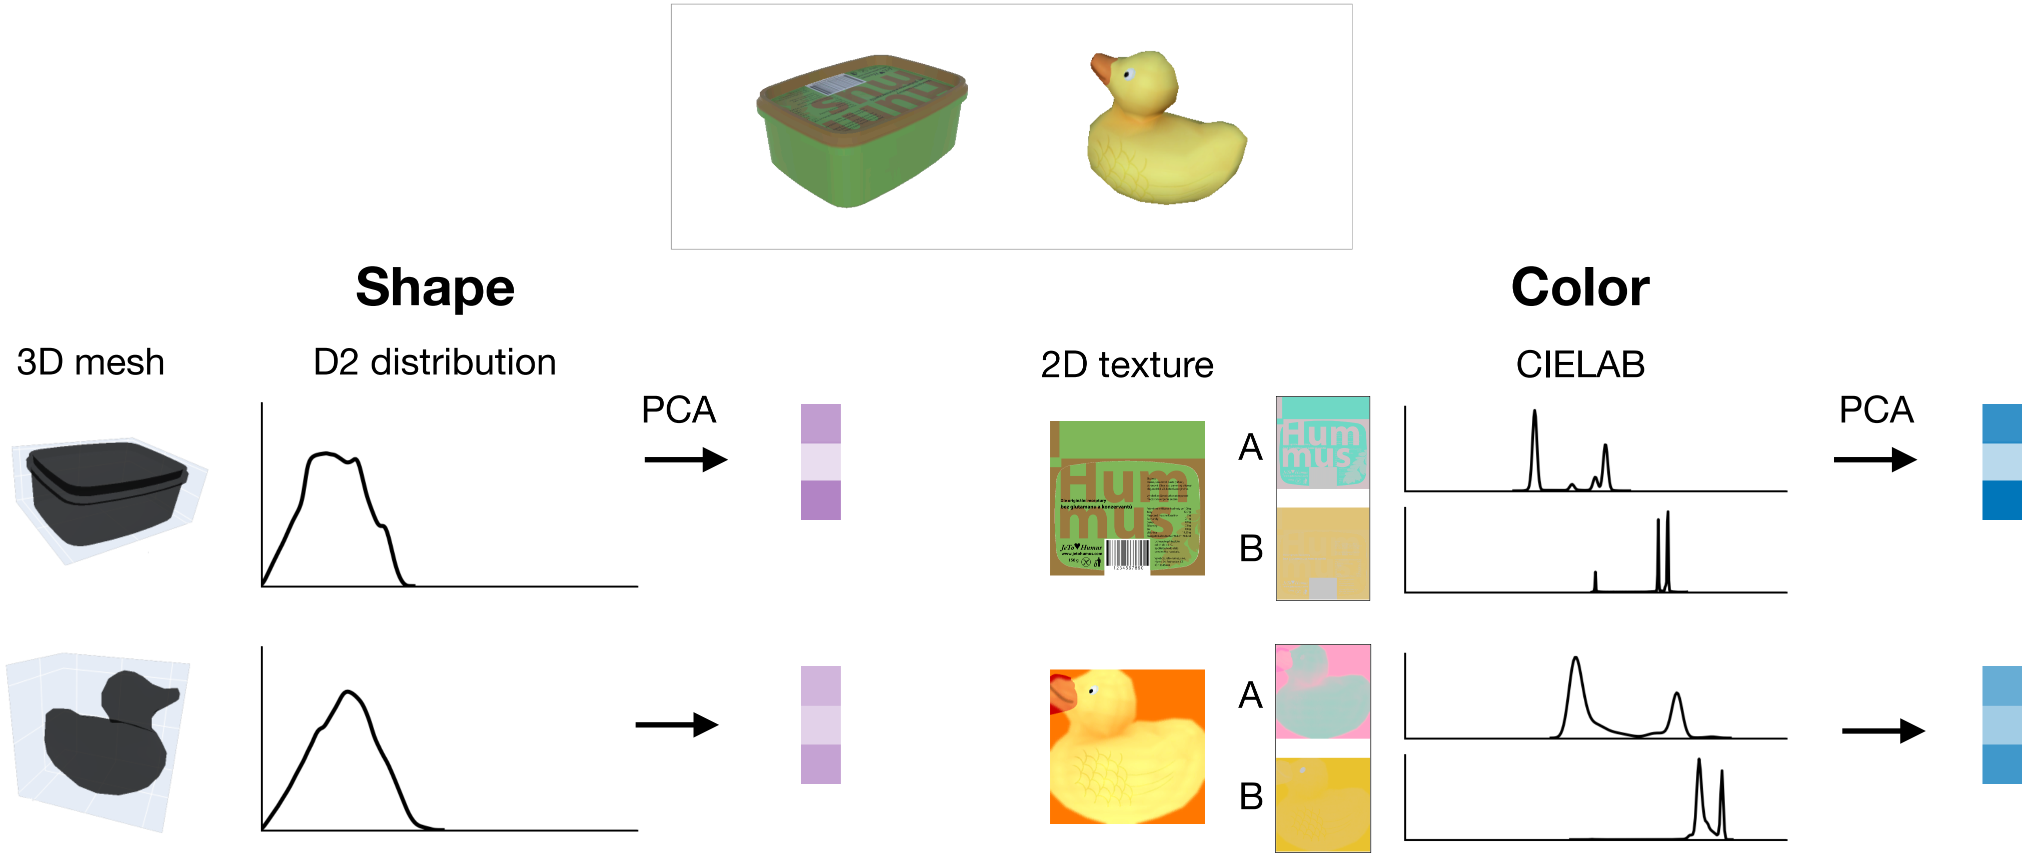

Which features to include?

