Estimating perceptual priors with finite experiments

Berkeley Optometry & Vision Science

Tyler S. Manning[1], Iona R. McLean[1], Benjamin Naecker[2], Jonathan W. Pillow[2], Bas Rokers[3], Emily A. Cooper[1]

[1]University of California, Berkeley, [2]Princeton University, [3]NYU Abu Dhabi

tmanning@berkeley.edu

This work was funded by the NIH (F32 EY032321 & T32 EY007043)

Prior knowledge of stimulus statistics influences perception

Biases in perceptual judgments well-modeled with a Bayesian ideal observer that combines noisy measurements with these priors (e.g. Weiss et al. 2002)

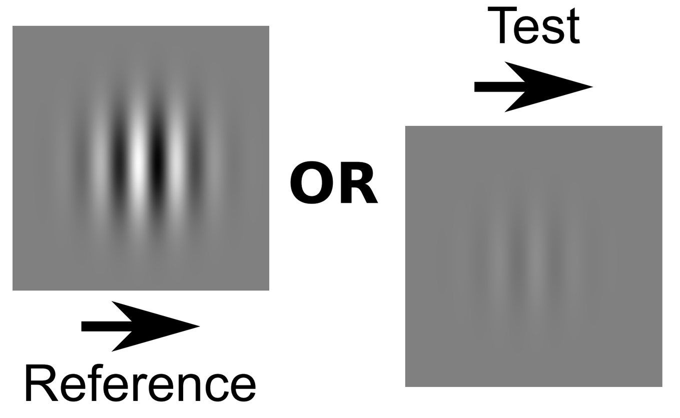

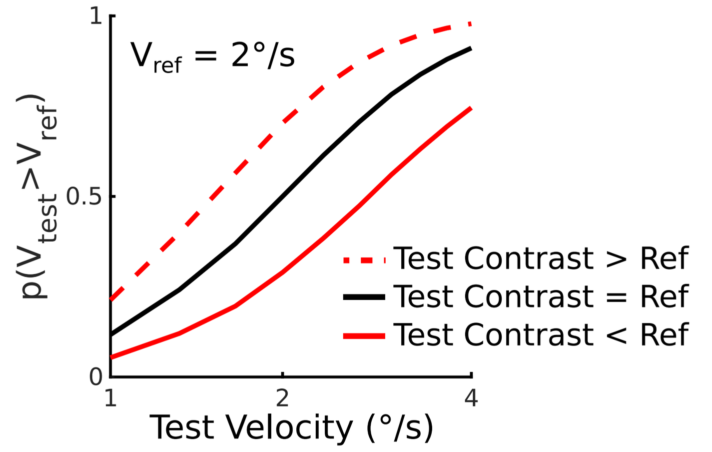

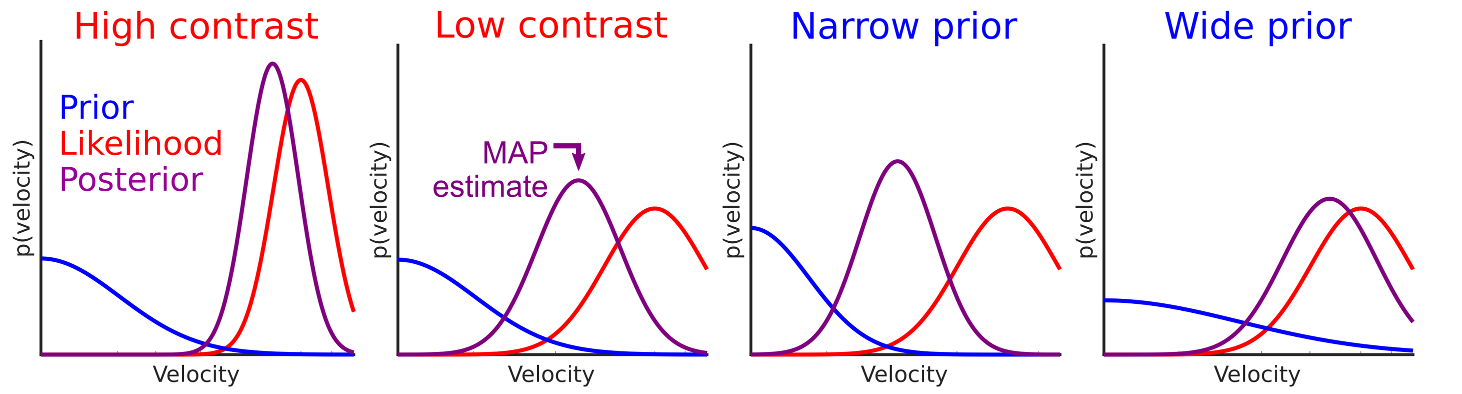

Priors should reflect natural statistics and may change during statistical learning

Visual motion priors assumed to be weighted towards slow speeds - reflecting world and retinal motion (Aytekin et al. 2014); though estimates of priors are variable between observers and conditions (Jogan & Stocker 2015, Rokers et al. 2018) making it difficult to link priors with natural statistics

Priors may change by manipulating stimulus statistics in an experimental setting (Sotiropoulos et al. 2011, Adams et al. 2004), but the principles with which they do is not well understood.

Aim:

Success in characterizing priors depends on the quality and quantity of experimental data and robustness of the models used to estimate priors from these data. Here, we investigate model performance within a synthetic experimental framework.

Prior estimation frameworks

We assess the performance of two methods: estimating the prior as a ***piecewise function*** or a ***mixture of Gaussians***

