# A neural basis of probabilistic computation in visual cortex

Edgar Y. Walker, Ph.D.

University of Tübingen

# A little bit about myself…

- Computational/theoretical systems neuroscience and machine learning researcher
- Post-doctoral fellow at Sinz Lab in University of Tübingen, Germany
- Developer of DataJoint and Co-founder of Vathes LLC

Edgar Y. Walker, Ph.D.

@eywalker

edgarwalker.com

How does the brain arrive at complex decisions and behaviors in response to sensory information with uncertainty?

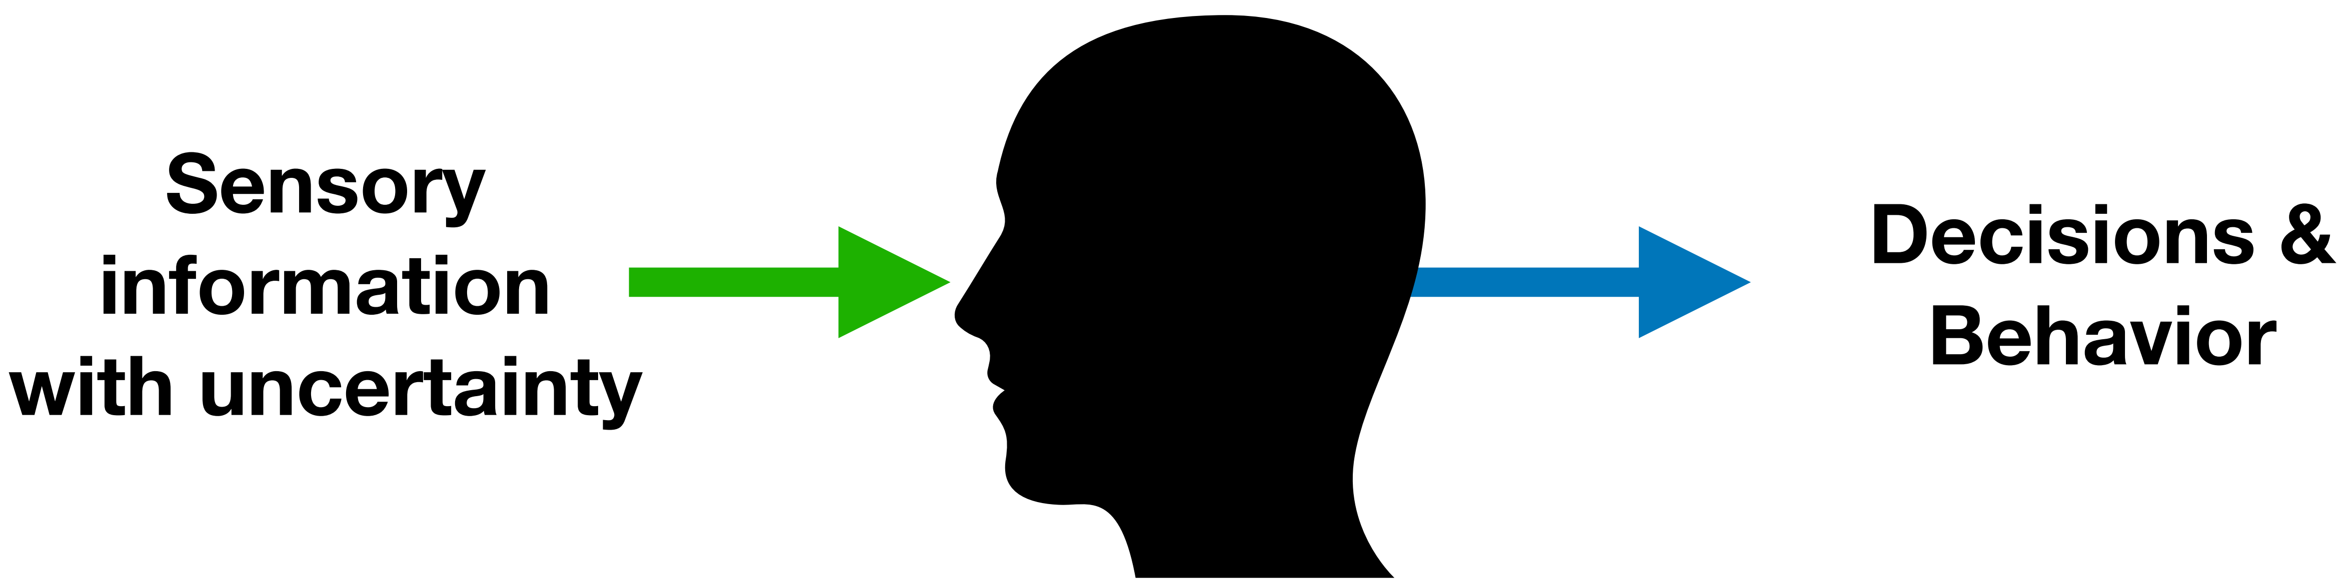

# Brain in action

# Brain in action

# Brain in action

# Flow of information


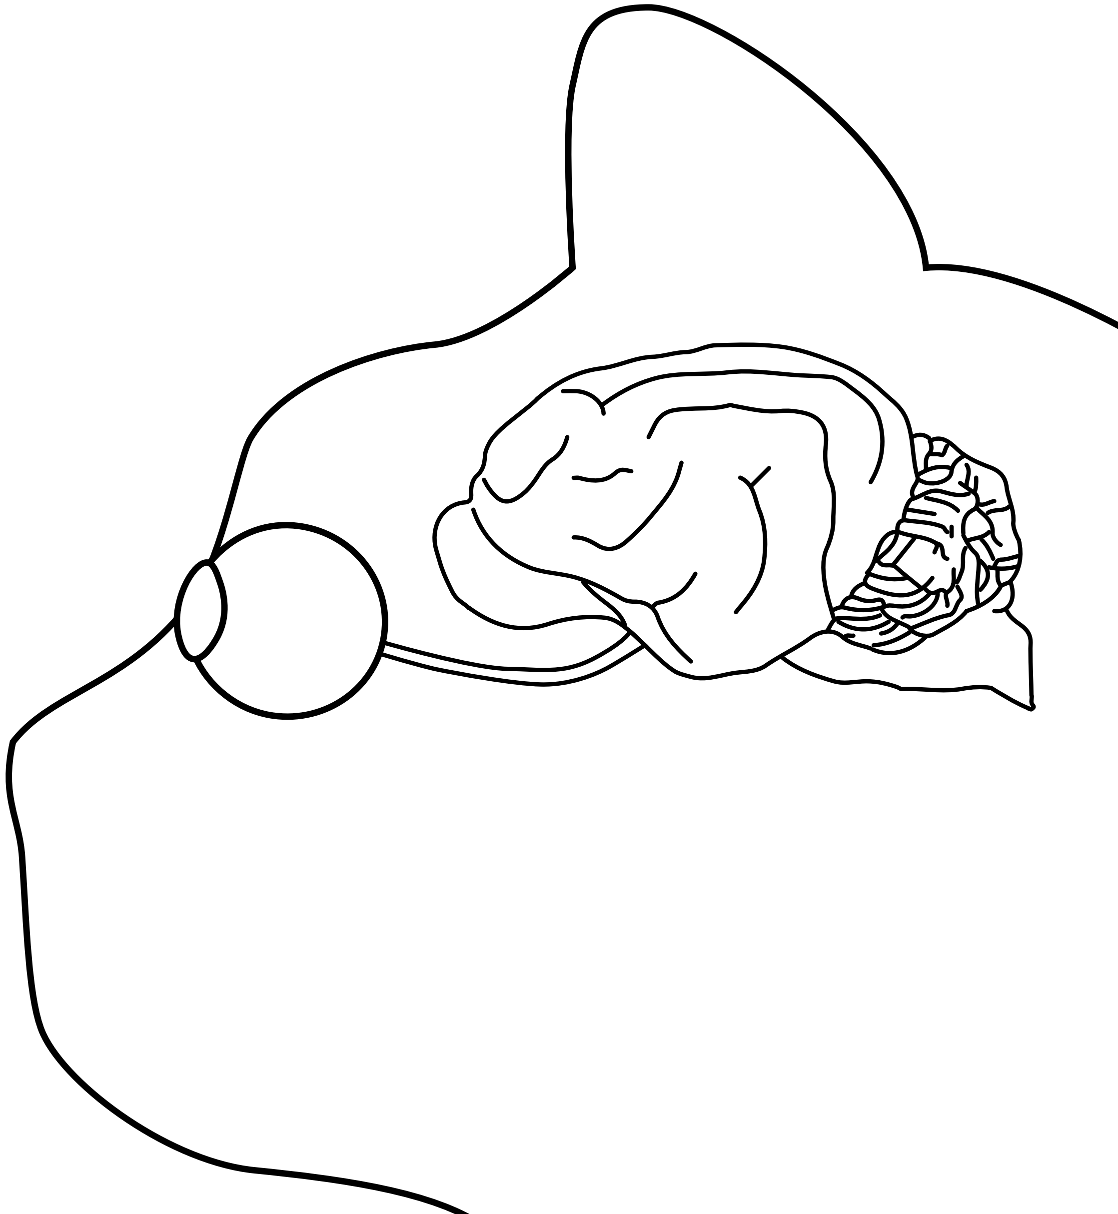

# Flow of information


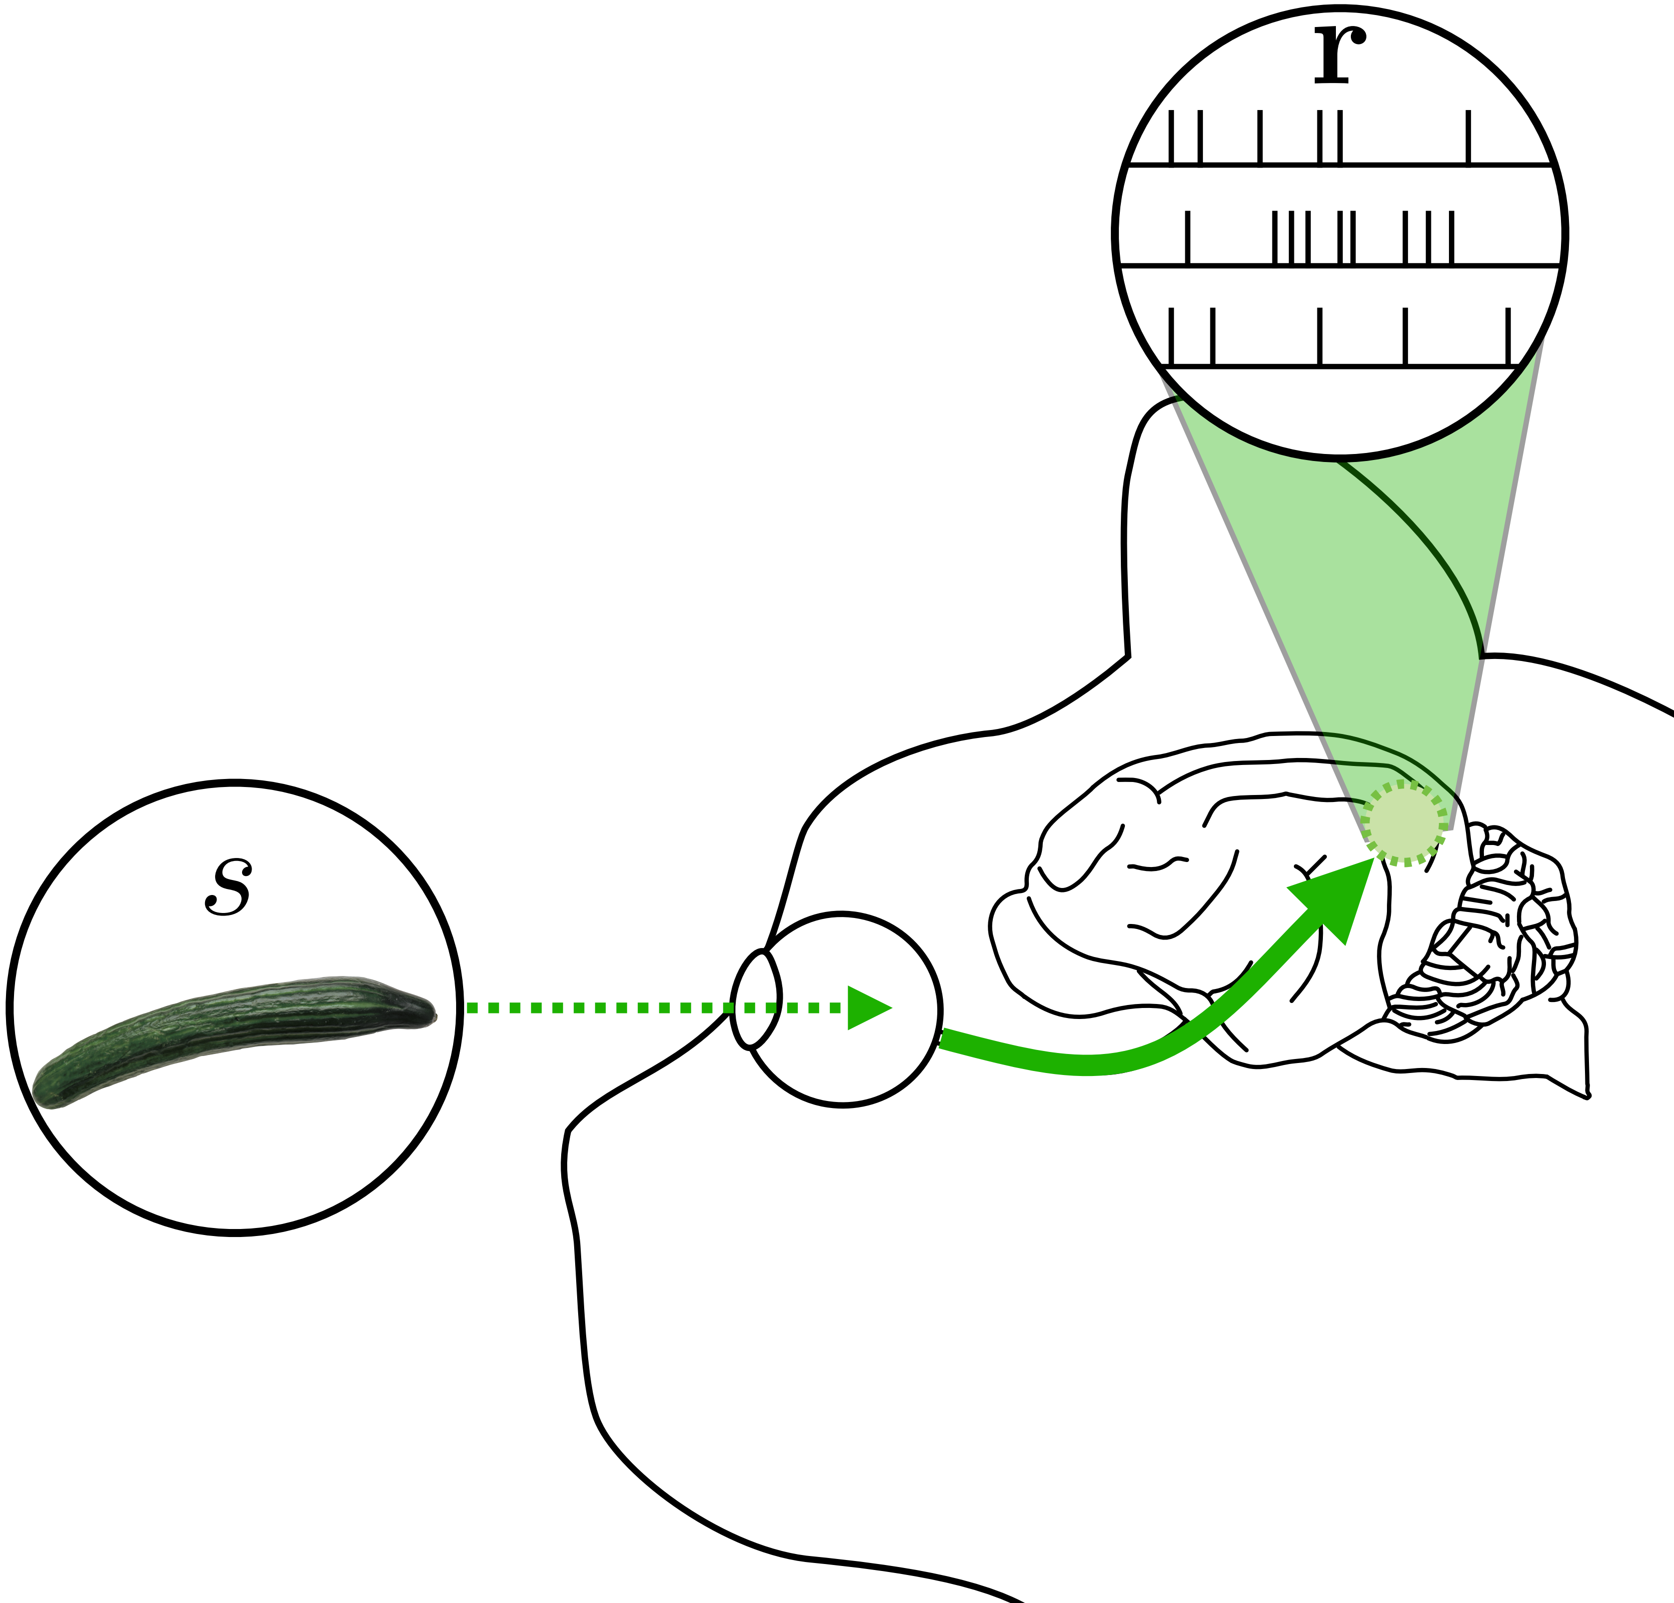