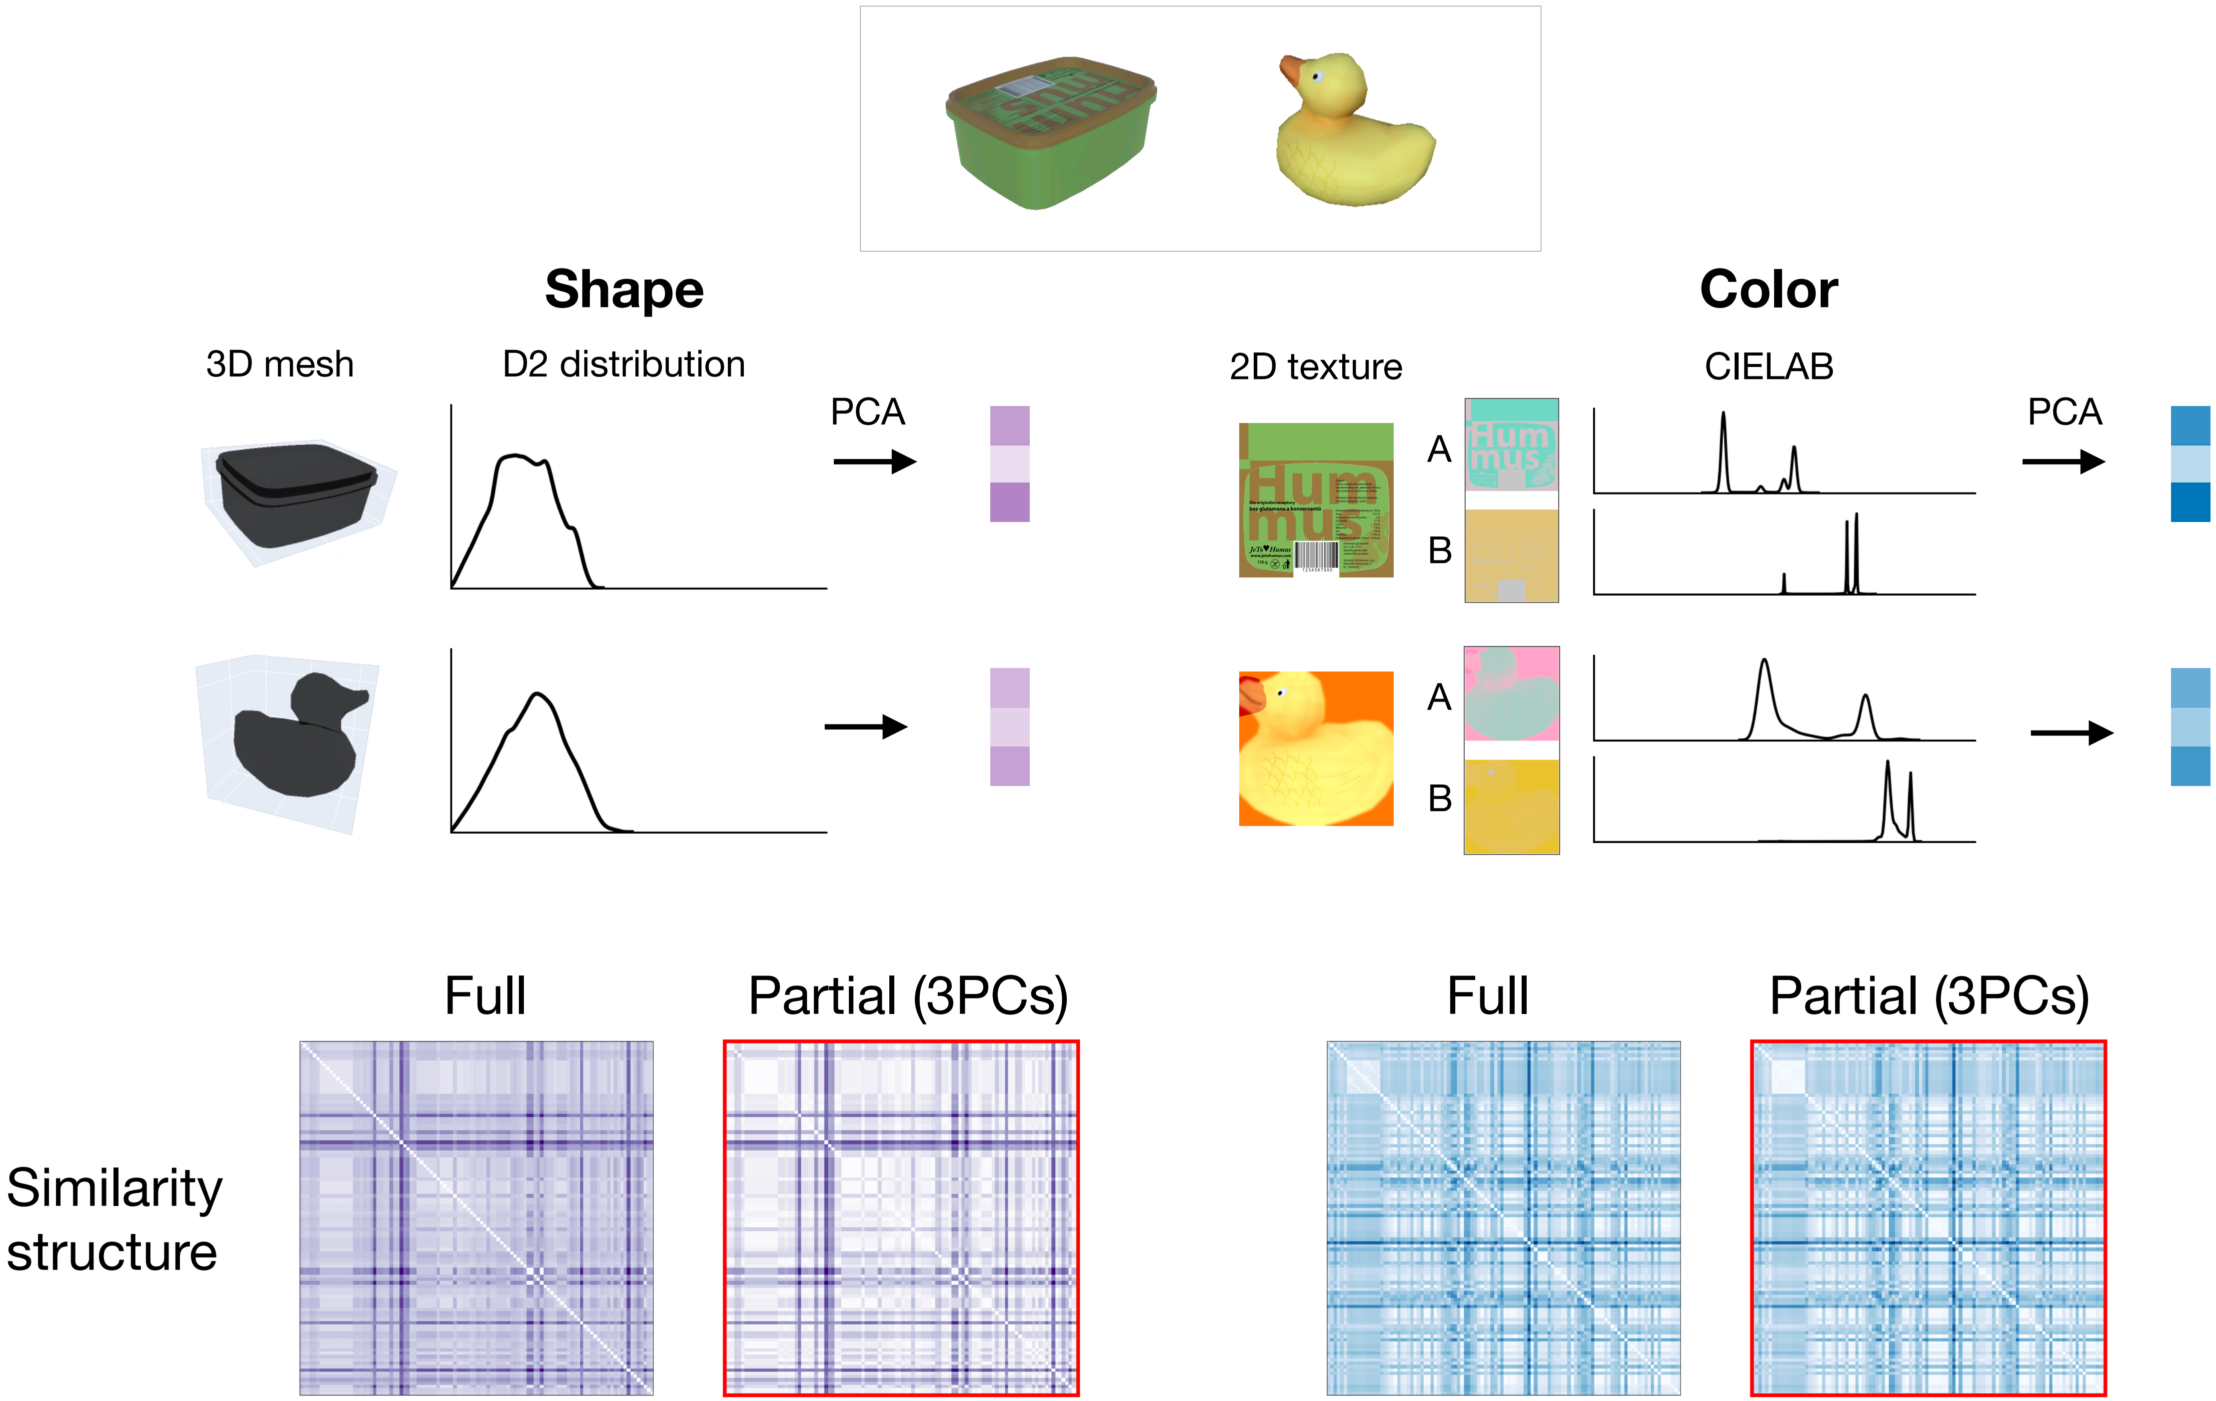

Shape and color predict gaze

Shape and color predict gaze

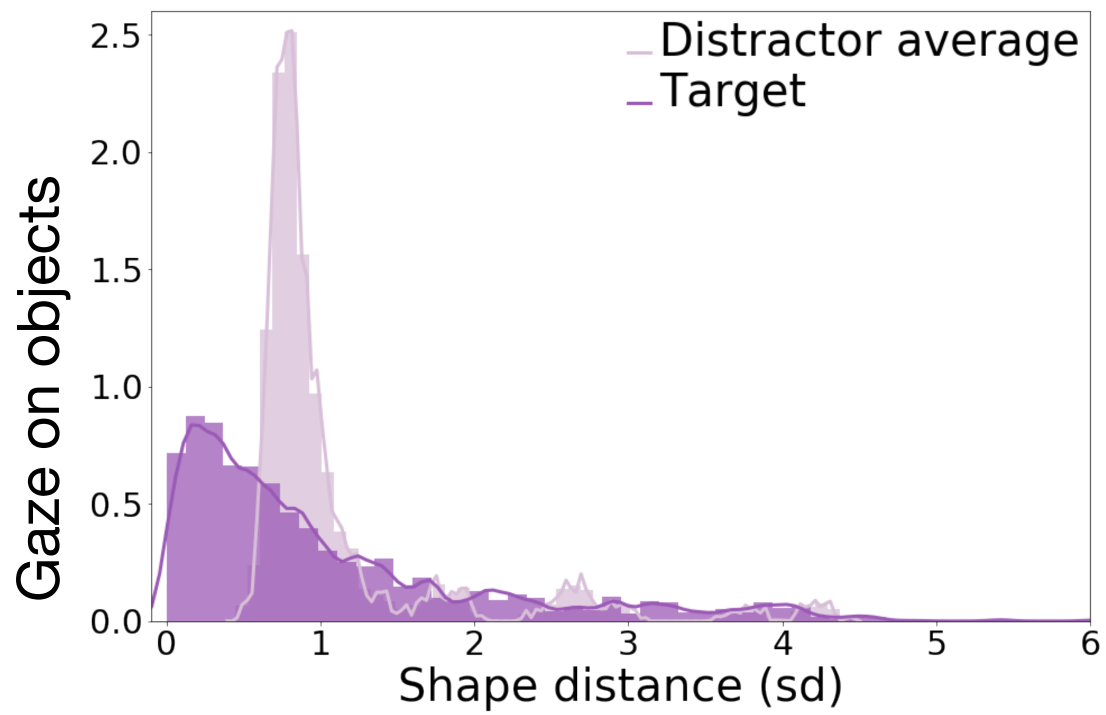

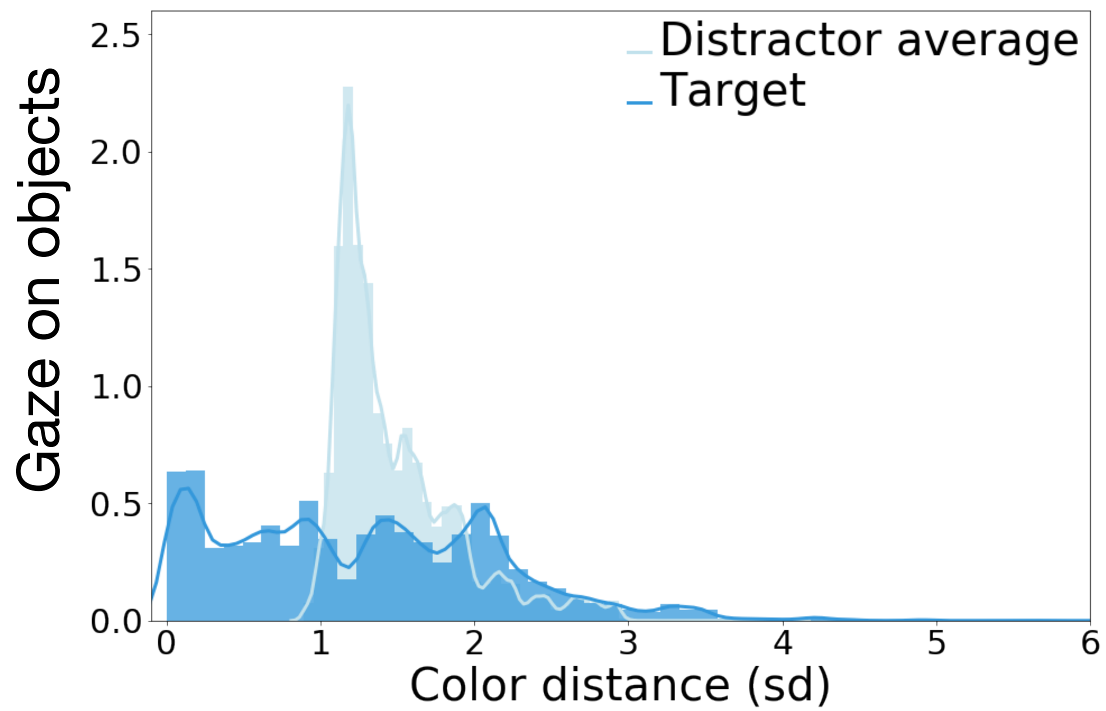

Shape and color predict gaze

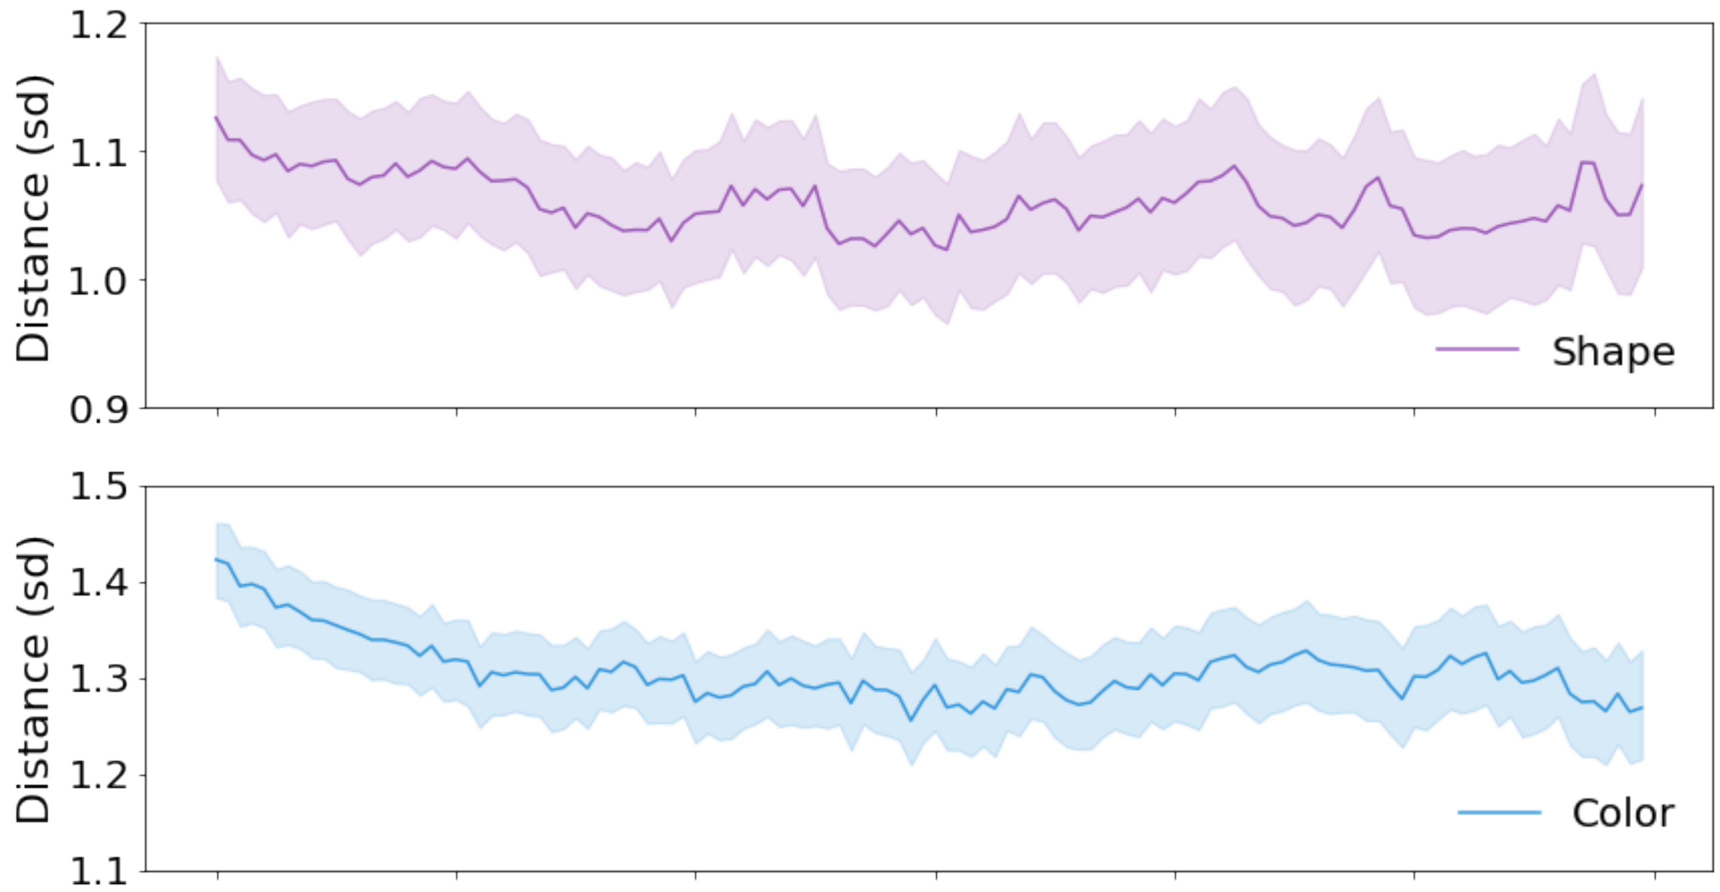