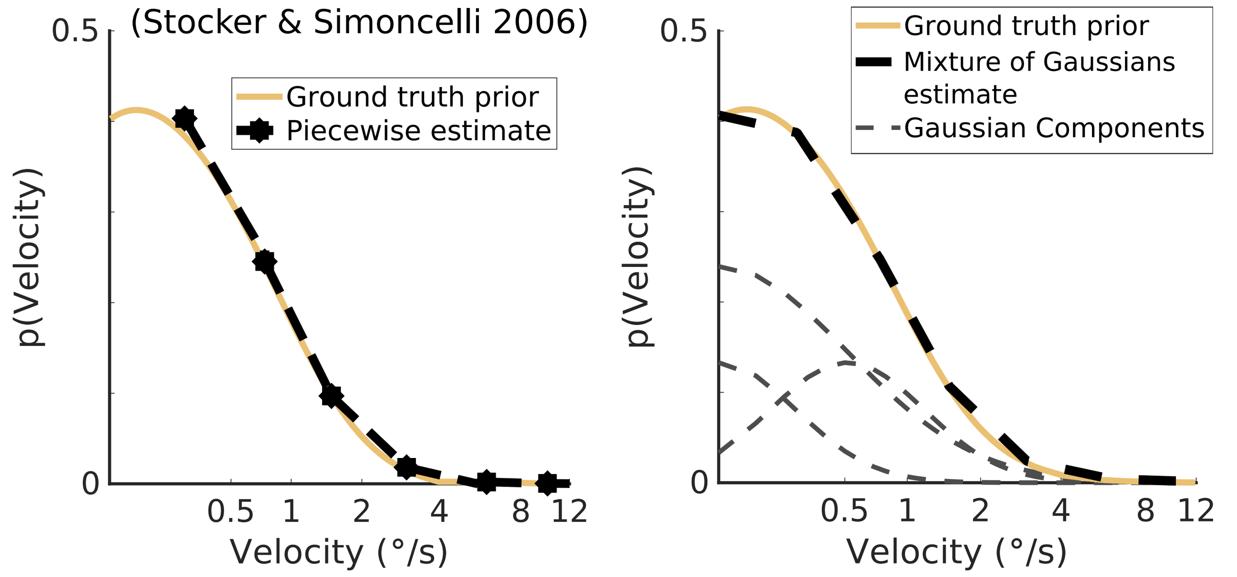

Simulating experiments using synthetic observers with known priors

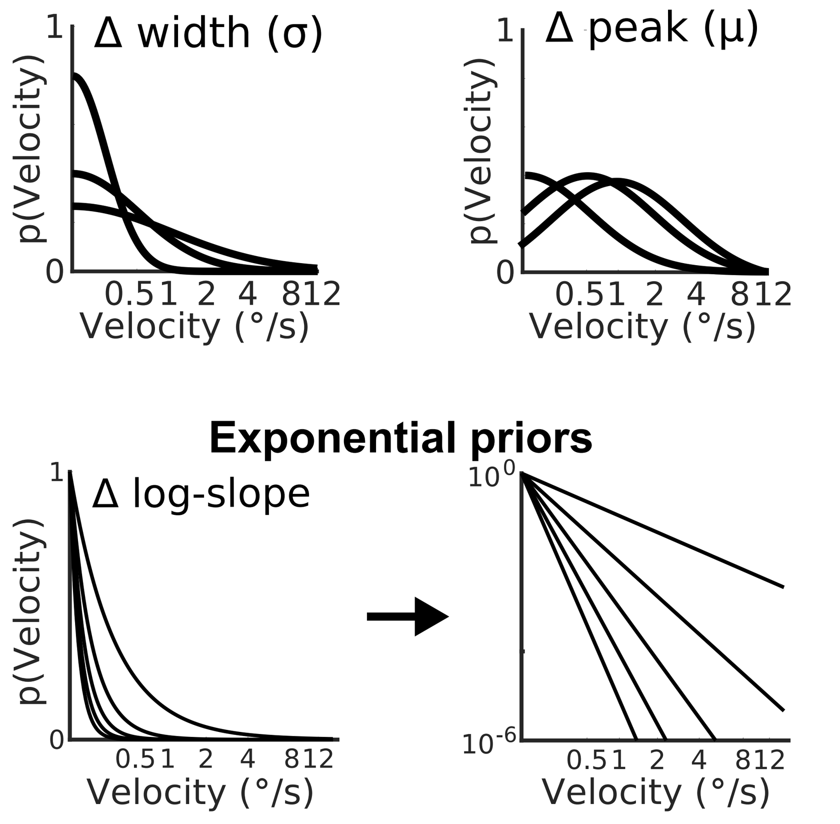

- 14 different observer priors (9 Gaussian, 5 exponential) that vary in *standard deviation* (σ), *maximum* (μ), and *kurtosis* (log-slope)