# Flow of information


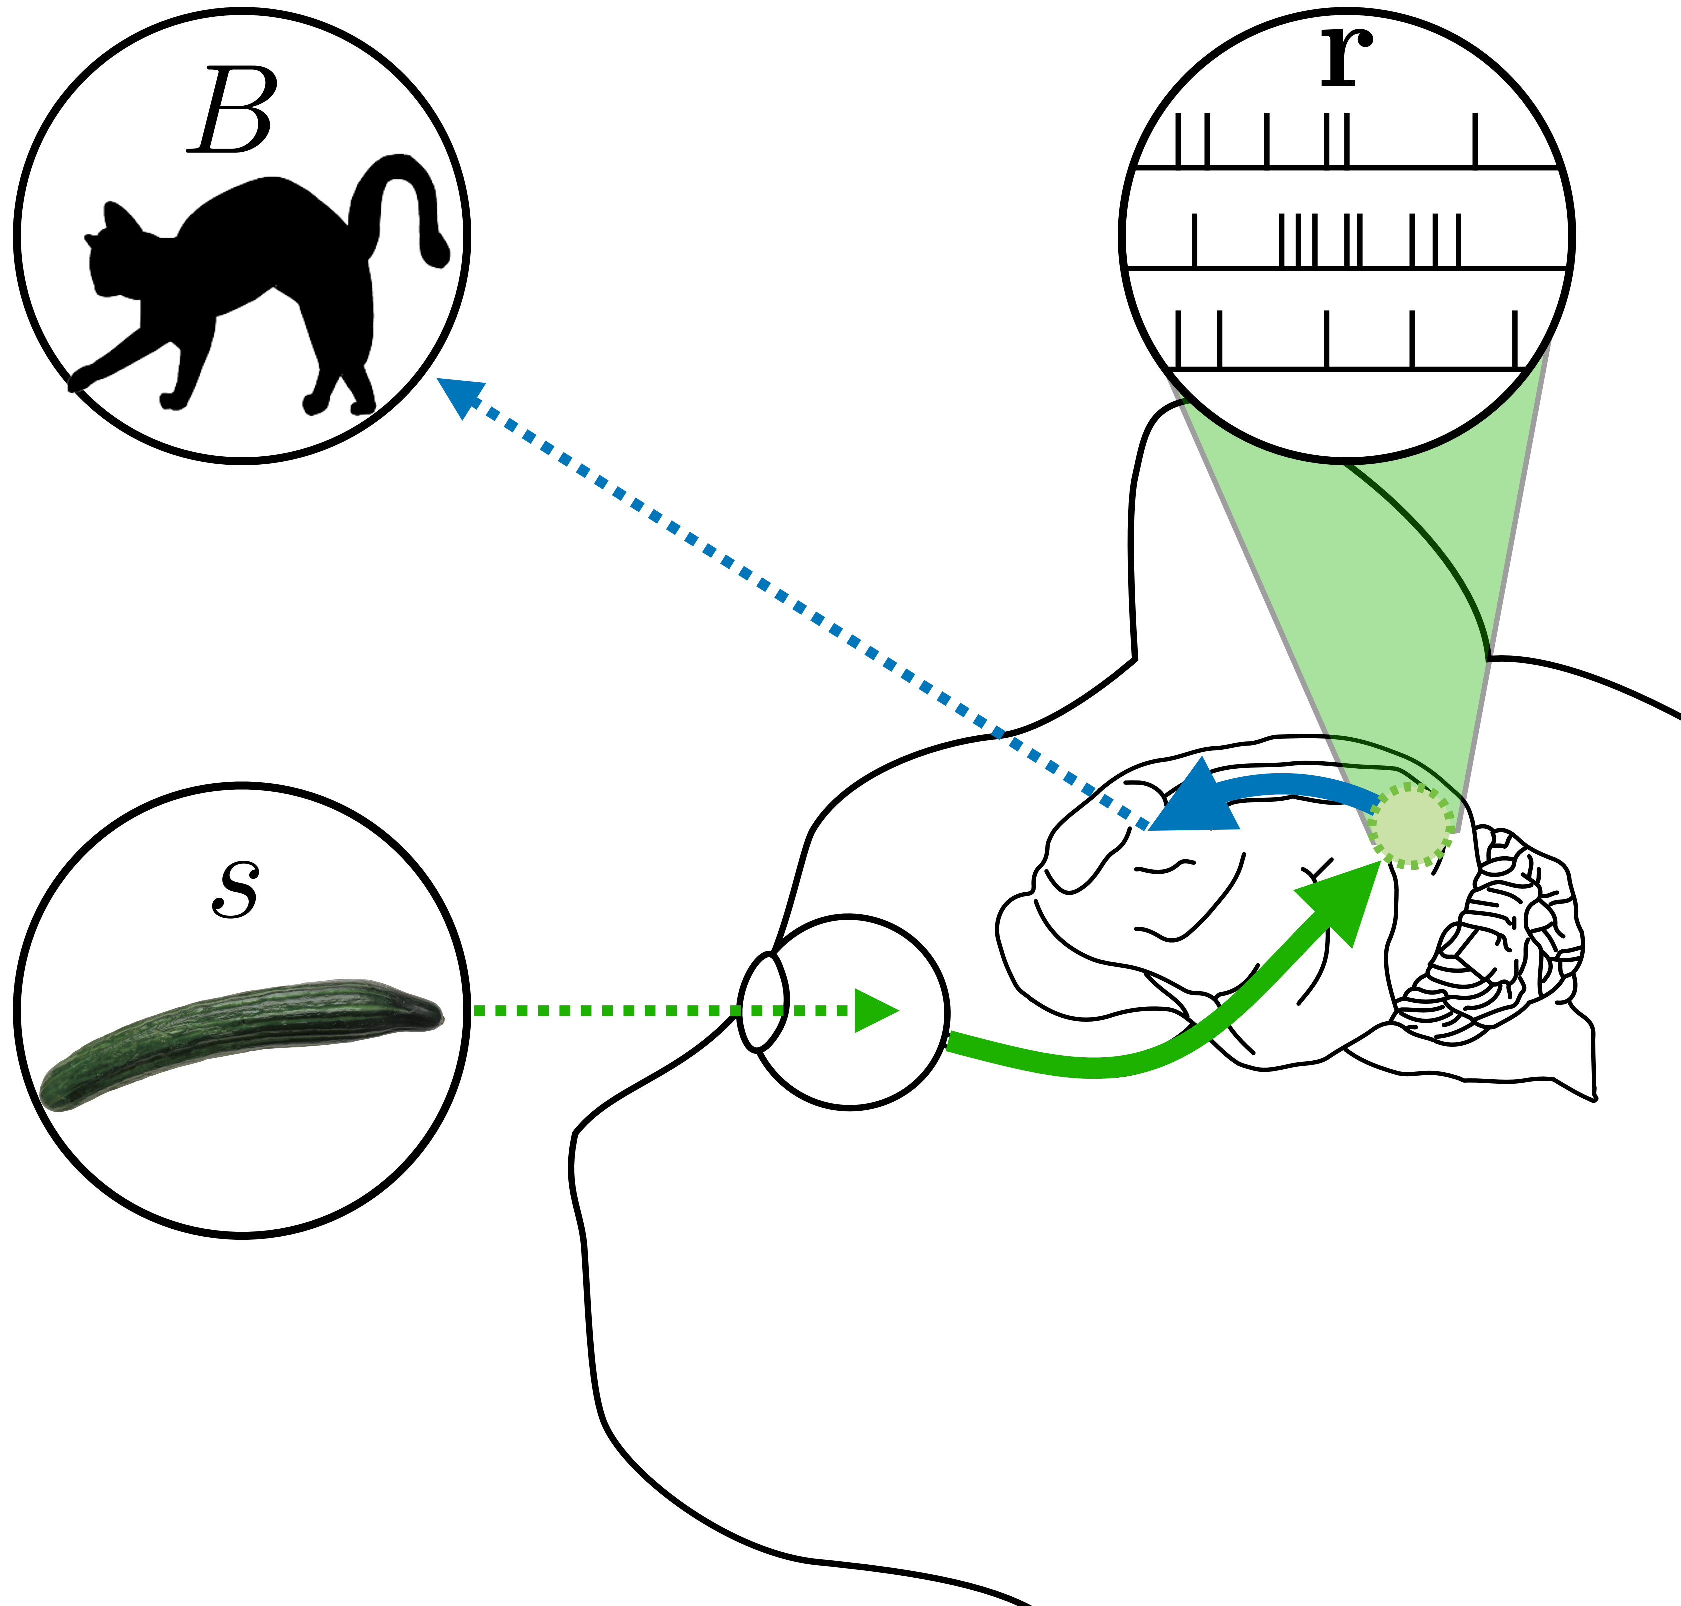

# Flow of information

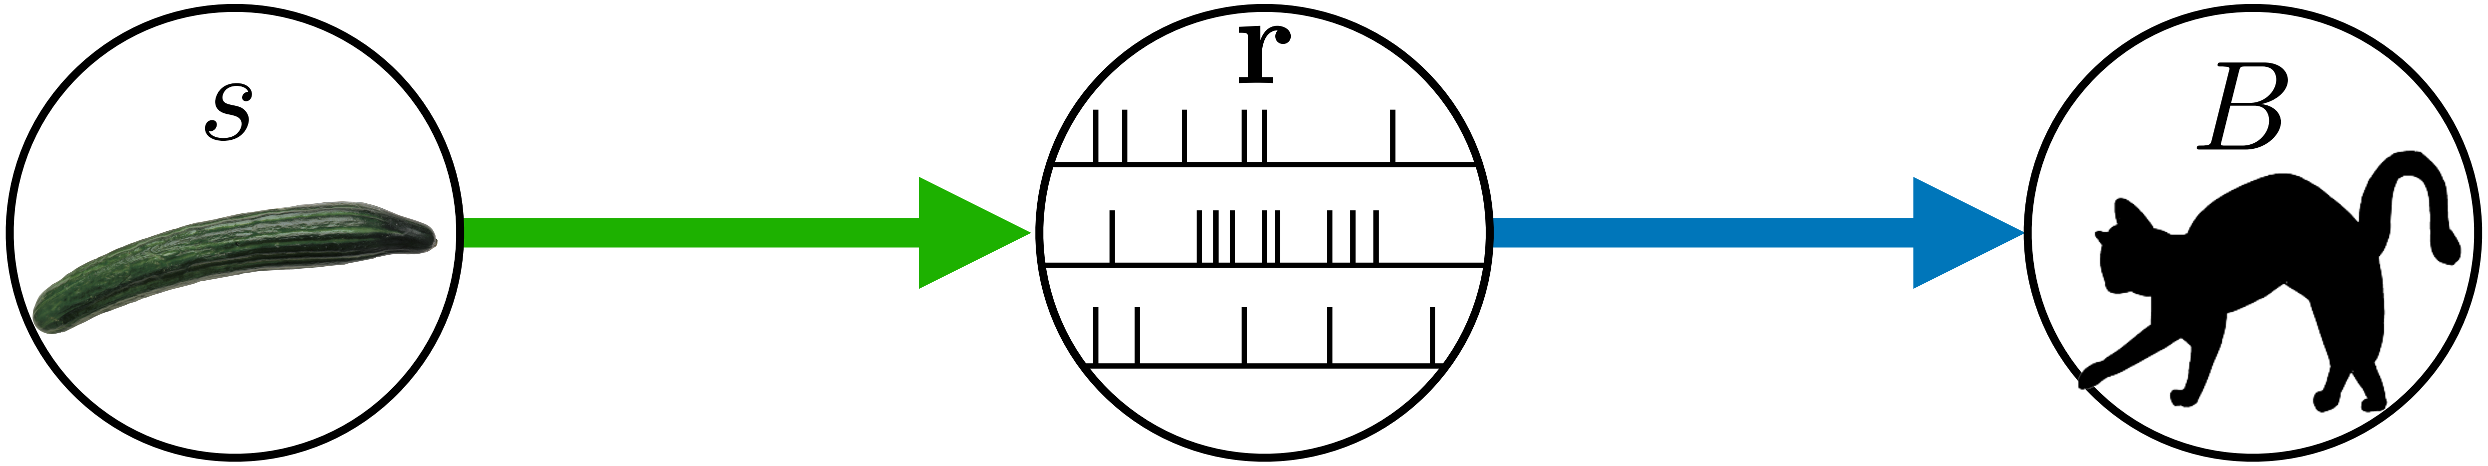

# Studying visual cortex

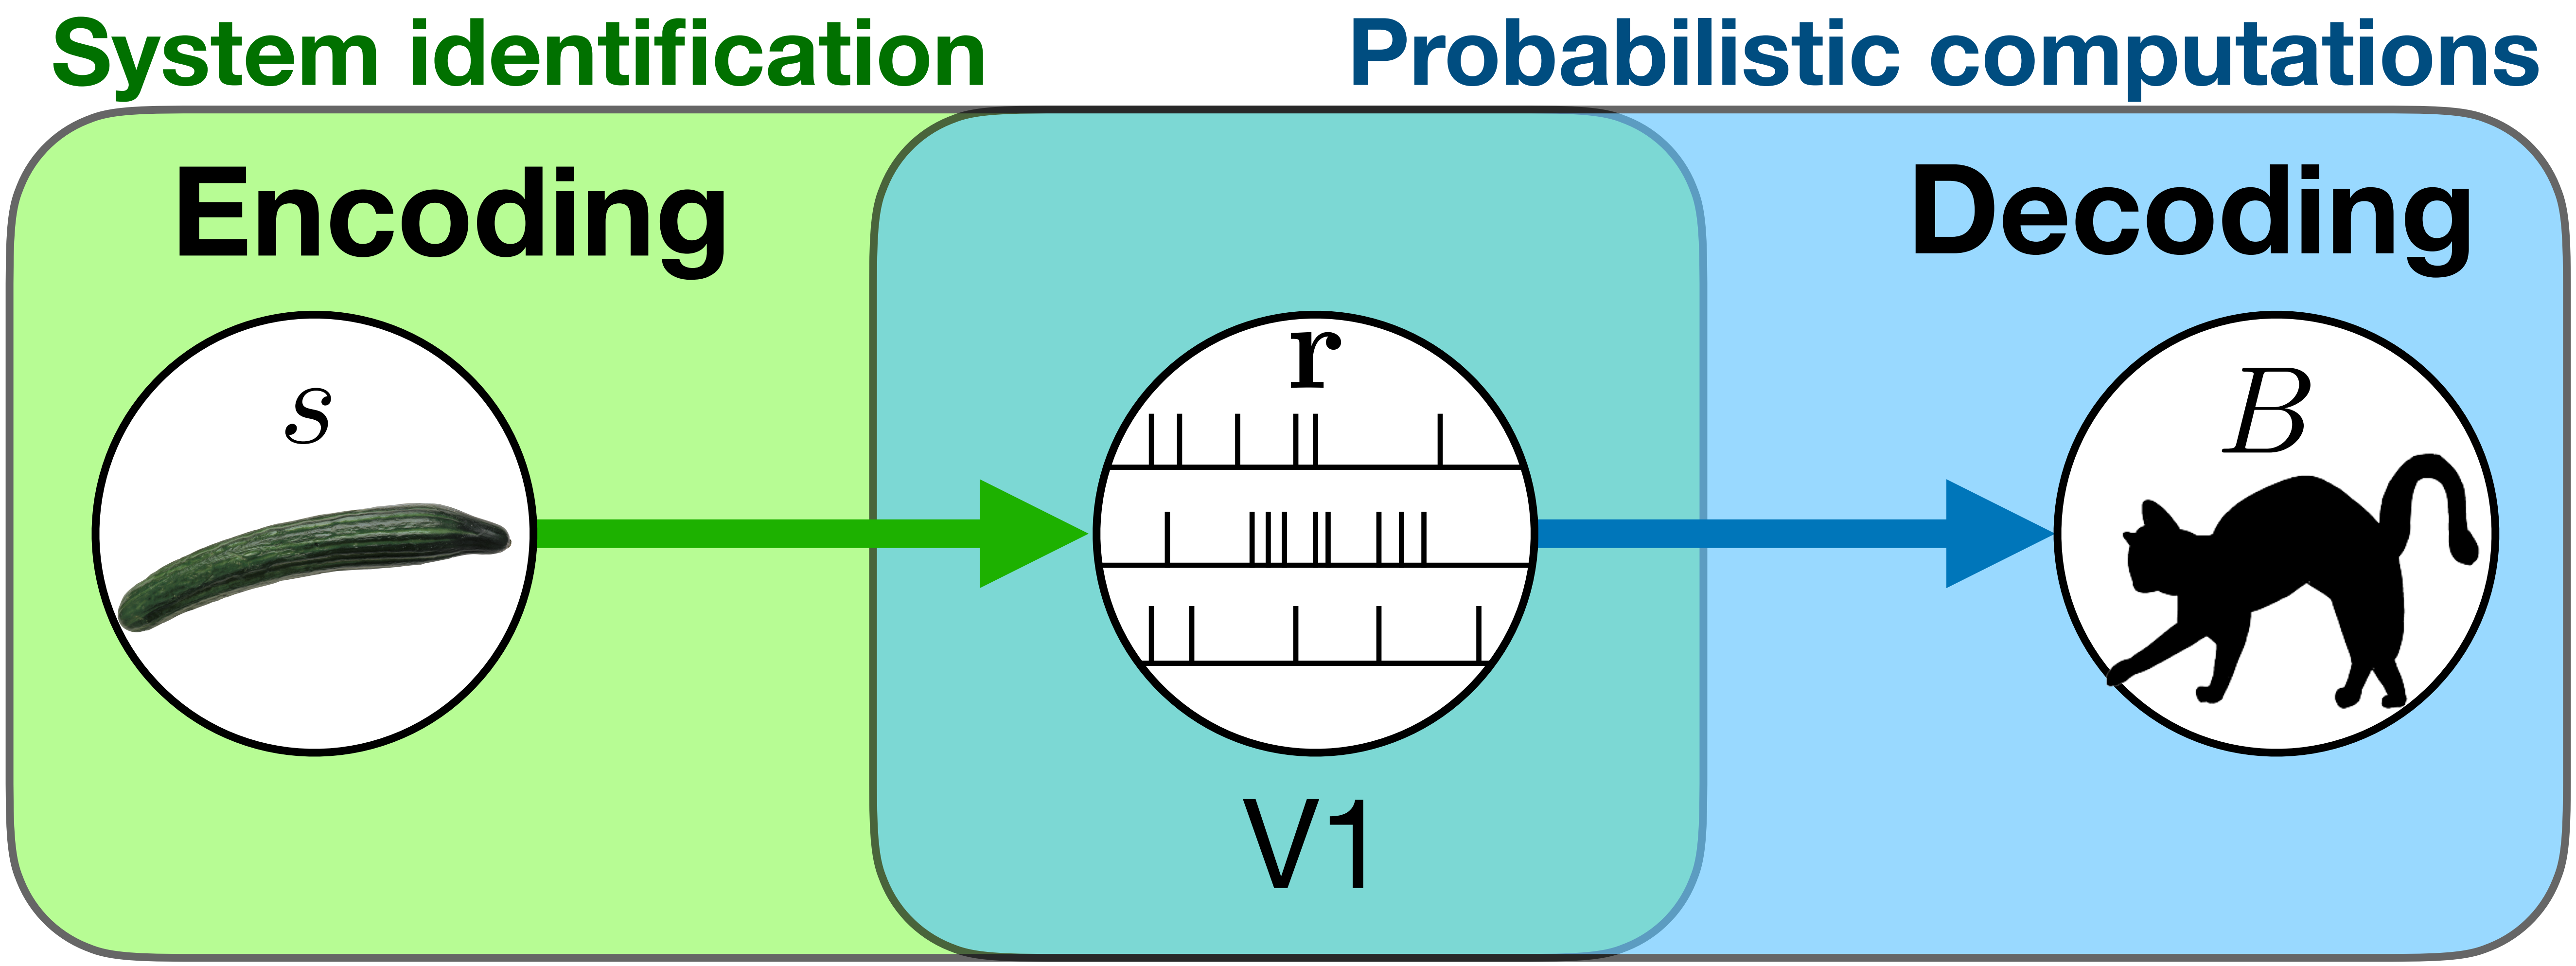

# Studying visual cortex

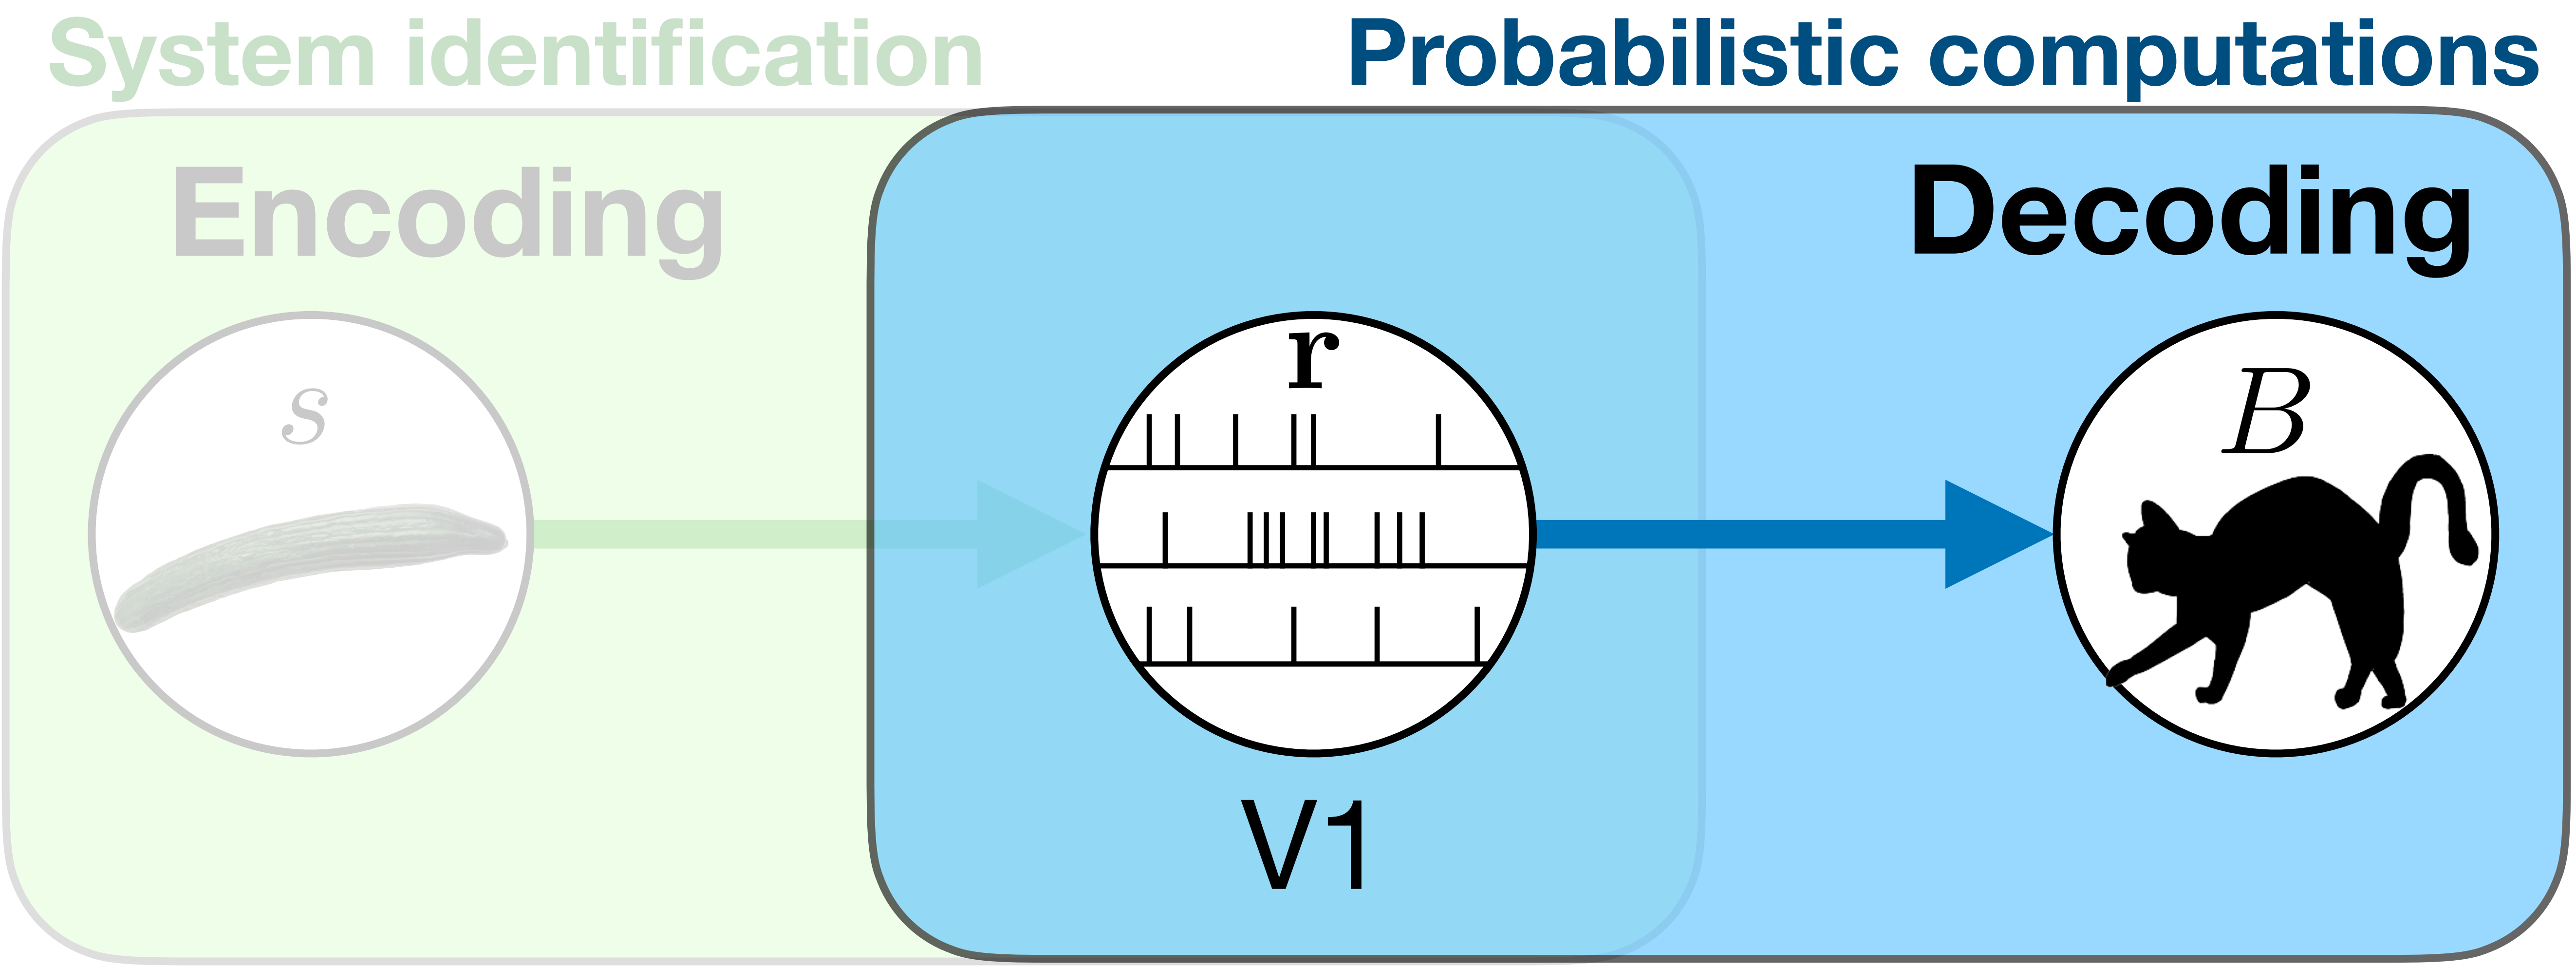


Walker*, Cotton*, Ma & Tolias. (2020) “A neural basis of probabilistic computation in visual cortex” *Nature Neuroscience*