Challenge I: representing the belief space

Challenge II: finding the optimal policy

"Ideal observer" model of visual search

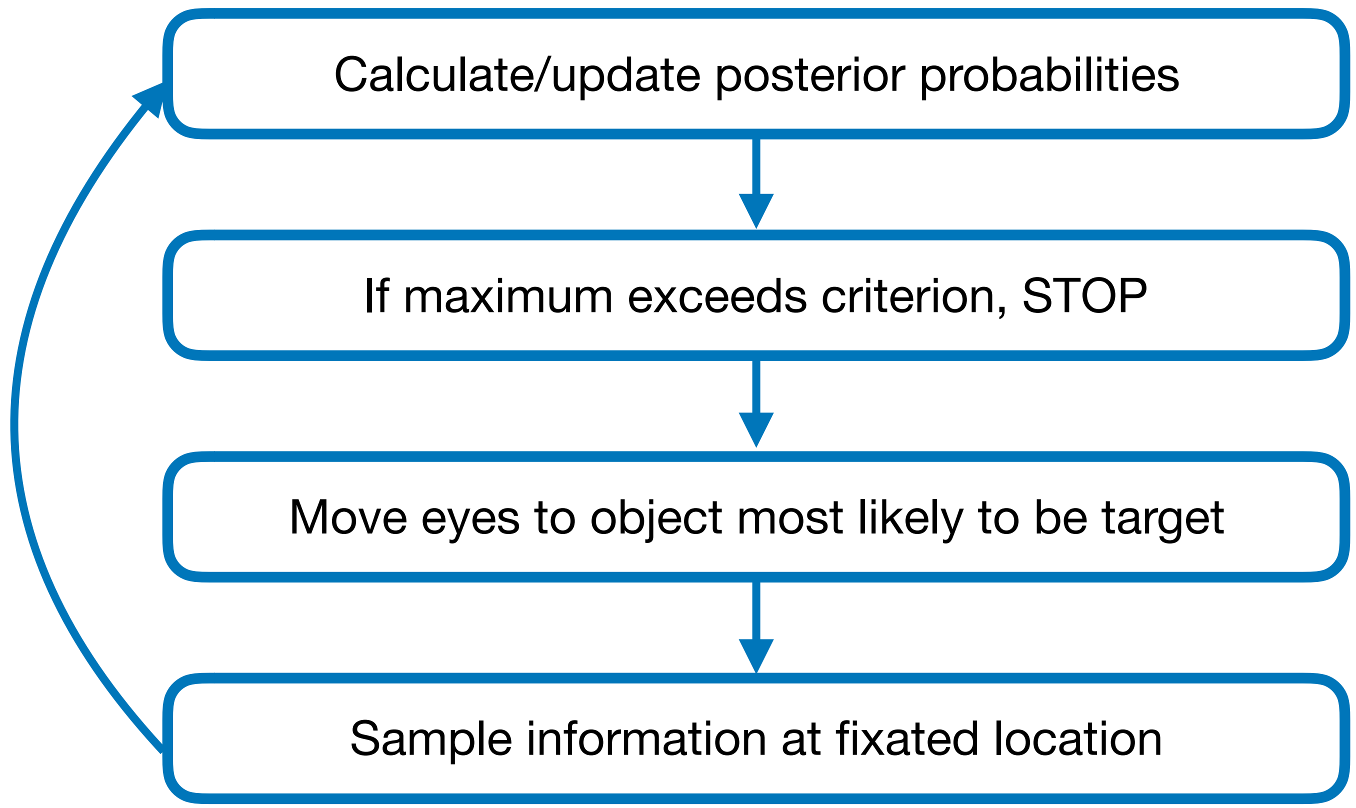

- Can be expressed as a policy in the meta-MDP, but not necessarily optimal

Najemnik and Geisler, 2005
Yang, Lengyel and Wolpert, 2017

Optimizing meta-level return with deep reinforcement learning

- Proximal Policy Optimization (PPO, Schulman, 2017), implemented with tf-agents
- 10 replications, manually tuned hyper-parameters
- Manual tweaking of input representation & initialization