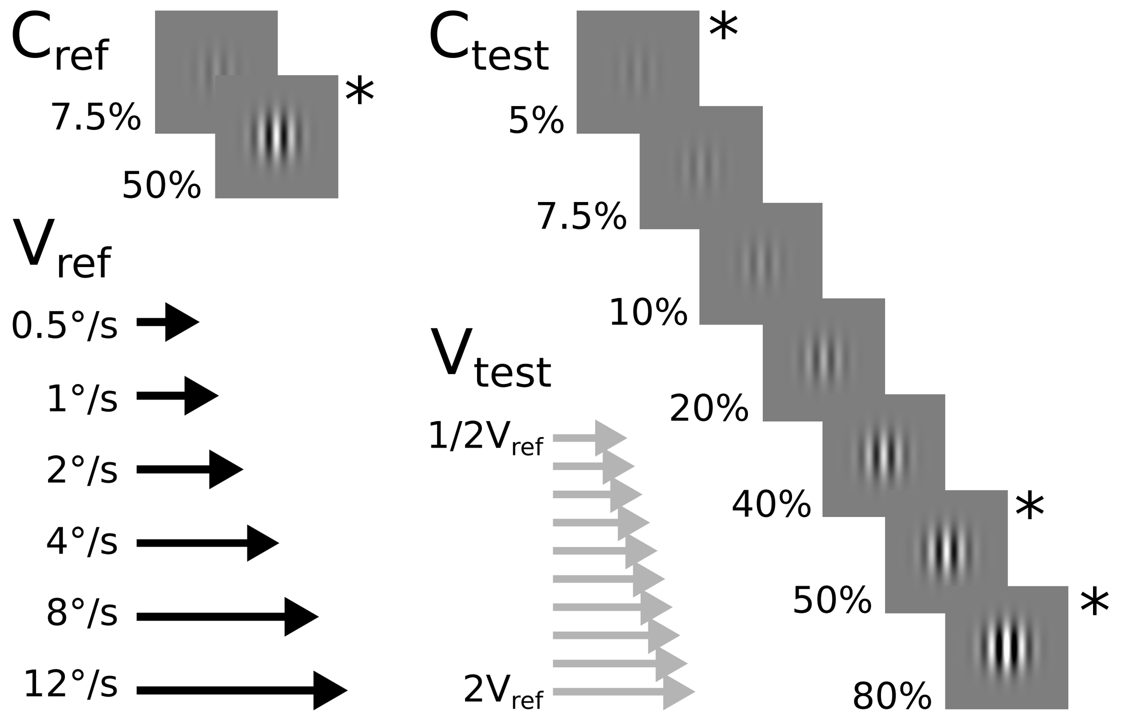

- 2AFC 1D speed judgment experiment (Stocker & Simoncelli 2006)
- 840 unique reference-test pairs, presented 5, 10, or 20 times (4200, 8400, or 16800 trials)