# (There will be slides with technical details) ★

- Some slides will contain extra technical details
- Technical details slides have titles in parenthesis and large **star** on the corner

# Why did the cat decide to jump?


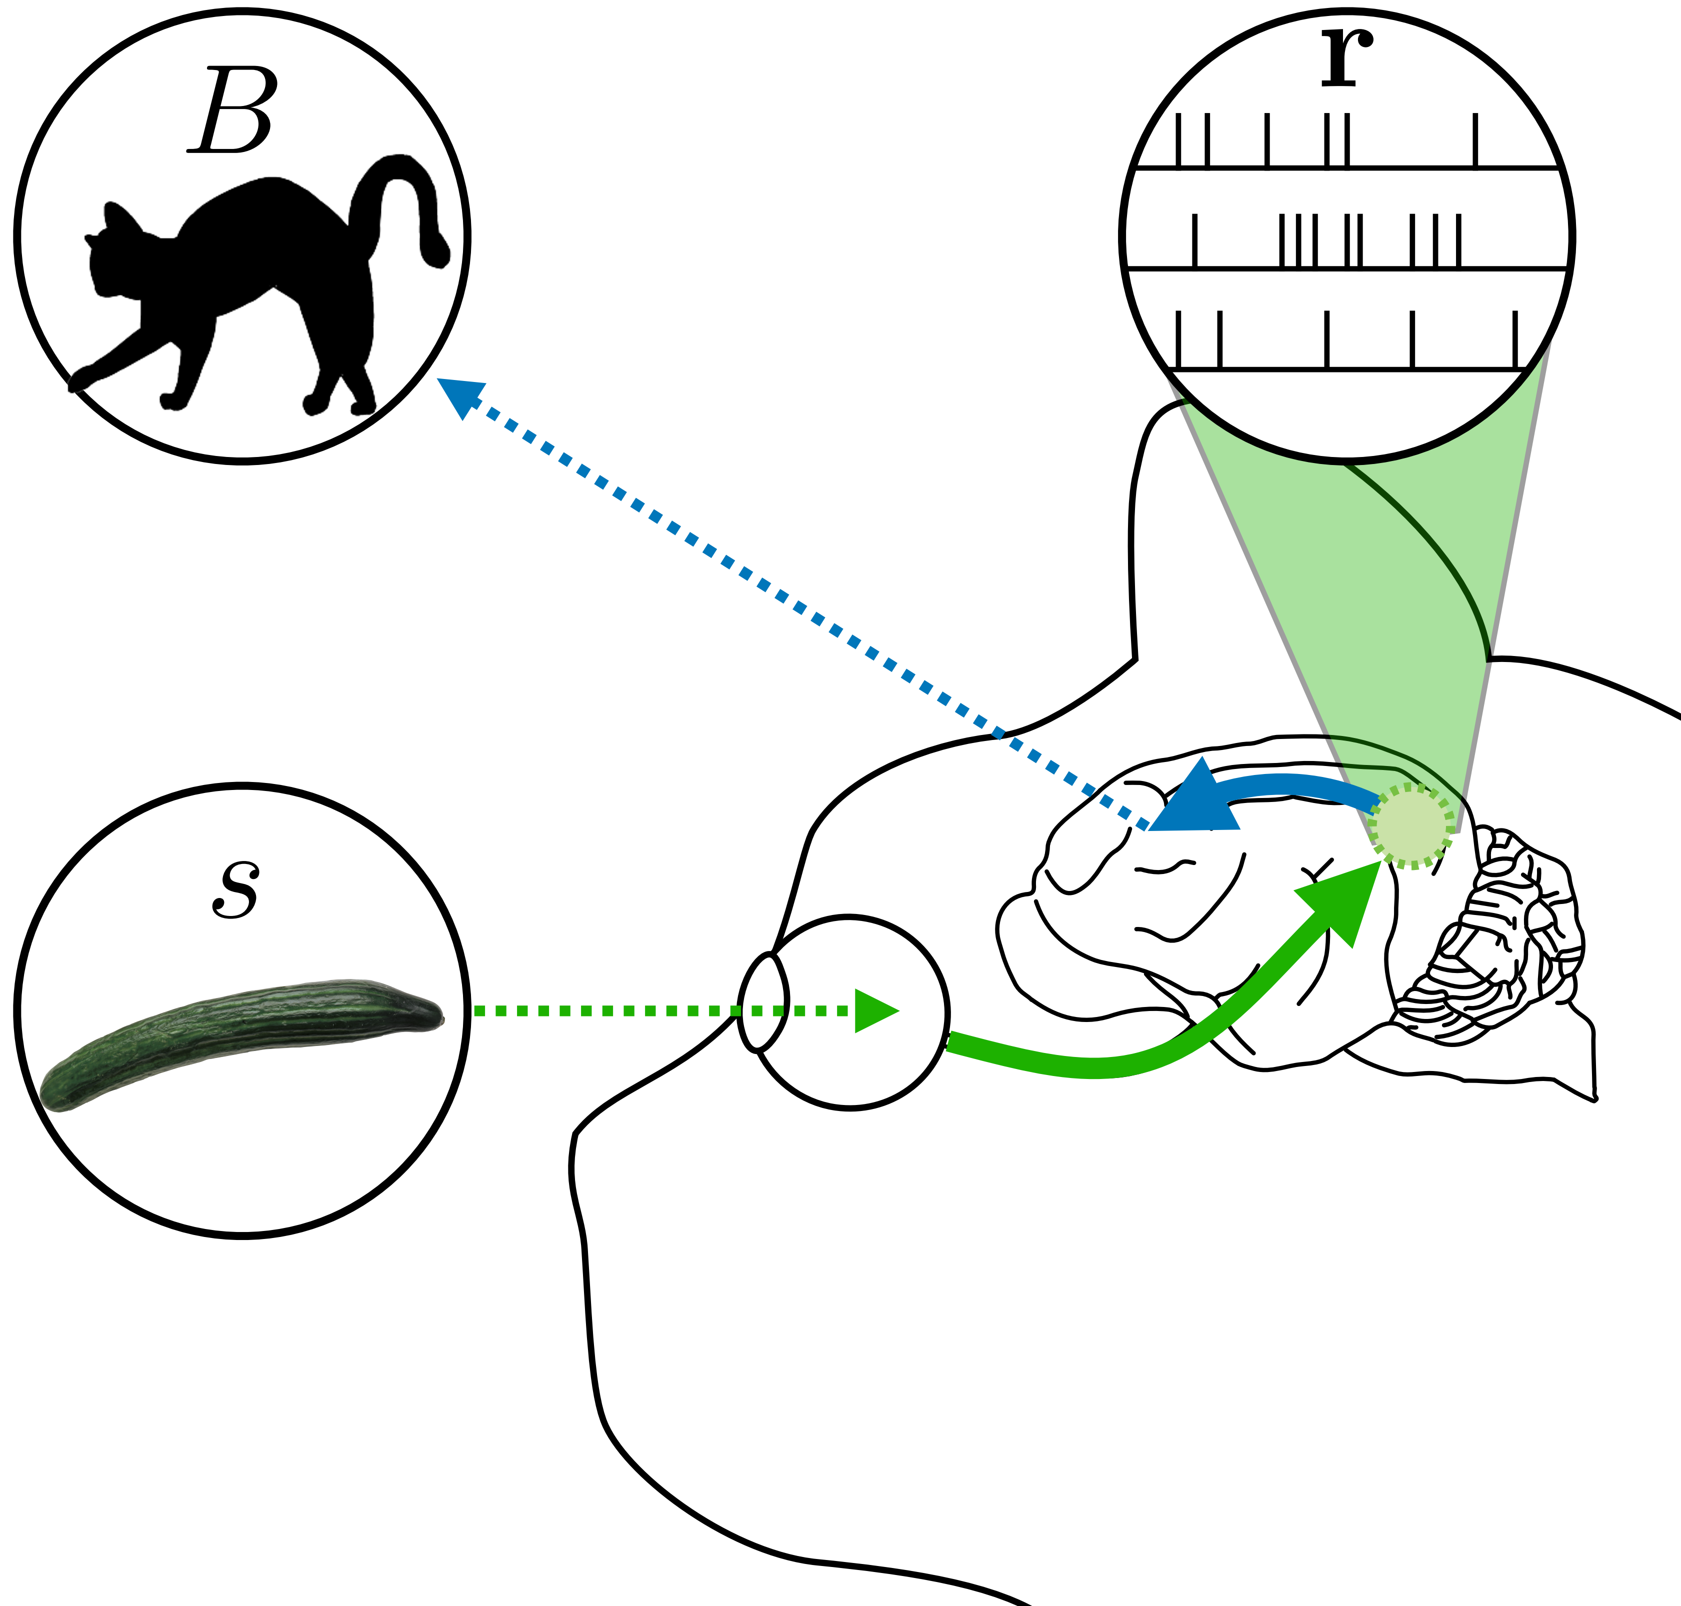

# This is a common behavior

# This is a common behavior

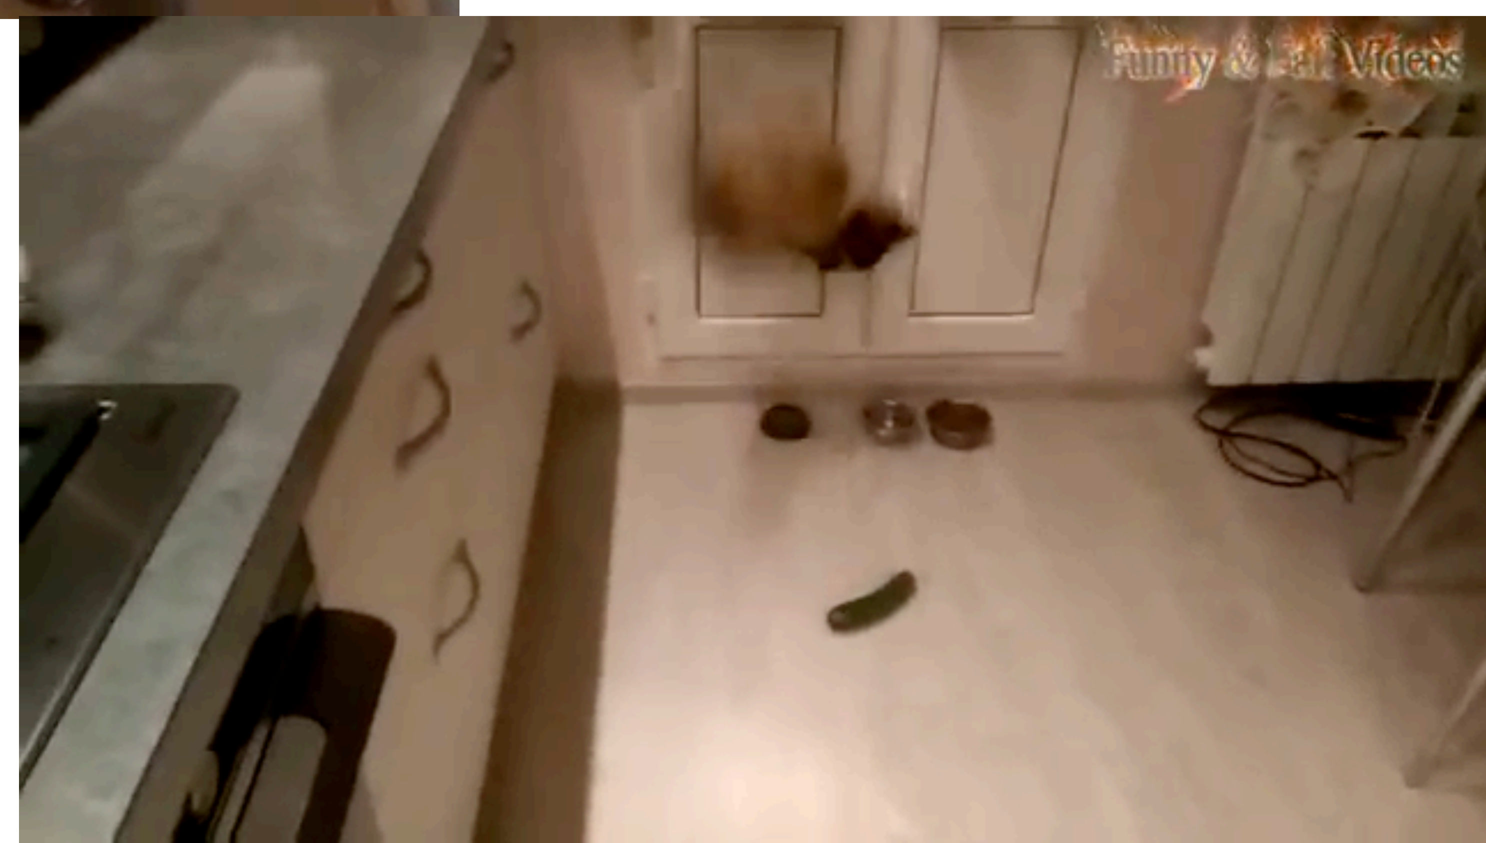

# Cucumber looks like…

# Cucumber looks like…

# …snake?

# Cats do not always jump to cucumber…

# Cats do not always jump to cucumber…

# What does the stimulus look like?

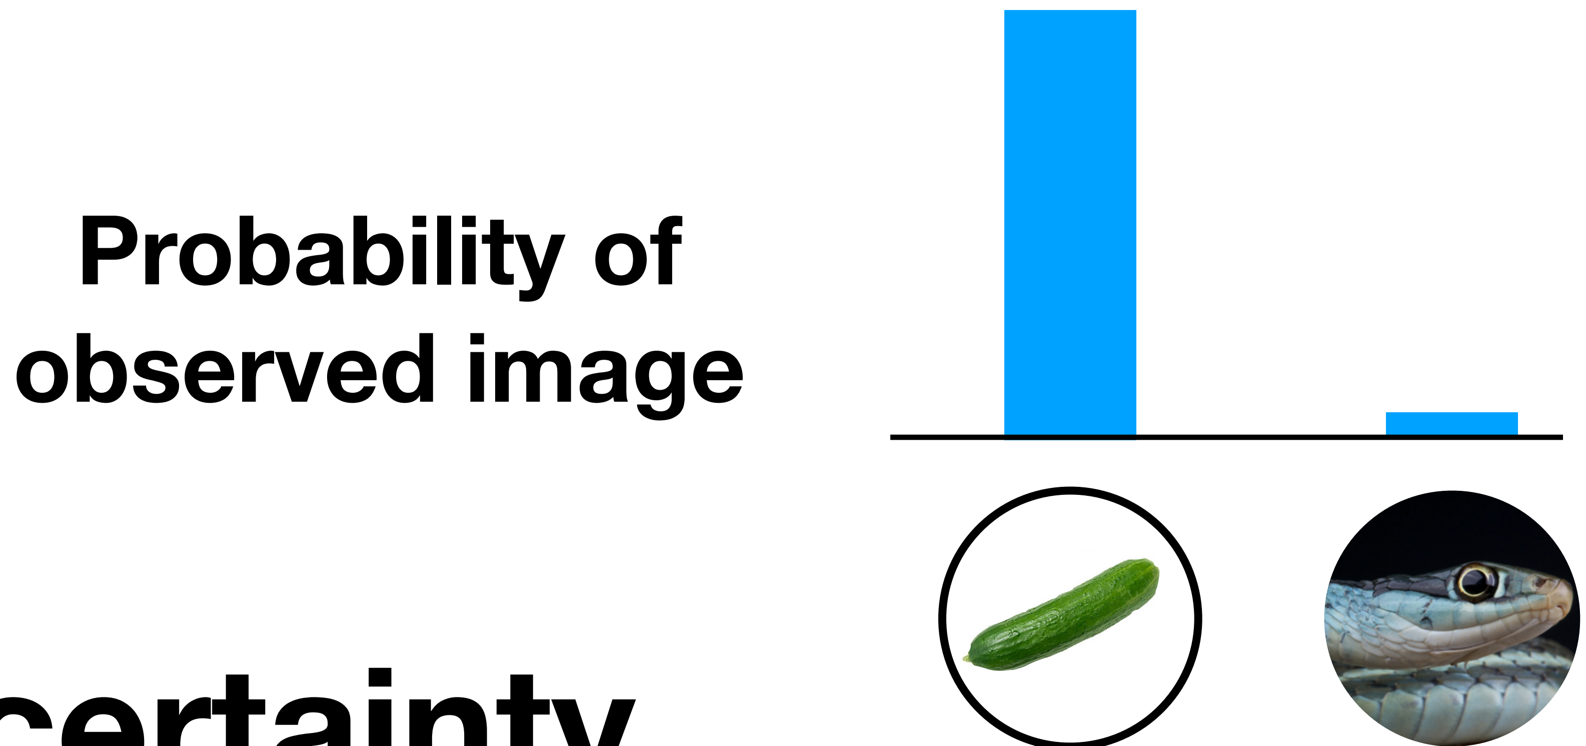