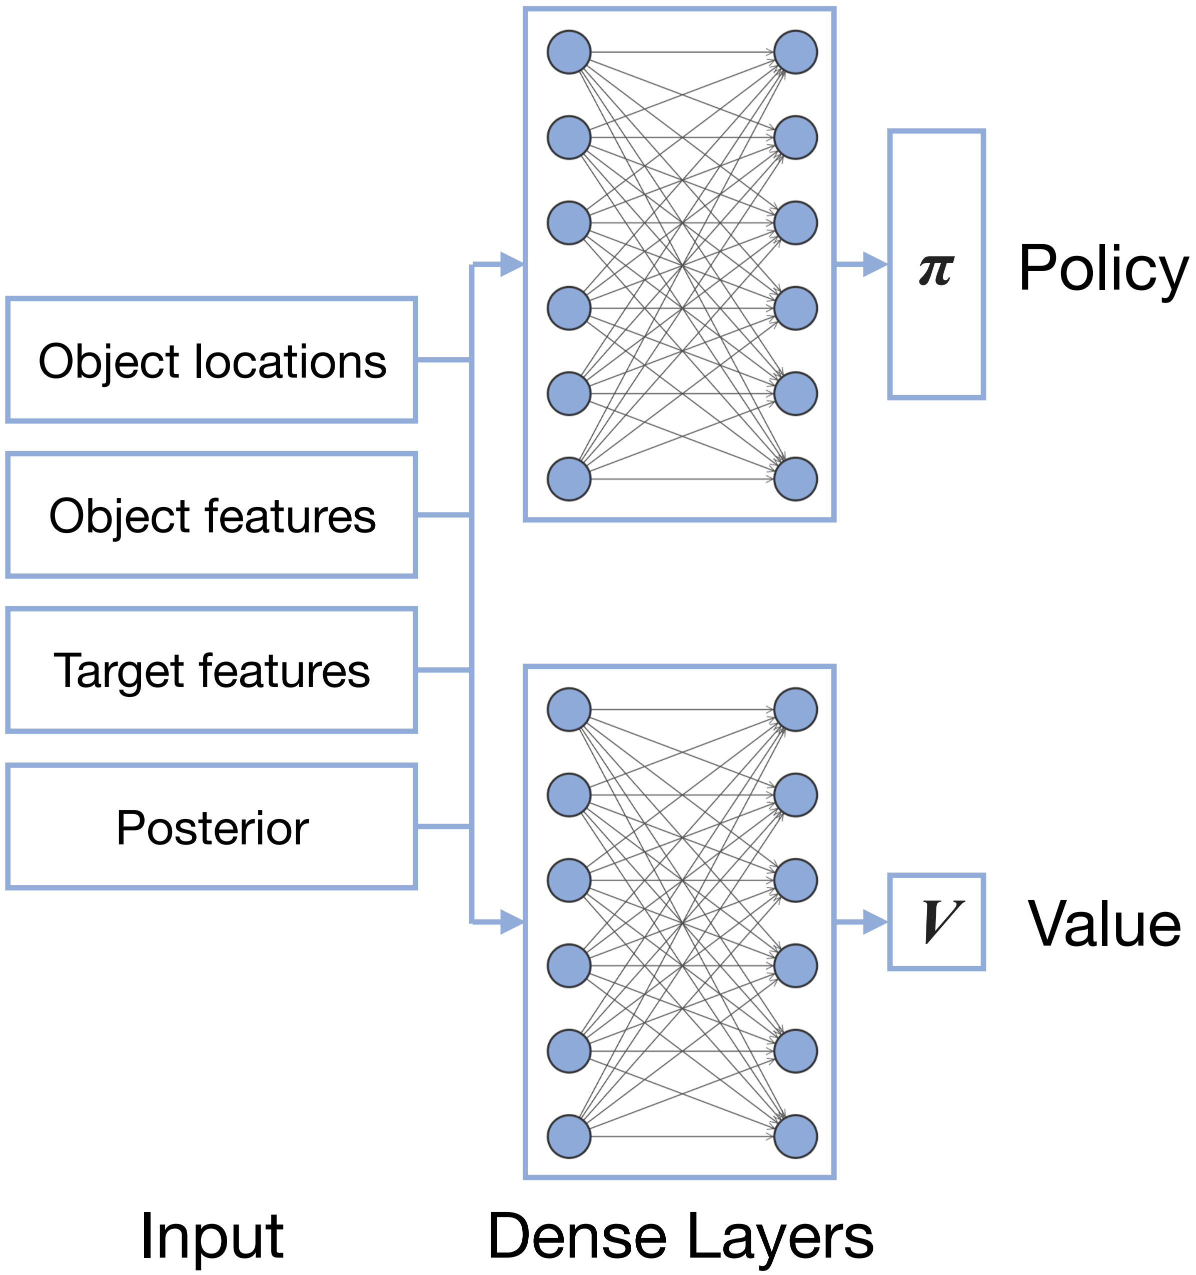

Optimizing meta-level return with deep reinforcement learning

- Proximal Policy Optimization (PPO, Schulman, 2017), implemented with tf-agents
- 10 replications, manually tuned hyper-parameters
- Manual tweaking of input representation & initialization

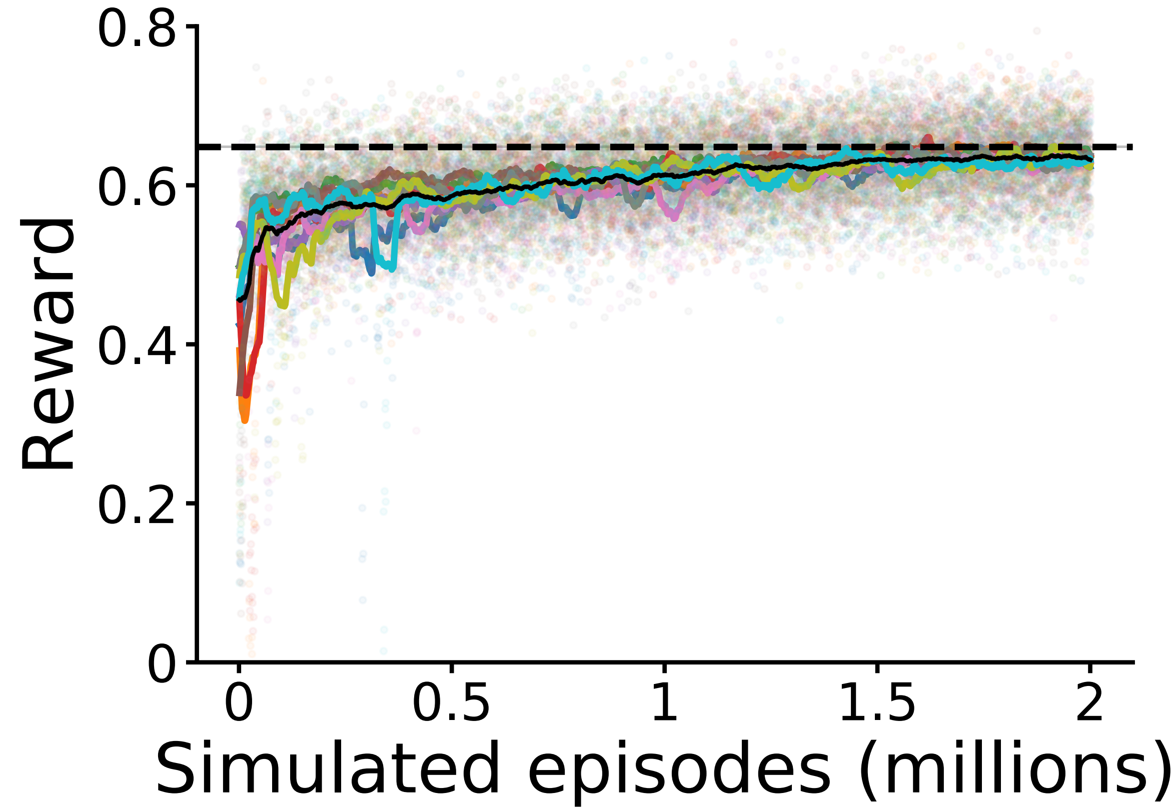

Optimizing meta-level return with deep reinforcement learning

- Proximal Policy Optimization (PPO, Schulman, 2017), implemented with tf-agents
- 10 replications, manually tuned hyper-parameters
- Manual tweaking of input representation & initialization