- One truncated experiment (900 trials) with reduced stimulus set (marked with *)
- Experiment repeated 25x per observer

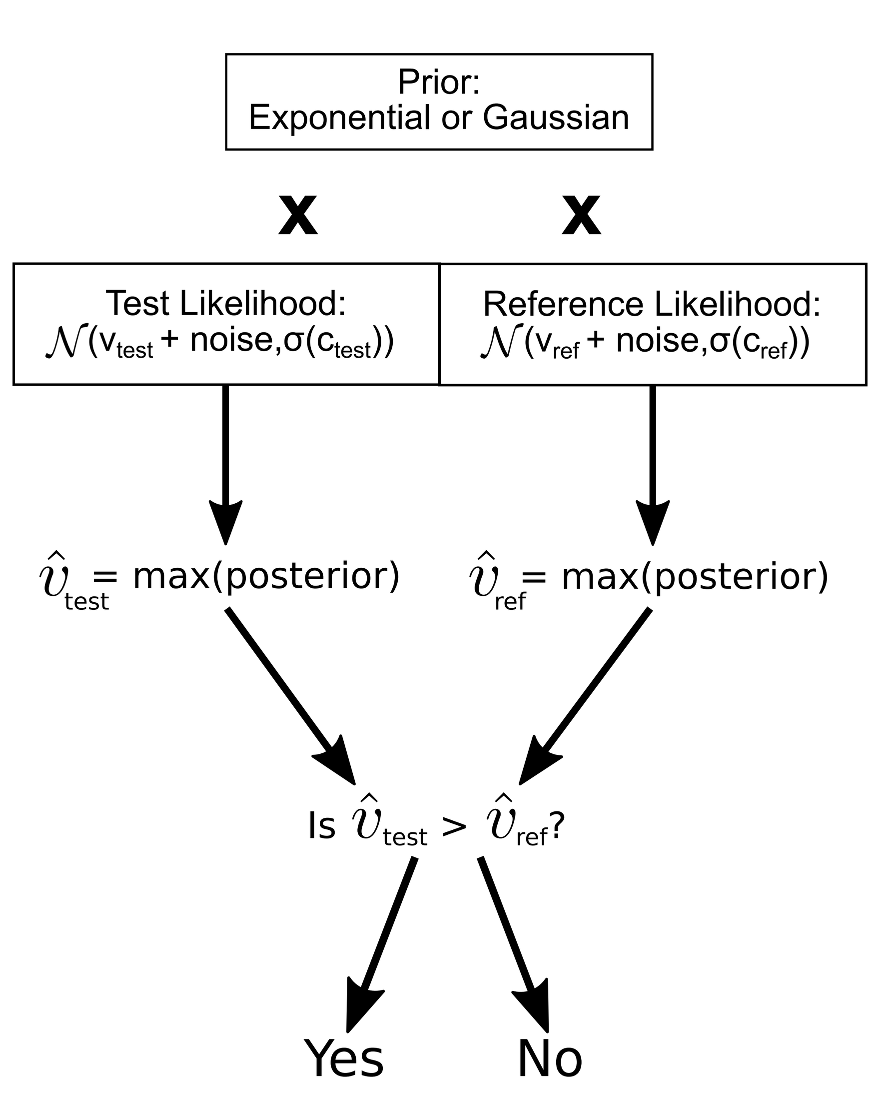

Analysis Pipeline

Accuracy (Jensen-Shannon Divergence)