**Low uncertainty**

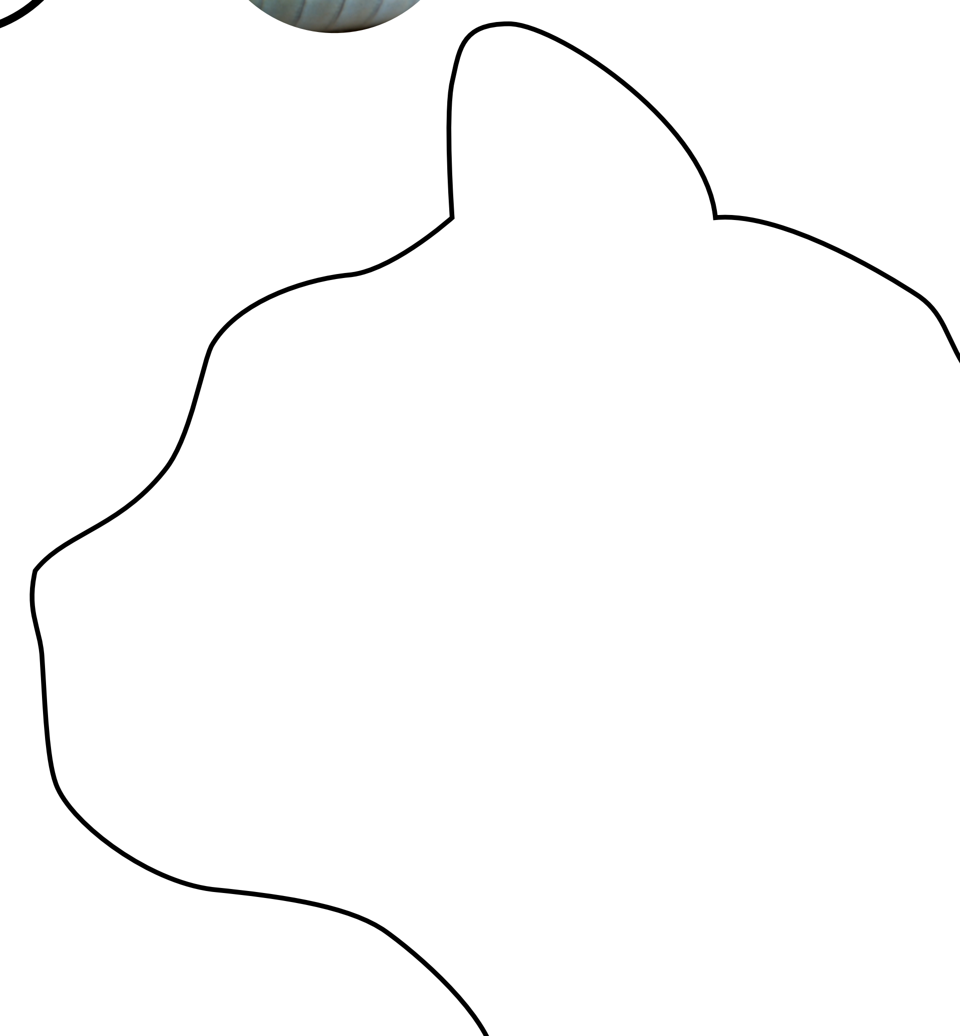

# What does the stimulus look like?

**Probability of observed image**

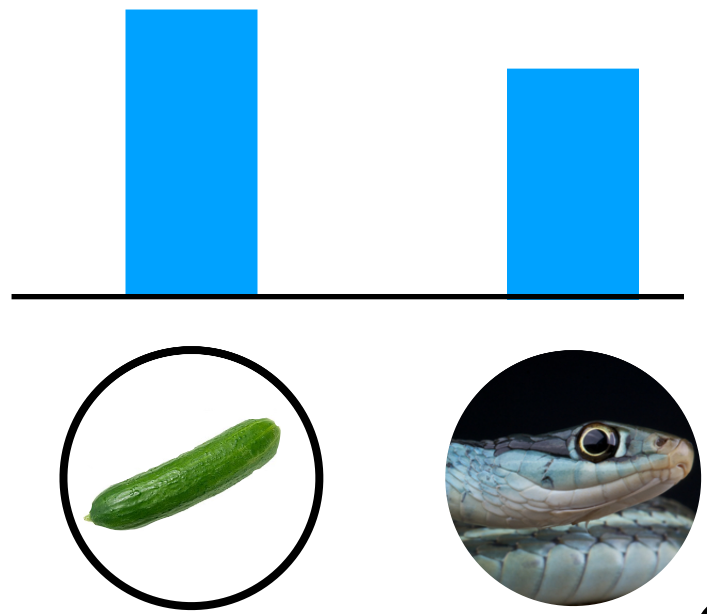

## High uncertainty

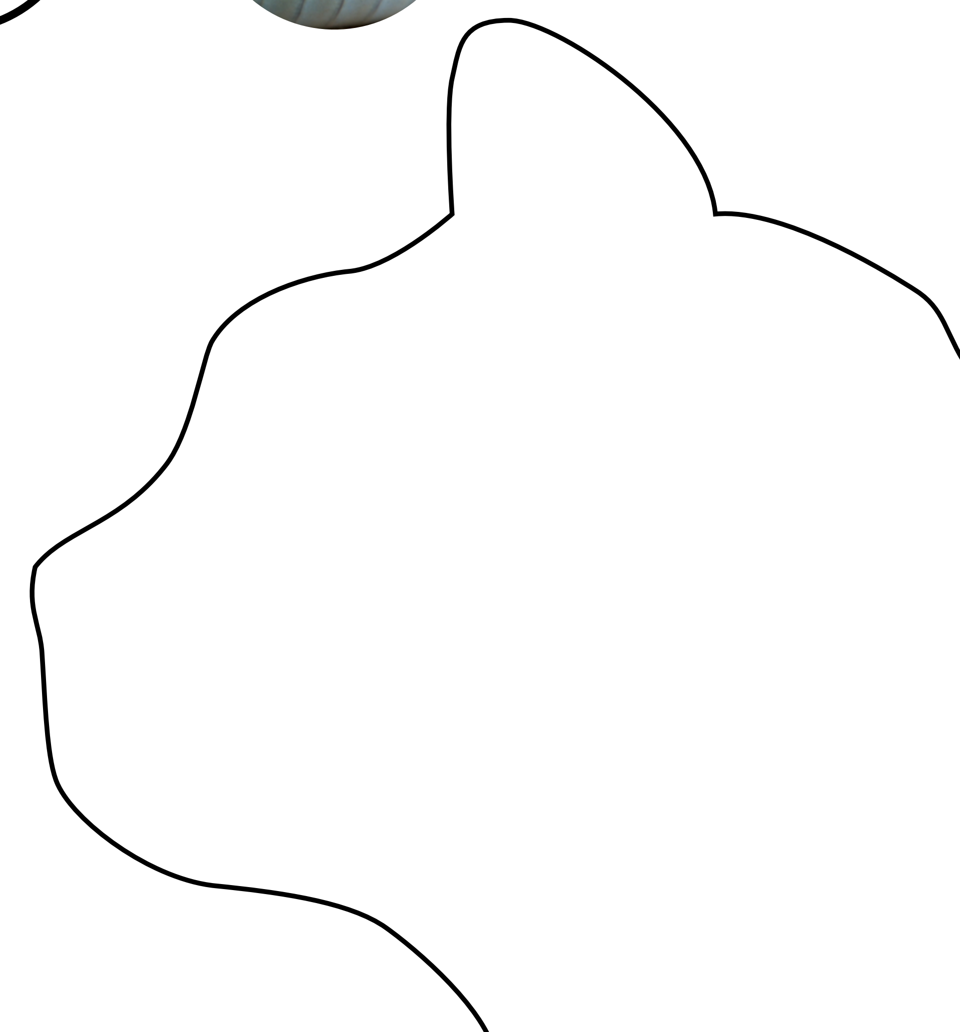

# Uncertainty affects decision!

- In multi-sensory integration, human and monkey combine information about the stimulus weighing each evidence by their reliability (e.g. Ernst and Banks, 2002 *Nature*)
- Uncertainty information is utilized on a trial-by-trial basis
- **But how is uncertainty represented by a population of neurons?**