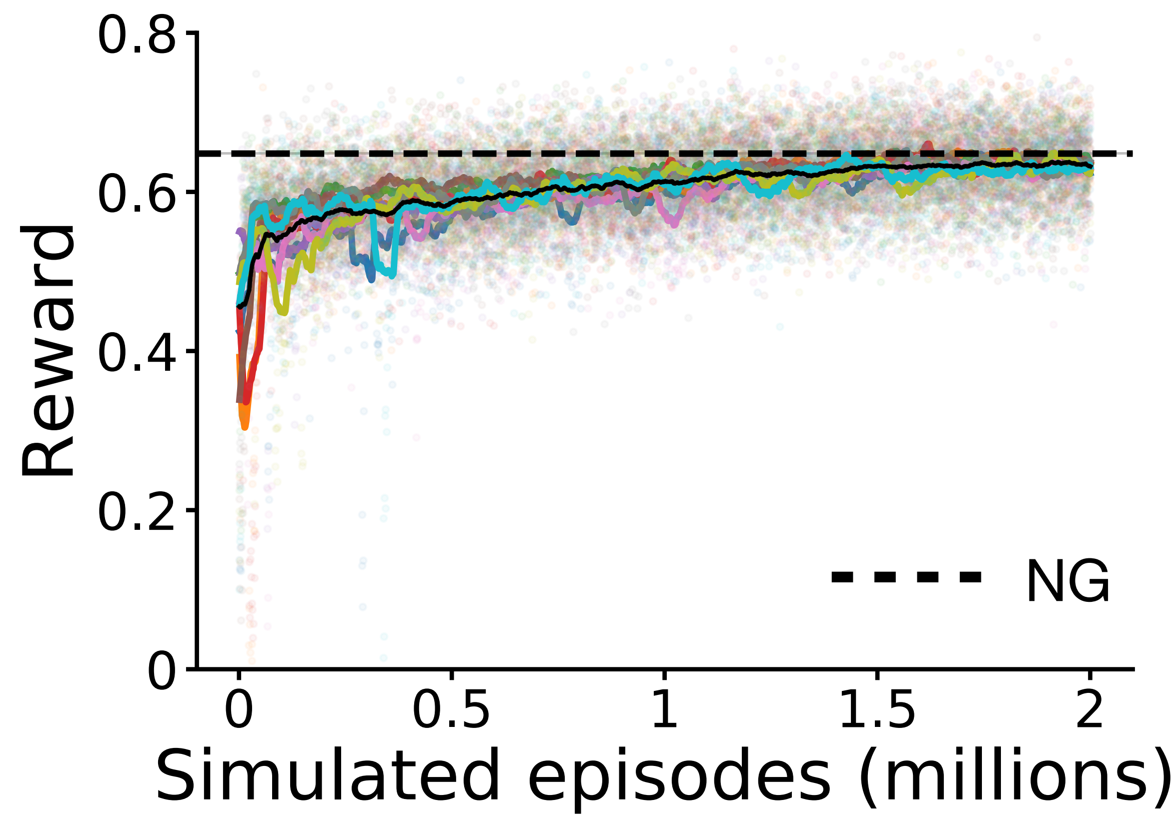

Does optimal policy match humans?

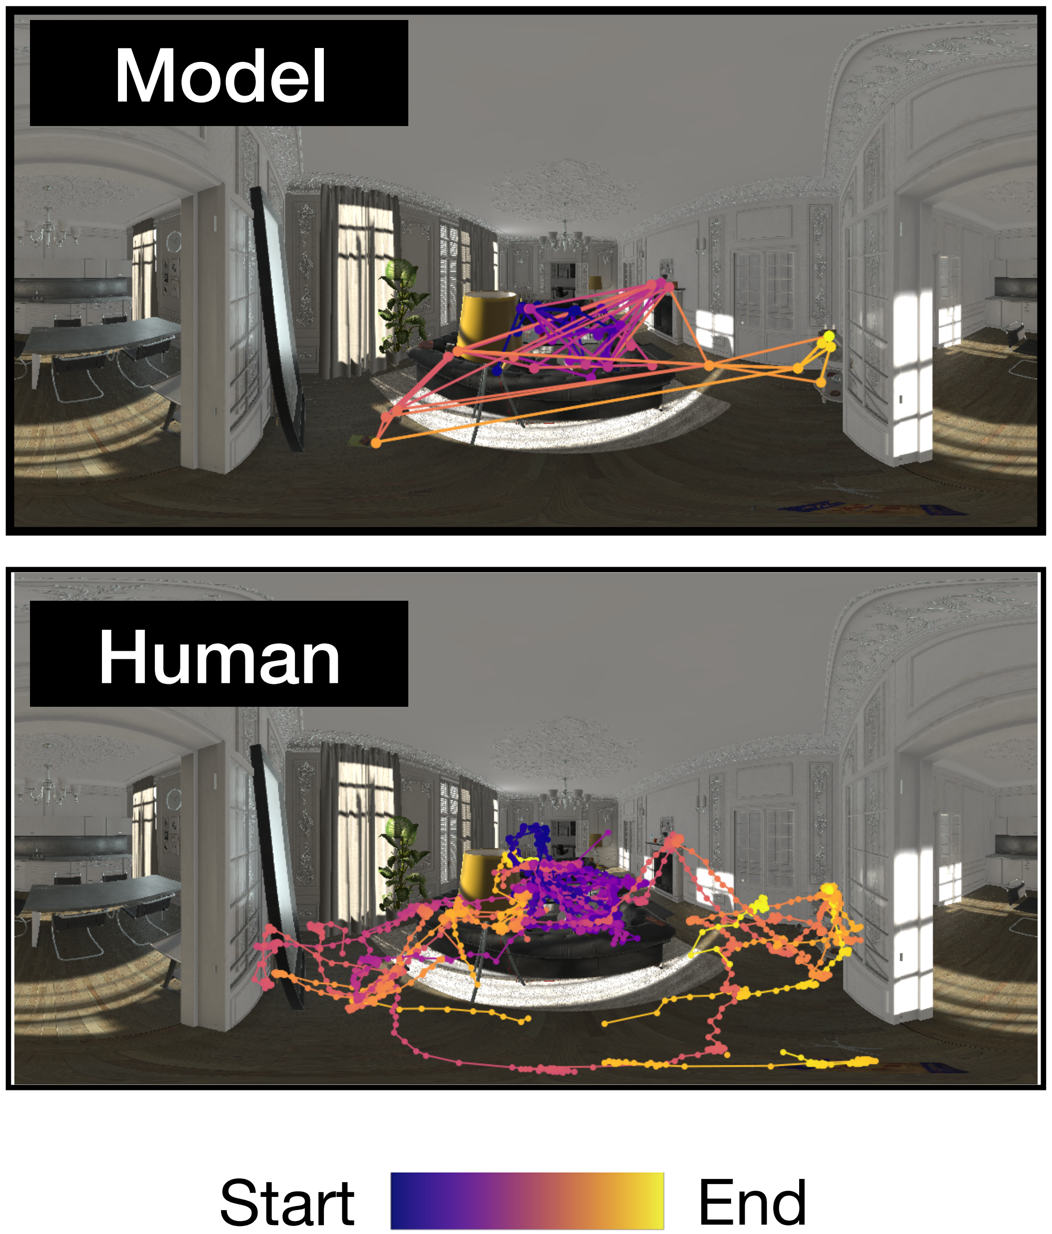

Does optimal policy match humans?