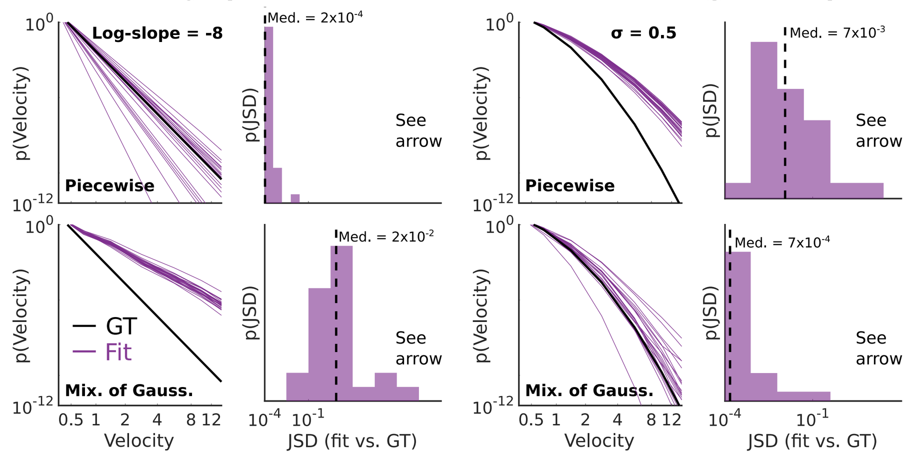

Discriminability (JSD-based precision)

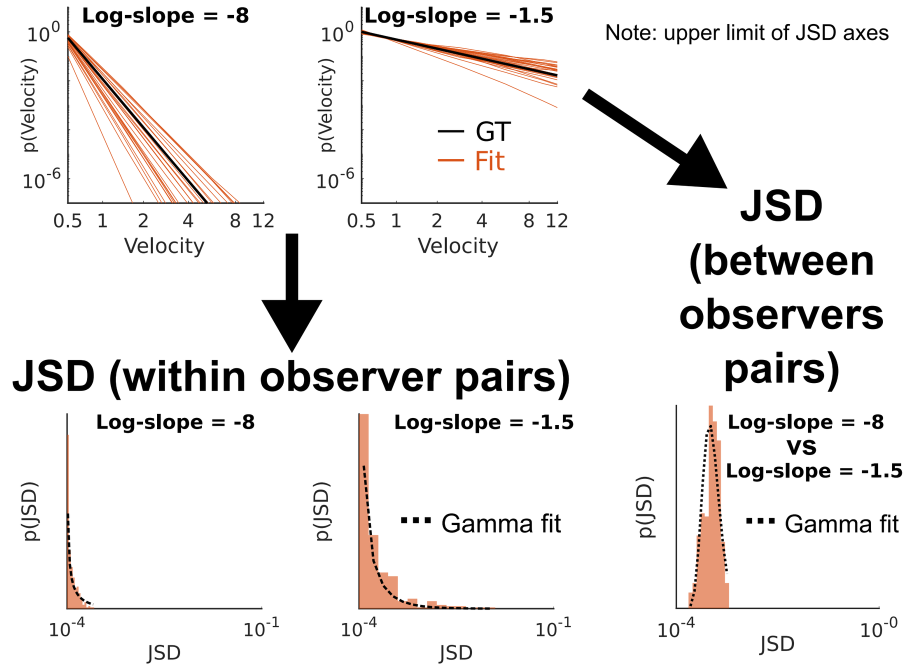

Results: How *accurate* are our prior estimates?