# Identical stimulus presentations result in variable spike counts

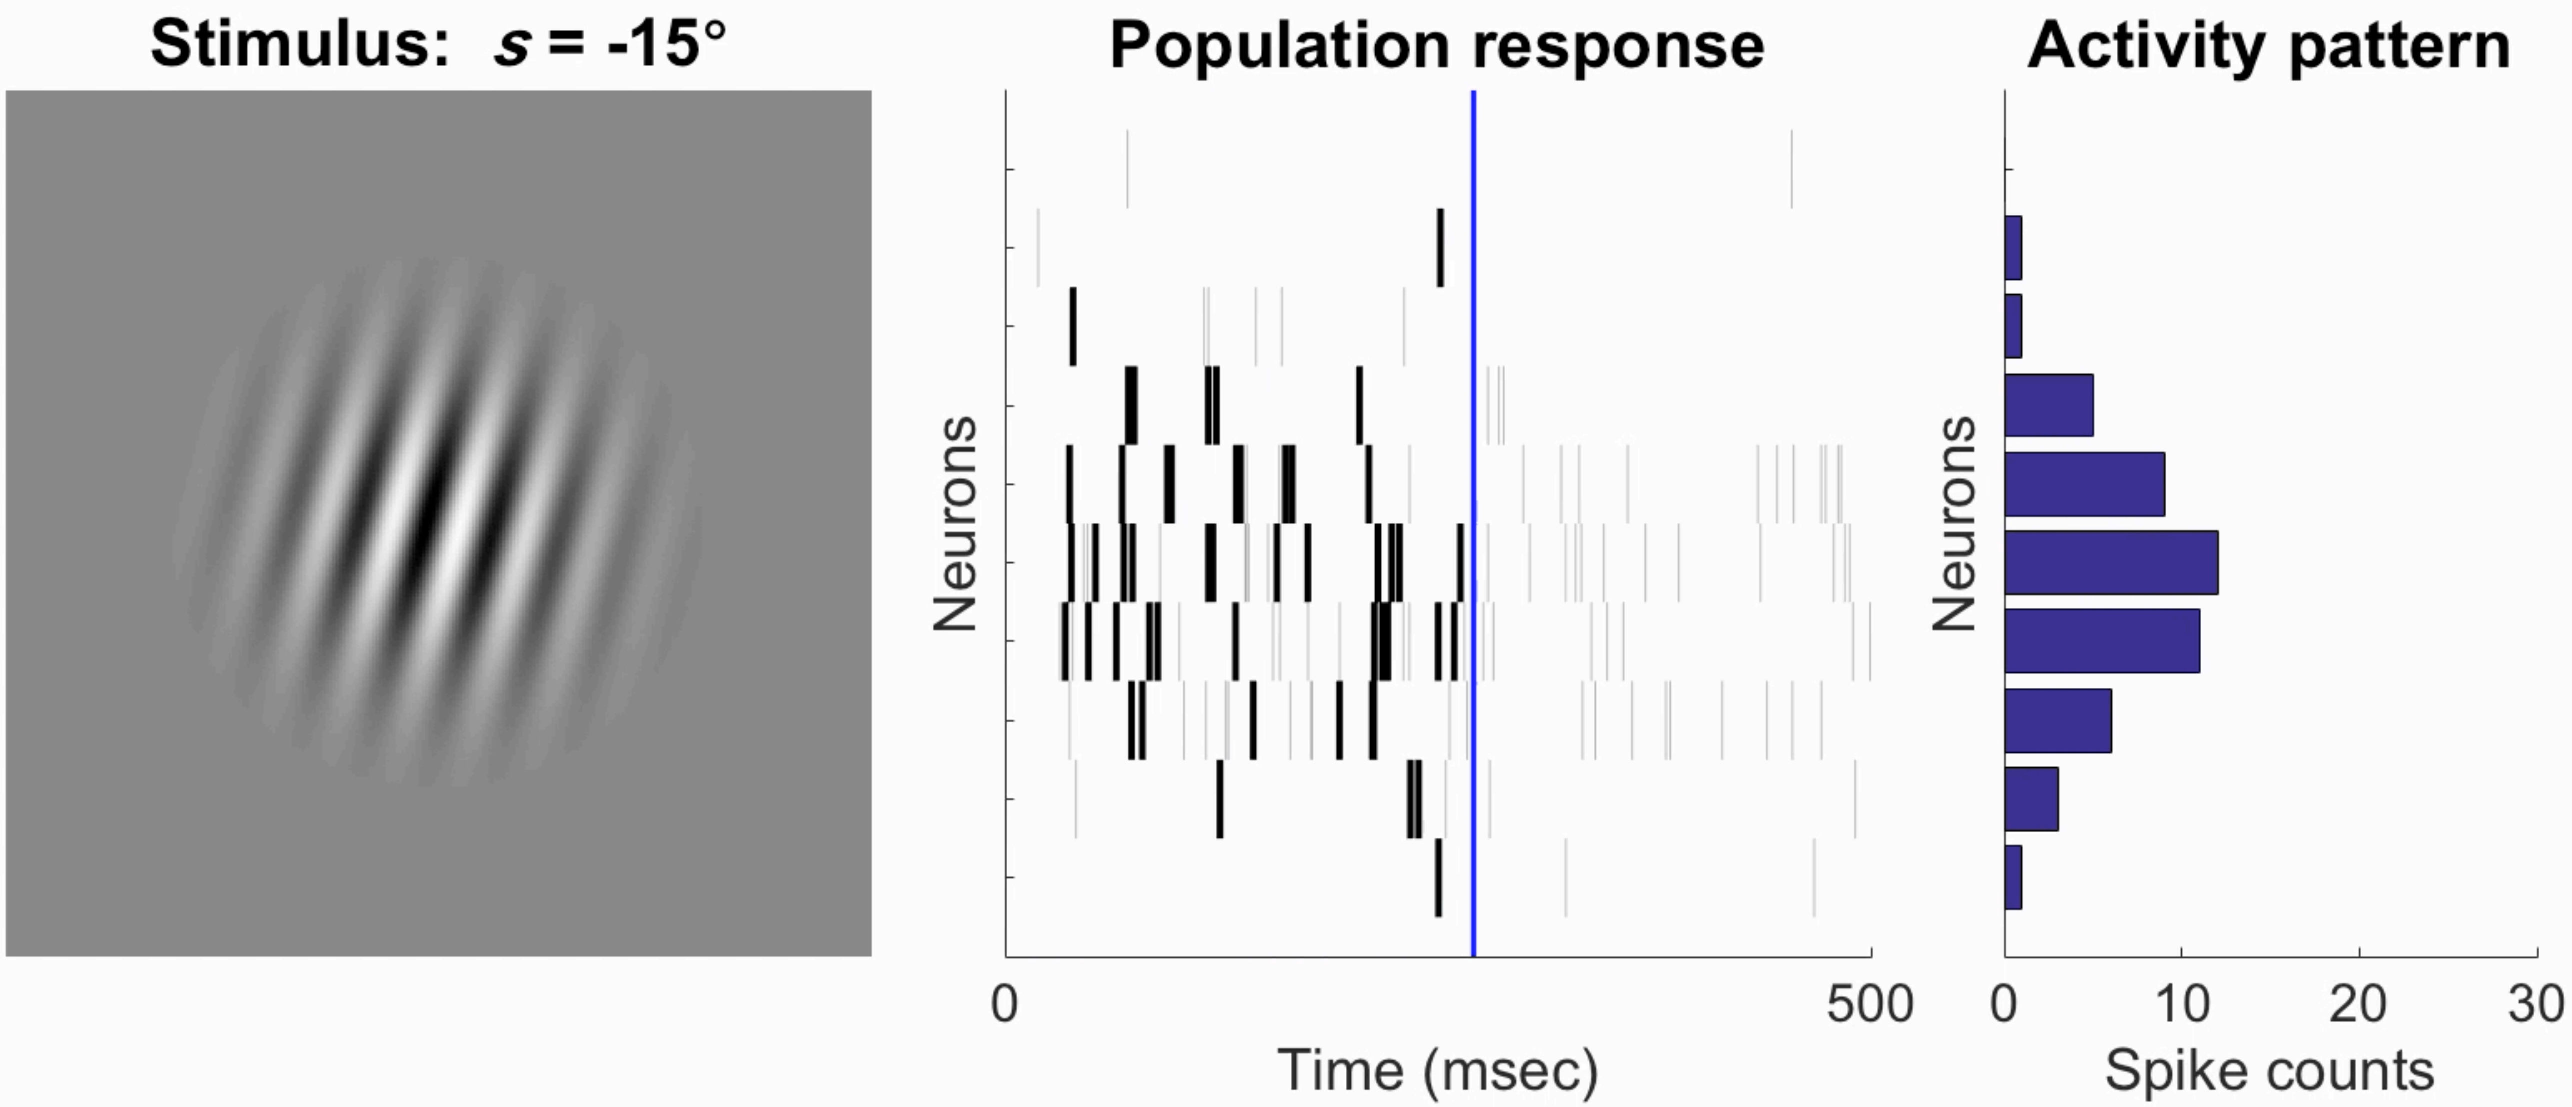


any particular spike count pattern may be consistent with multiple stimulus values!

# Identical stimulus presentations result in variable spike counts

High contrast

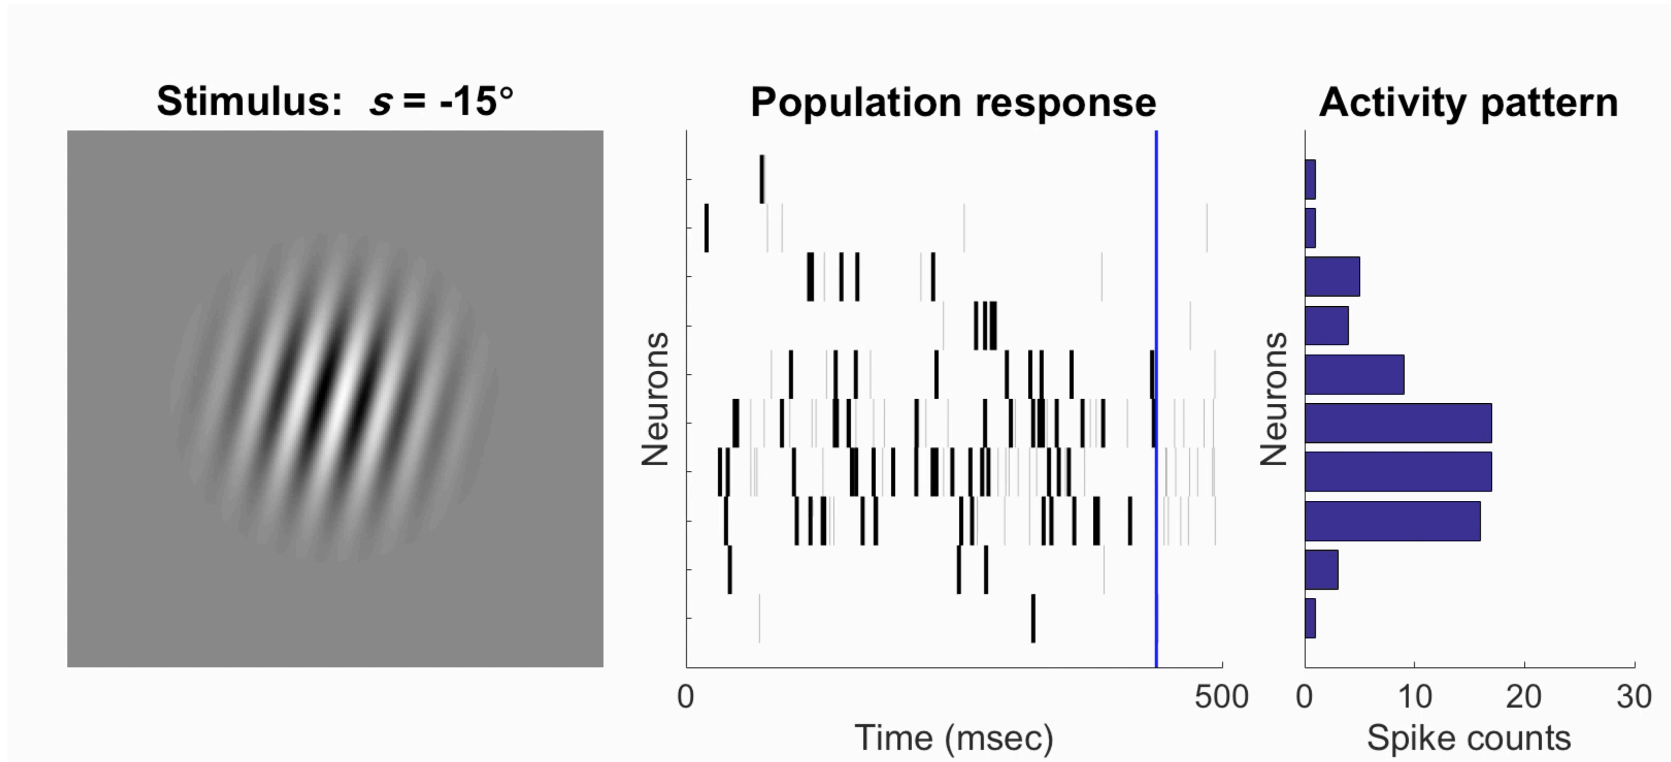


Low contrast

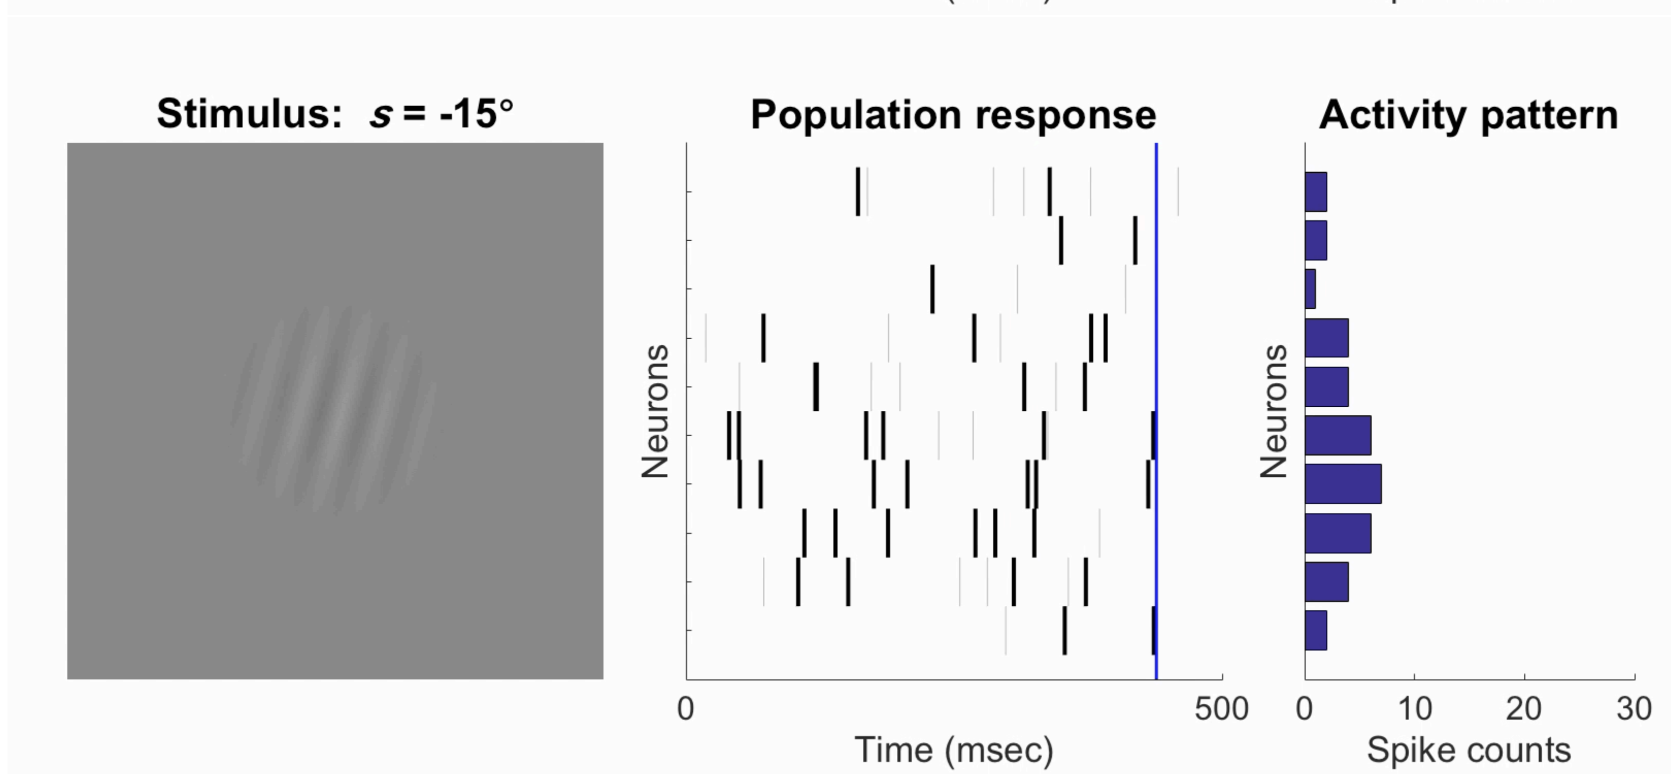

# Likelihood function captures uncertainty

Likelihood function = P(activity pattern | hypothesized orientation)