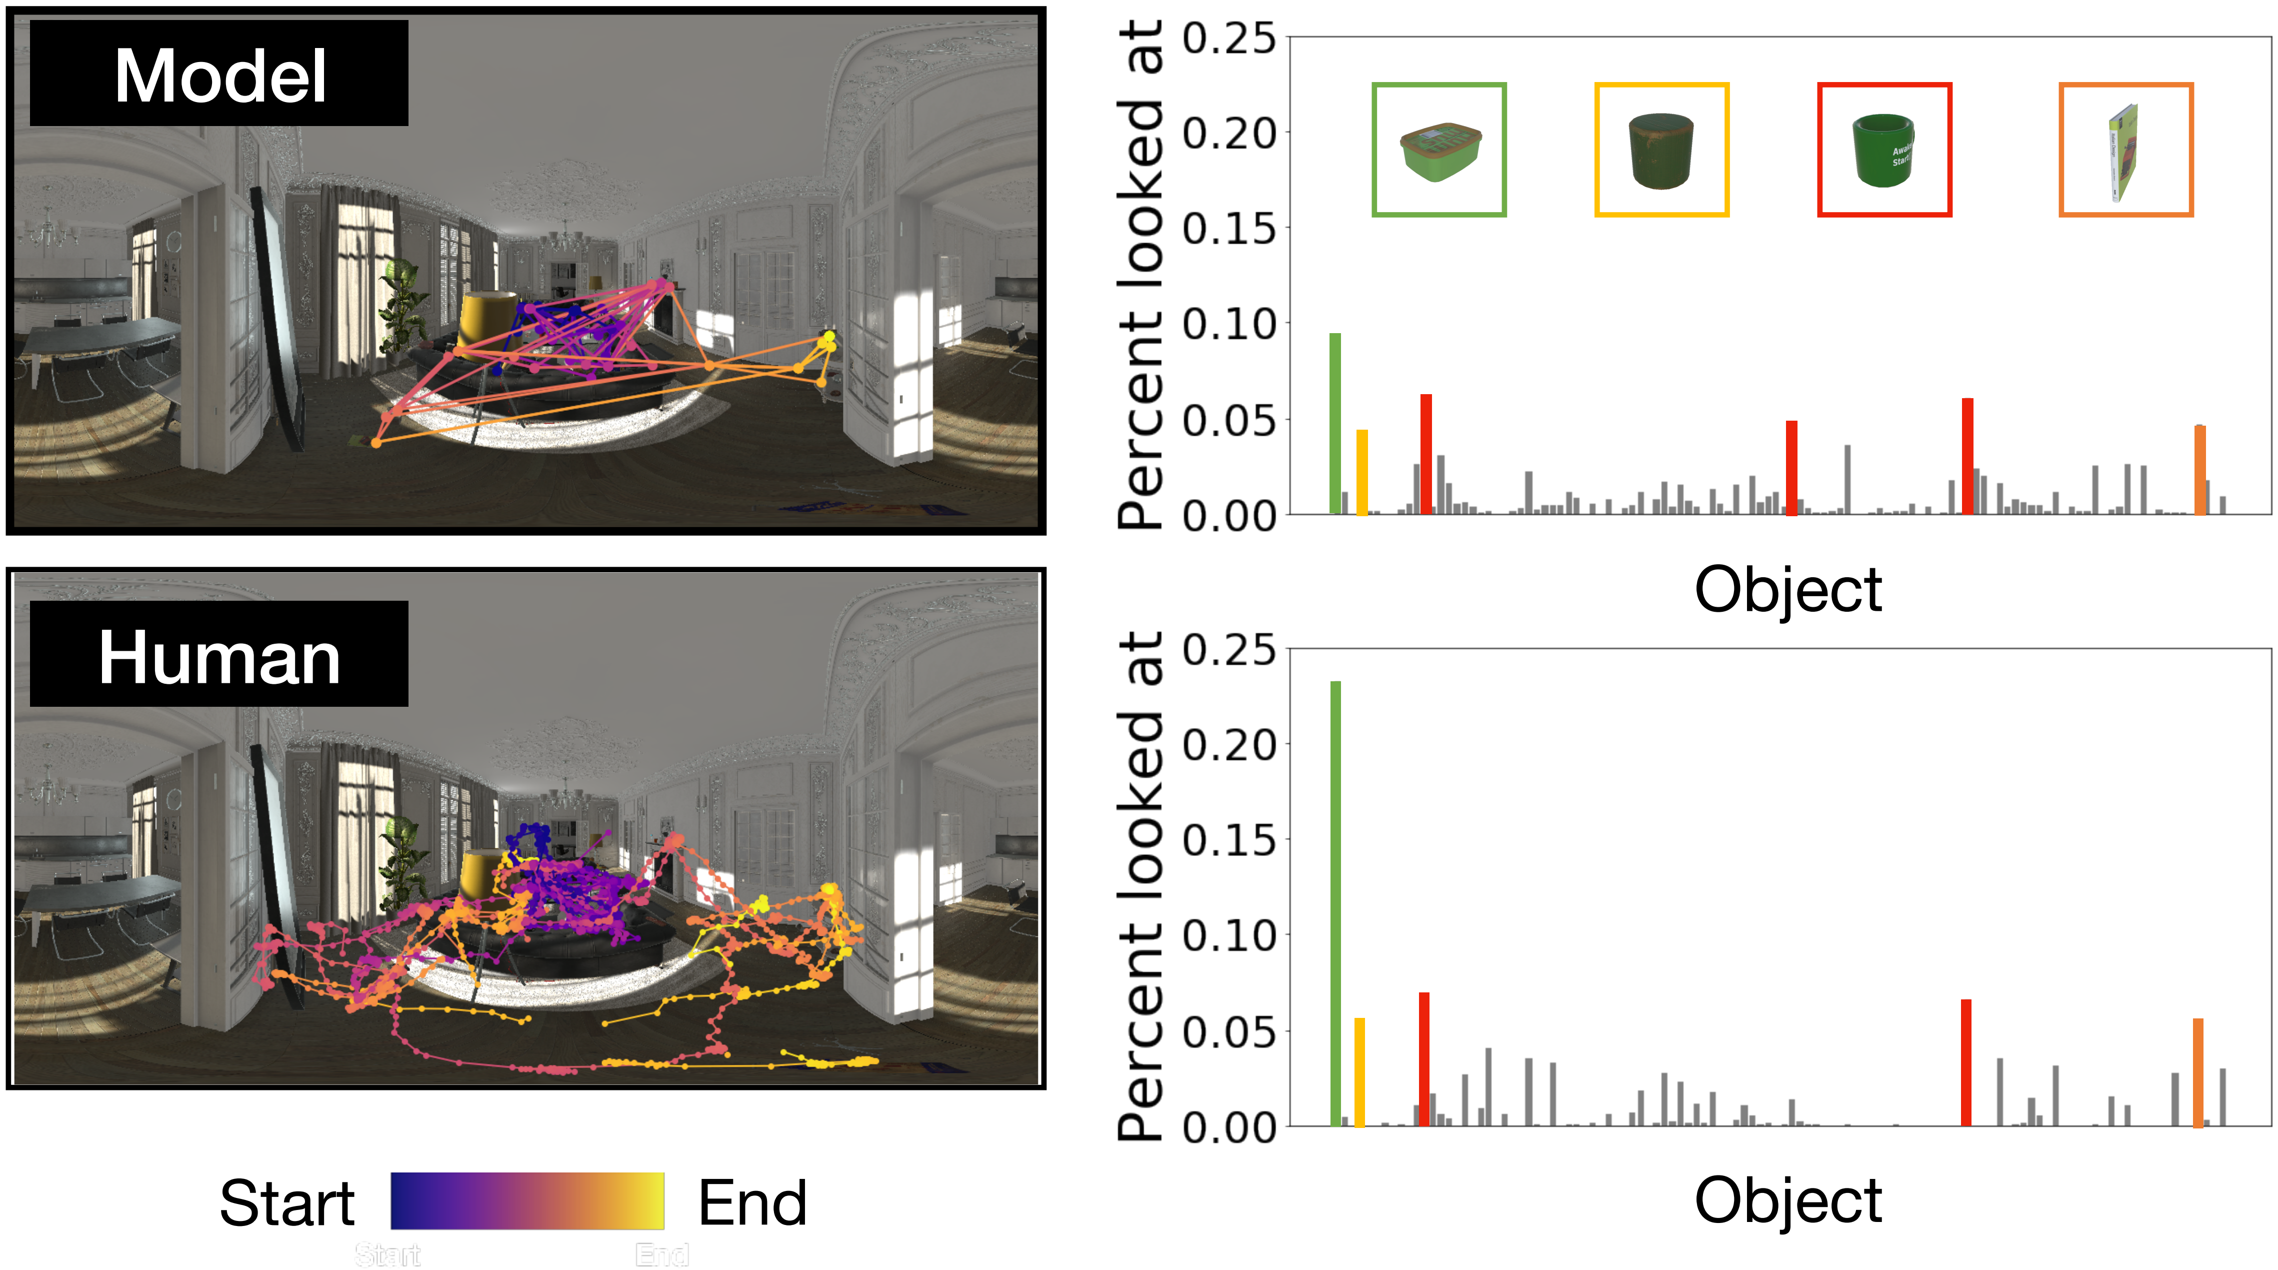

Which features drive human search?