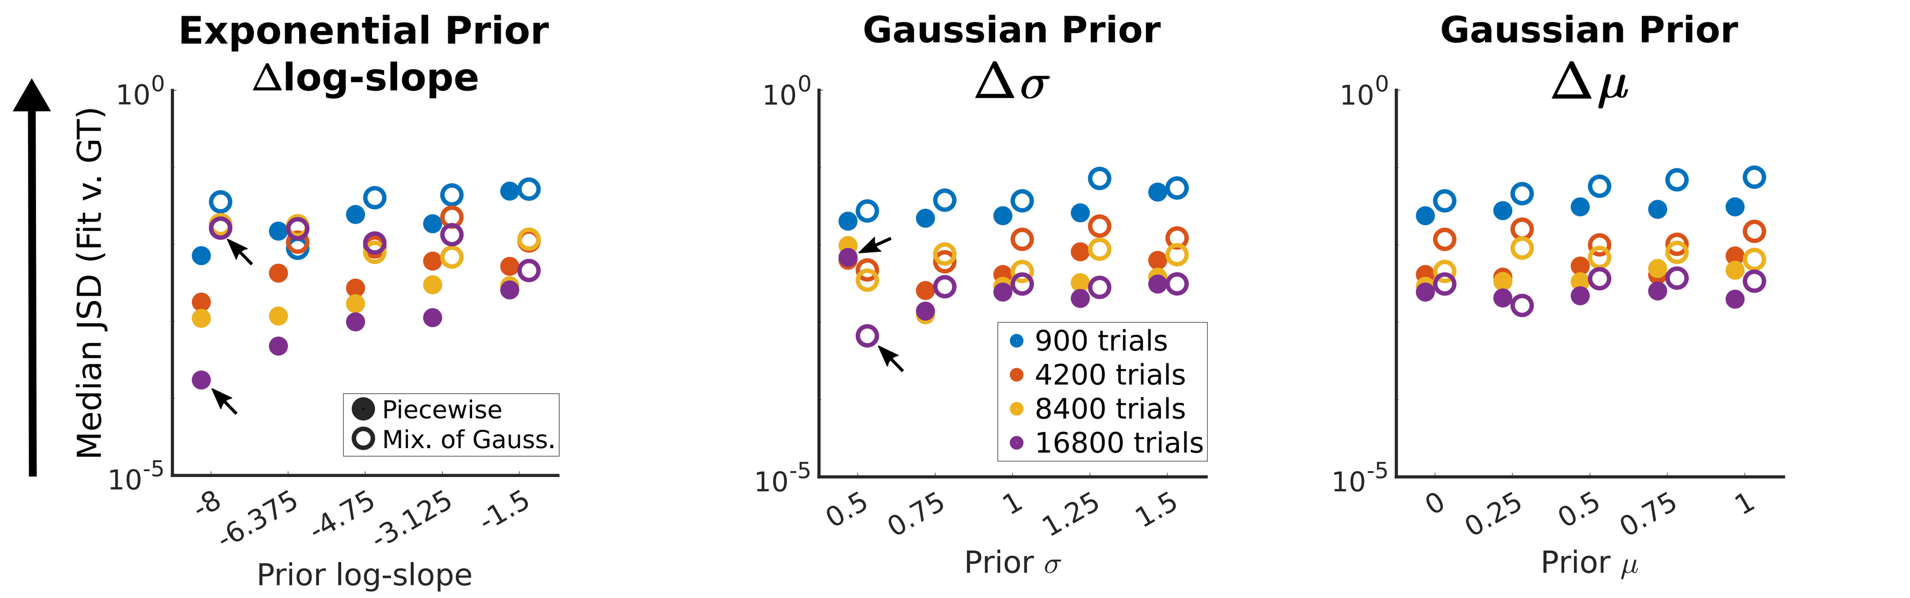

Results: How *different* must priors be to discriminate them?