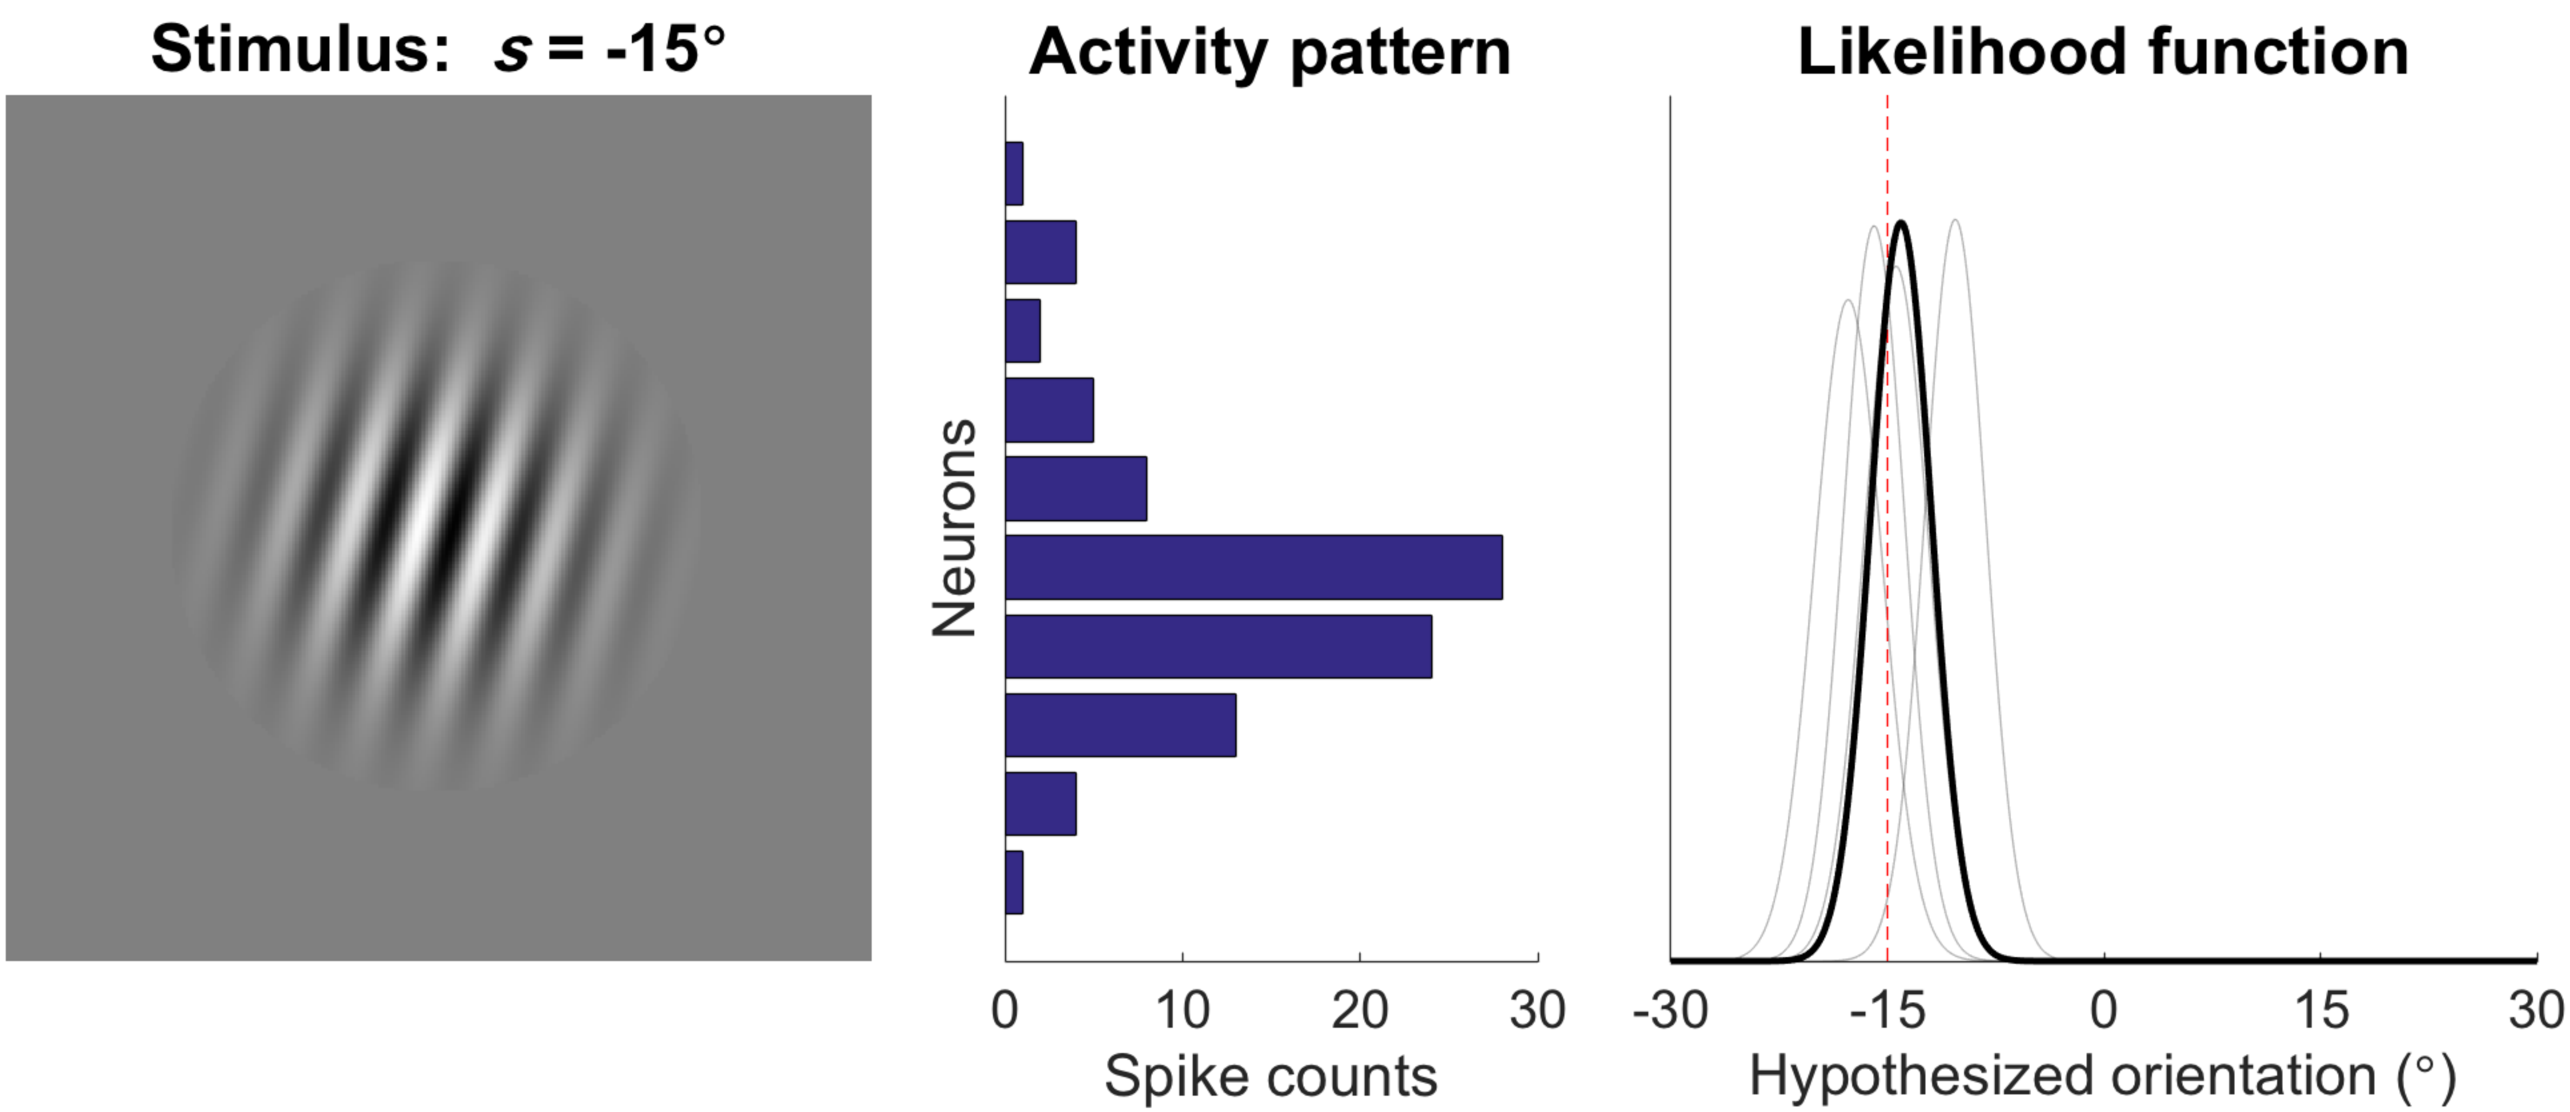


Pouget et al. 2000; Ma et al. 2006

# Likelihood function captures uncertainty

High contrast

Low contrast

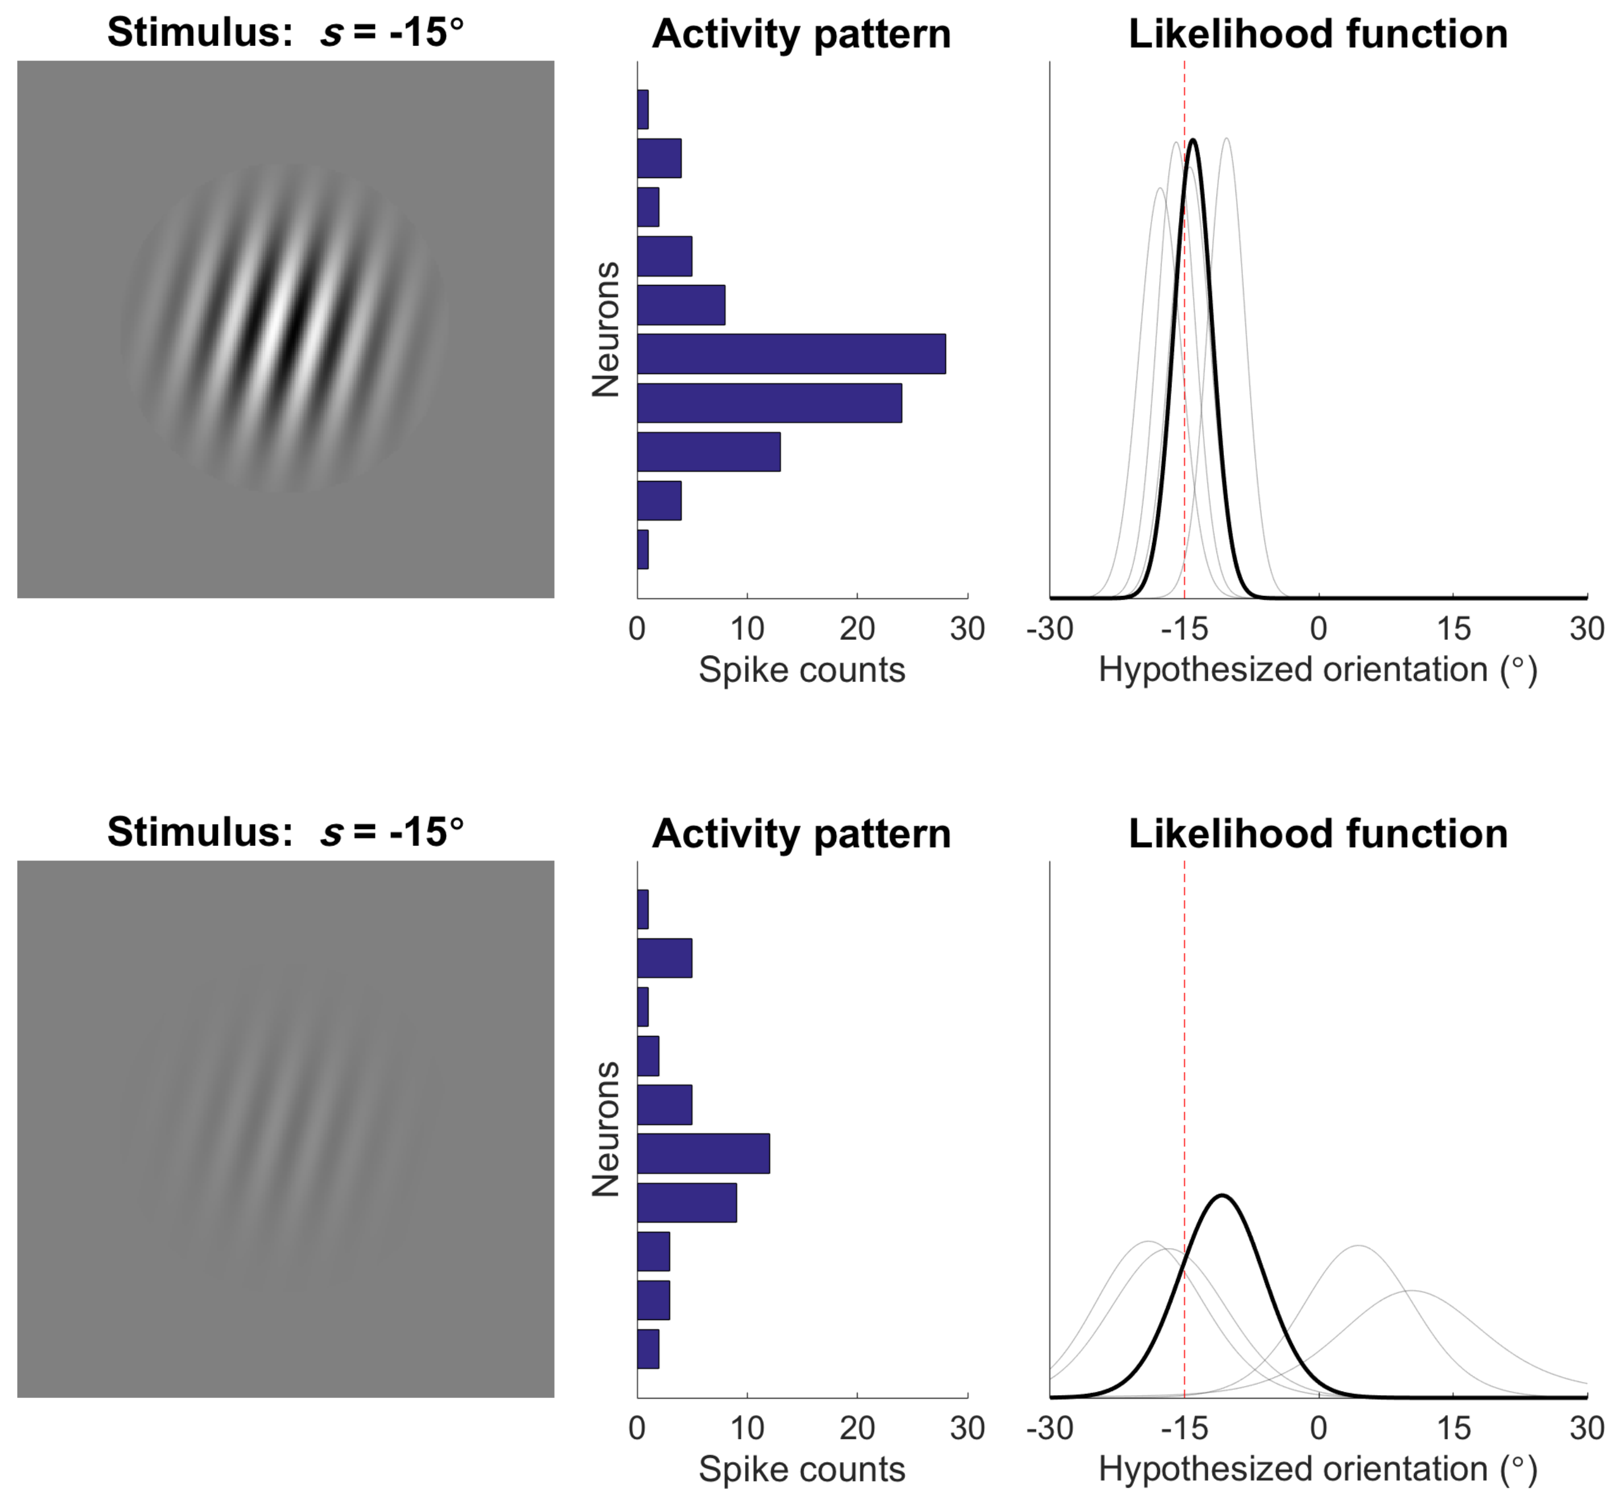

# Likelihood function captures uncertainty

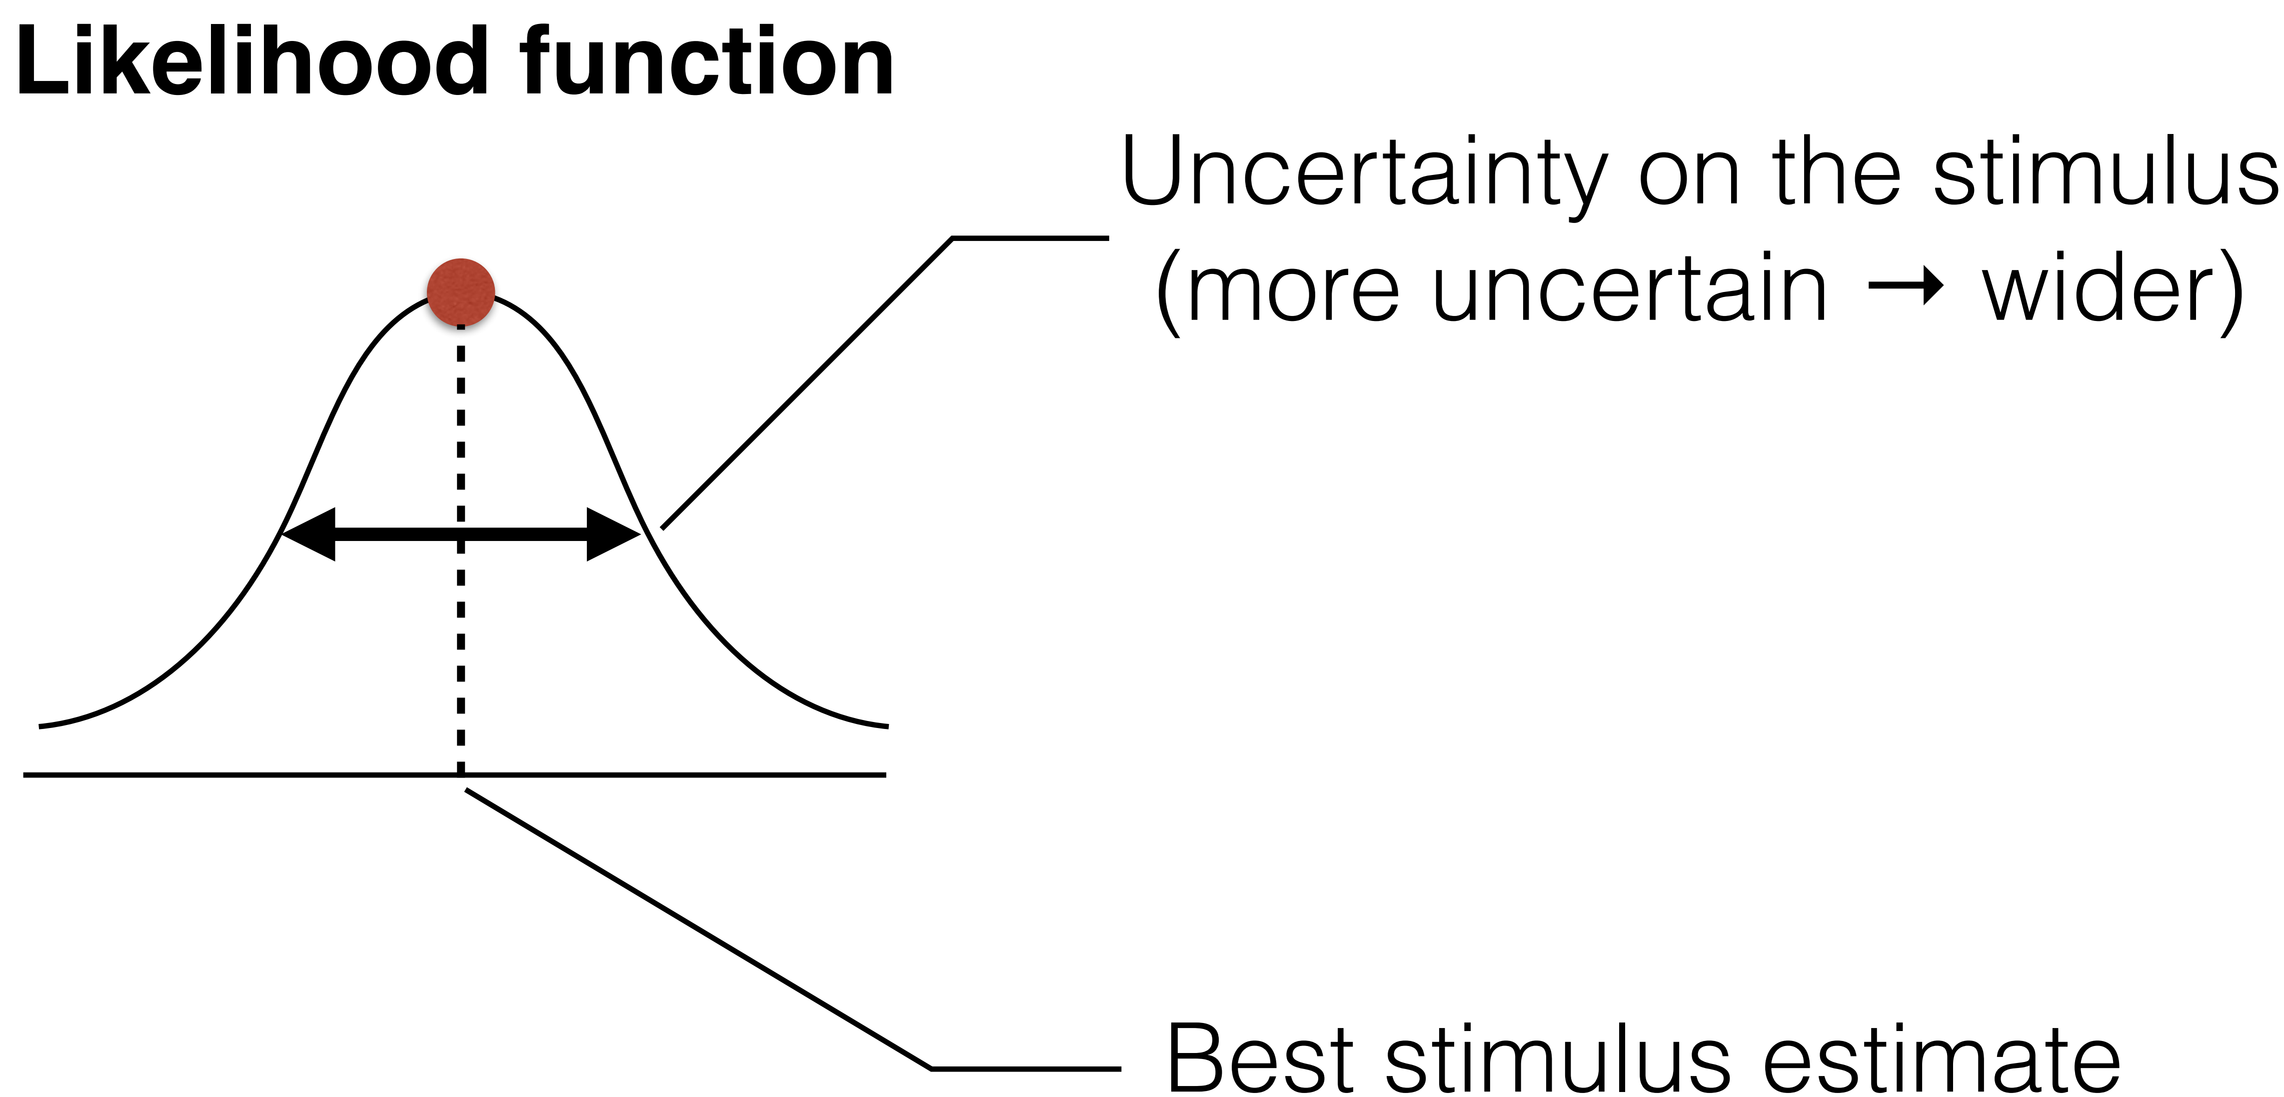


Pouget et al. 2000; Ma et al. 2006

# Probabilistic Population Code (PPC)

Populations of sensory neurons represents the best