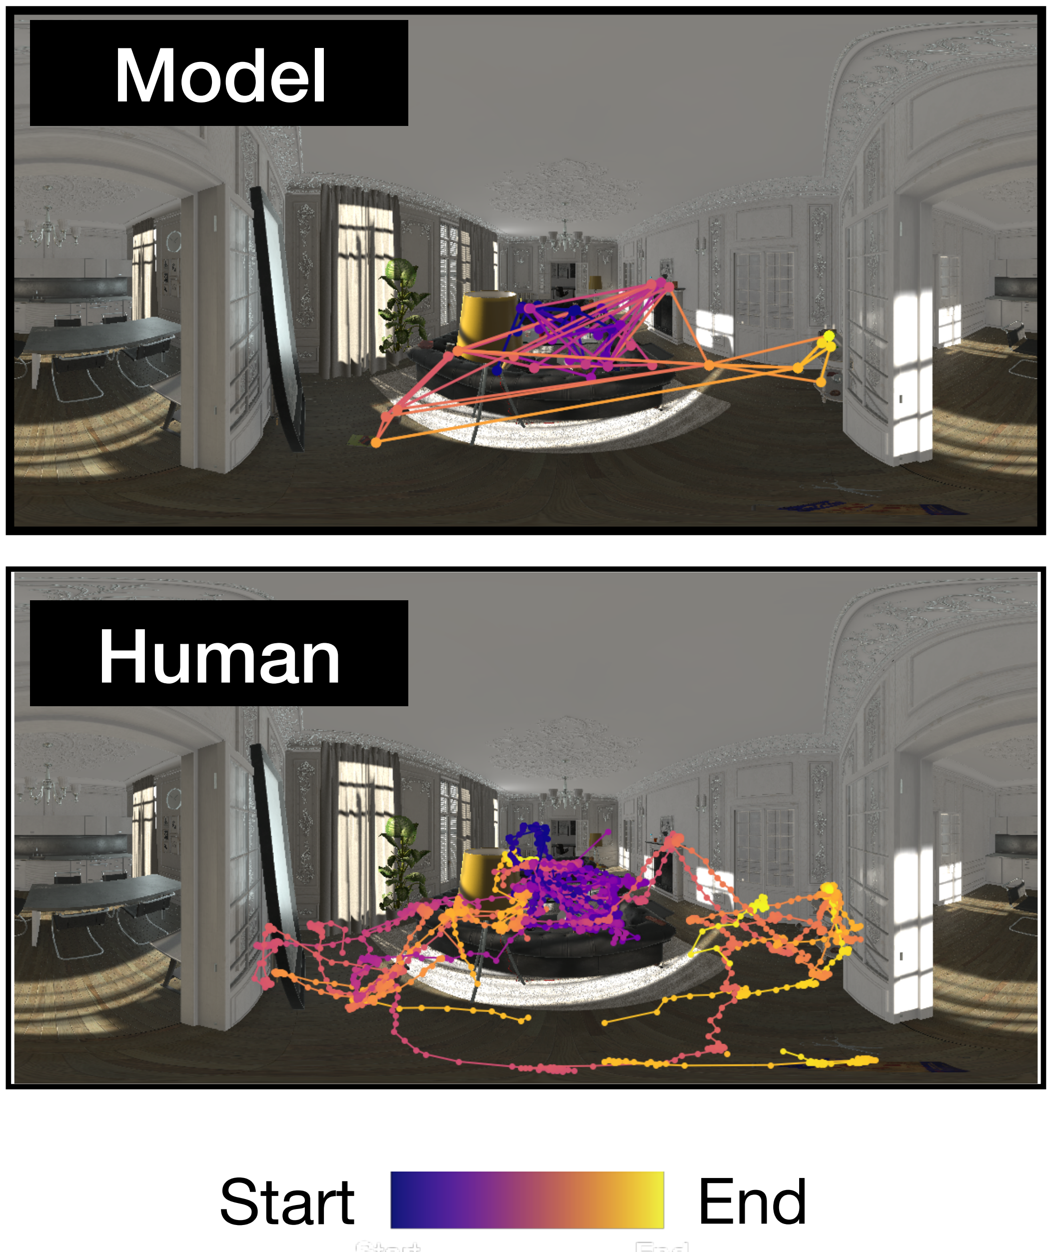

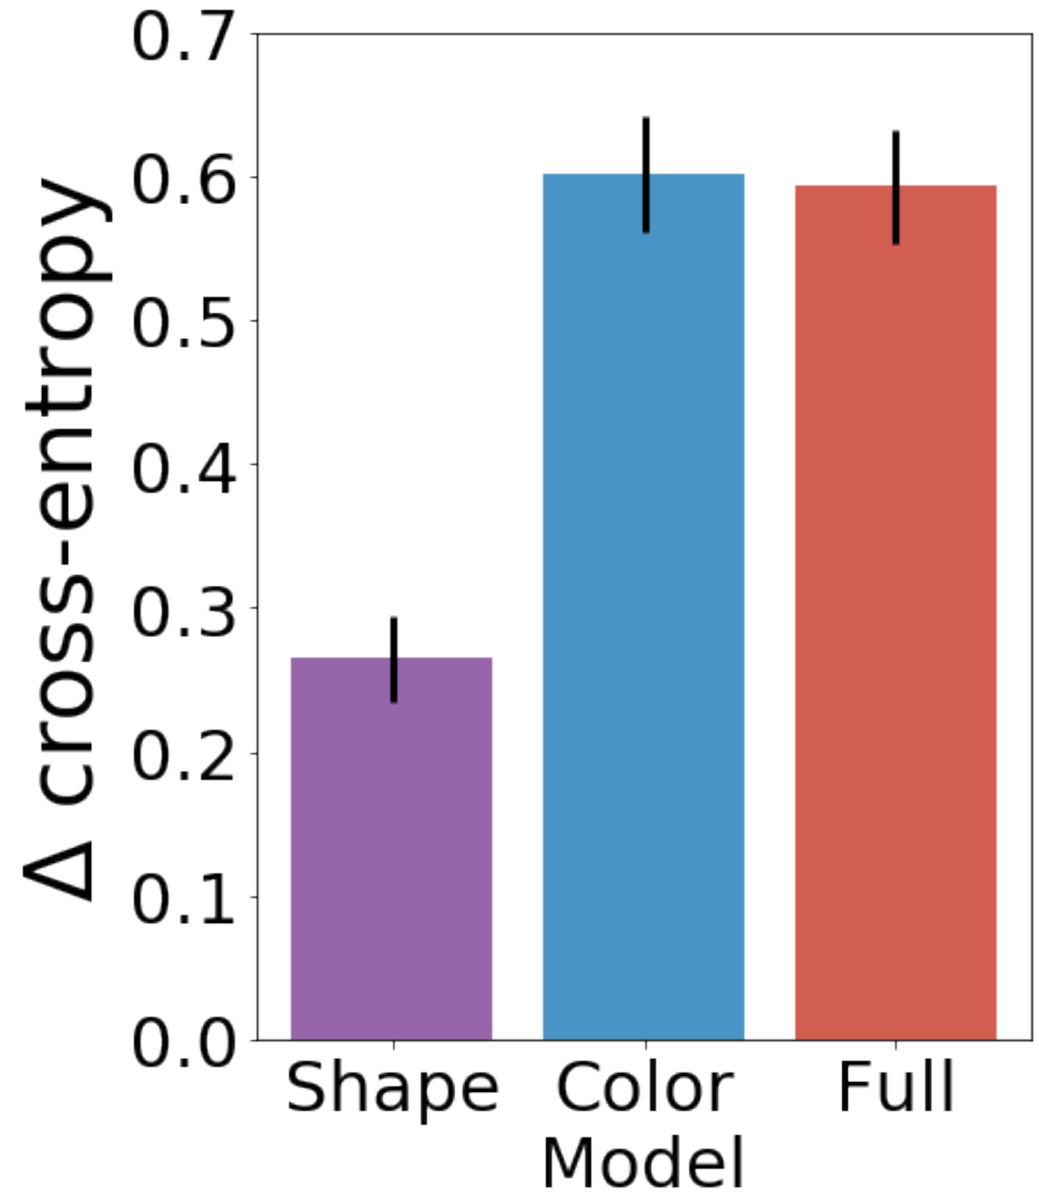

Ongoing work

- Alternative schemes for **extracting low-dimensional feature representations of objects**
    - Deep convolutional neural network models of human ventral visual stream (Yamins et al. 2014, Fan et al. 2019)
    - MeshNet model of 3D shape representation (Feng et al. 2018)

Ongoing work

- Alternative schemes for **extracting low-dimensional feature representations of objects**
    - Deep convolutional neural network models of human ventral visual stream (Yamins et al. 2014, Fan et al. 2019)
    - MeshNet model of 3D shape representation (Feng et al. 2018)
- Investigating the learned policy
    - Is it optimal?

Thank you!