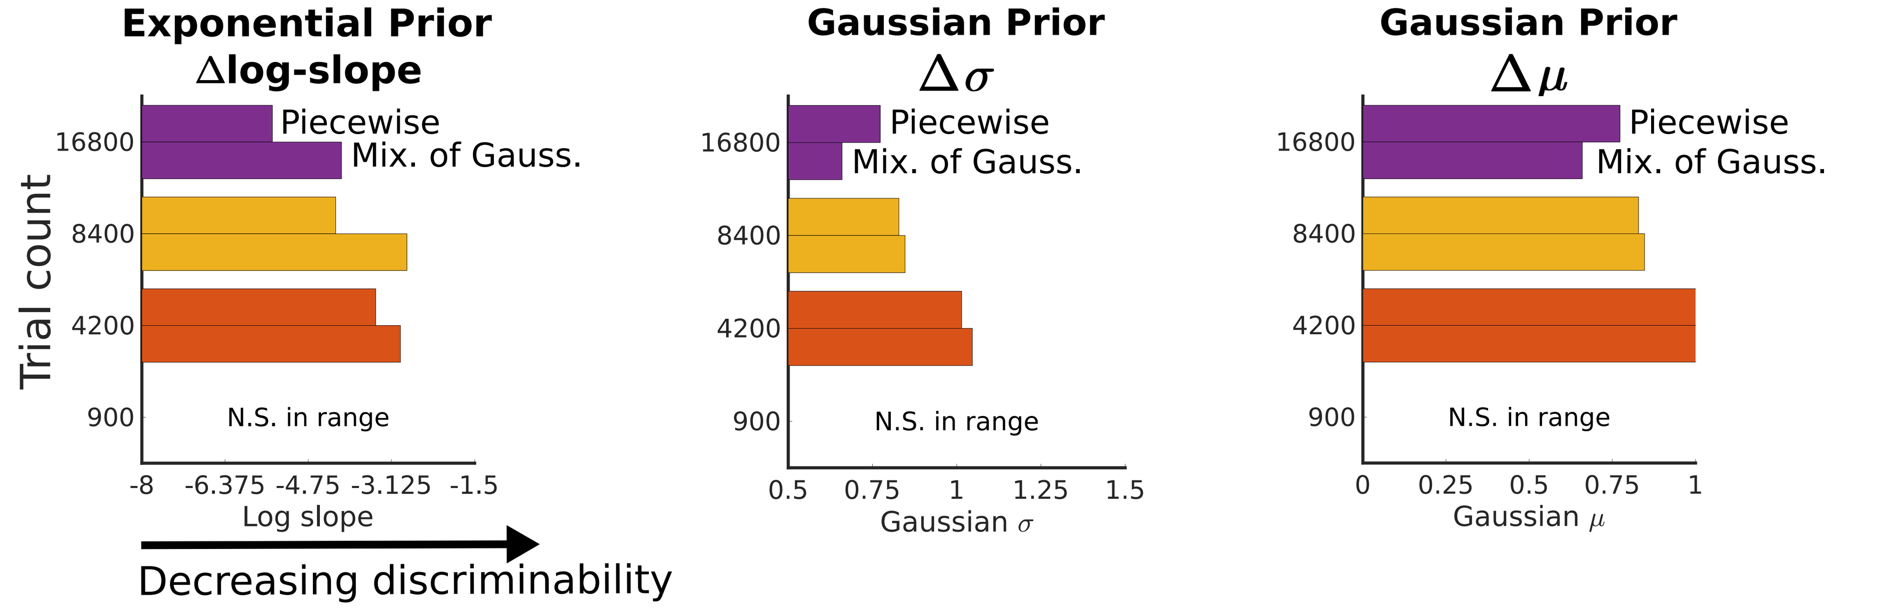

Methods details

Model Parameters:

Piecewise model:
13 (7 likelihood widths and 6 piecewise component log-slopes)

Mixture of Gaussian model:
16 (7 likelihood widths and 3 parameters for each of 3 Gaussian components)

Model Fitting:

p("$v_{test} > v_{ref}$"$|v_{test}, v_{ref}$) determined with signal detection theory; parameters of Bayesian models fit to trial-by-trial data via MLE and constrained nonlinear optimization

Prior Fit Accuracy:

Assessed with Jensen-Shannon divergence (JSD) between fit and ground truth (0->1, lower is more accurate)

Prior Discriminability:

Are two prior estimates different due to variability between samples from the same observer or different observers?

→ Compared to minimum log-slope/σ/μ value, compute parameter where *$p(JSD_{different} > JSD_{same})$* = 0.95 using gamma fit to JSD distributions (within and between observers), linear regression, and signal detection theory

Results: Comparison of model accuracy

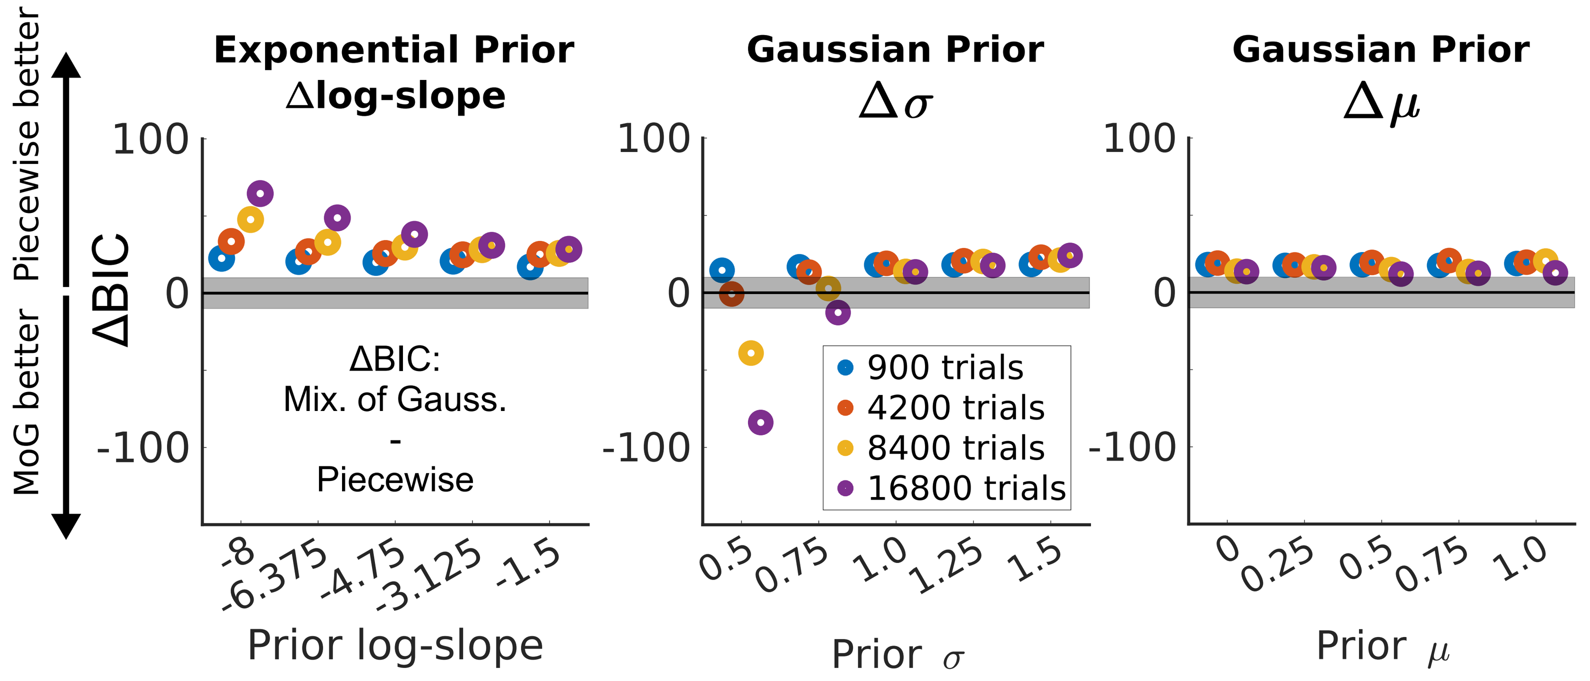

Conclusions

Developed a framework for guiding the experimental design of studies investigating perceptual priors

Long experiments required to accurately estimate priors and to confidently distinguish observers with different priors

Piecewise estimation method is biased towards recovering long-tailed priors; Mixture of Gaussians model easier to implement, but fits long-tails less well

Can use this framework to identify improvements to experimental efficiency based on specific hypotheses about changes in priors