stimulus estimate and the associated uncertainty simultaneously by **representing likelihood functions**

Ma et al. 2006 *Nature Neuro*

# Hypothesis under Probabilistic Population Code

Ma et al. 2006 *Nature Neuro*

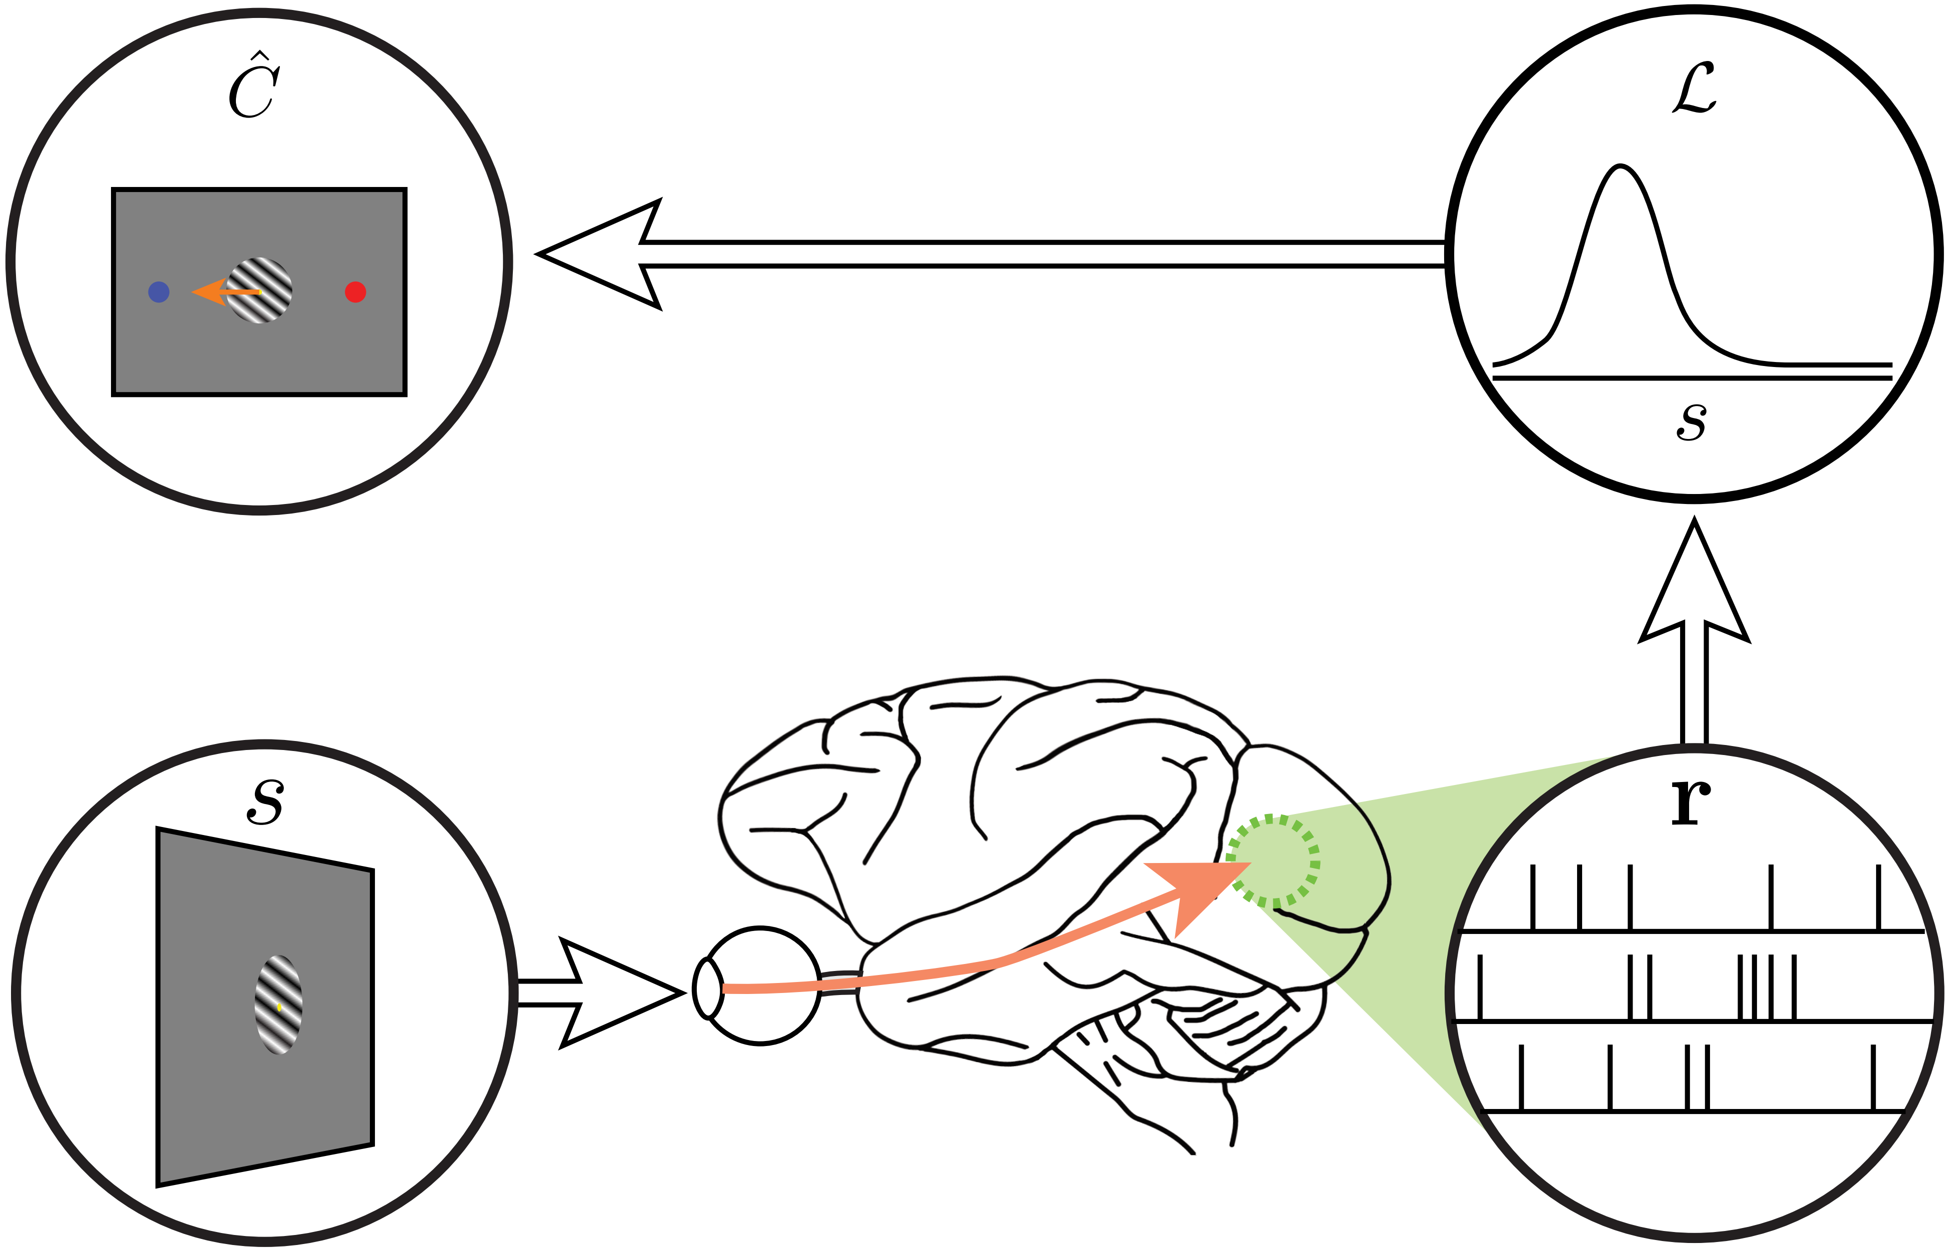


***Best stimulus estimate and uncertainty is simultaneously represented by V1 population as a likelihood function***

# Alternative hypothesis to PPC

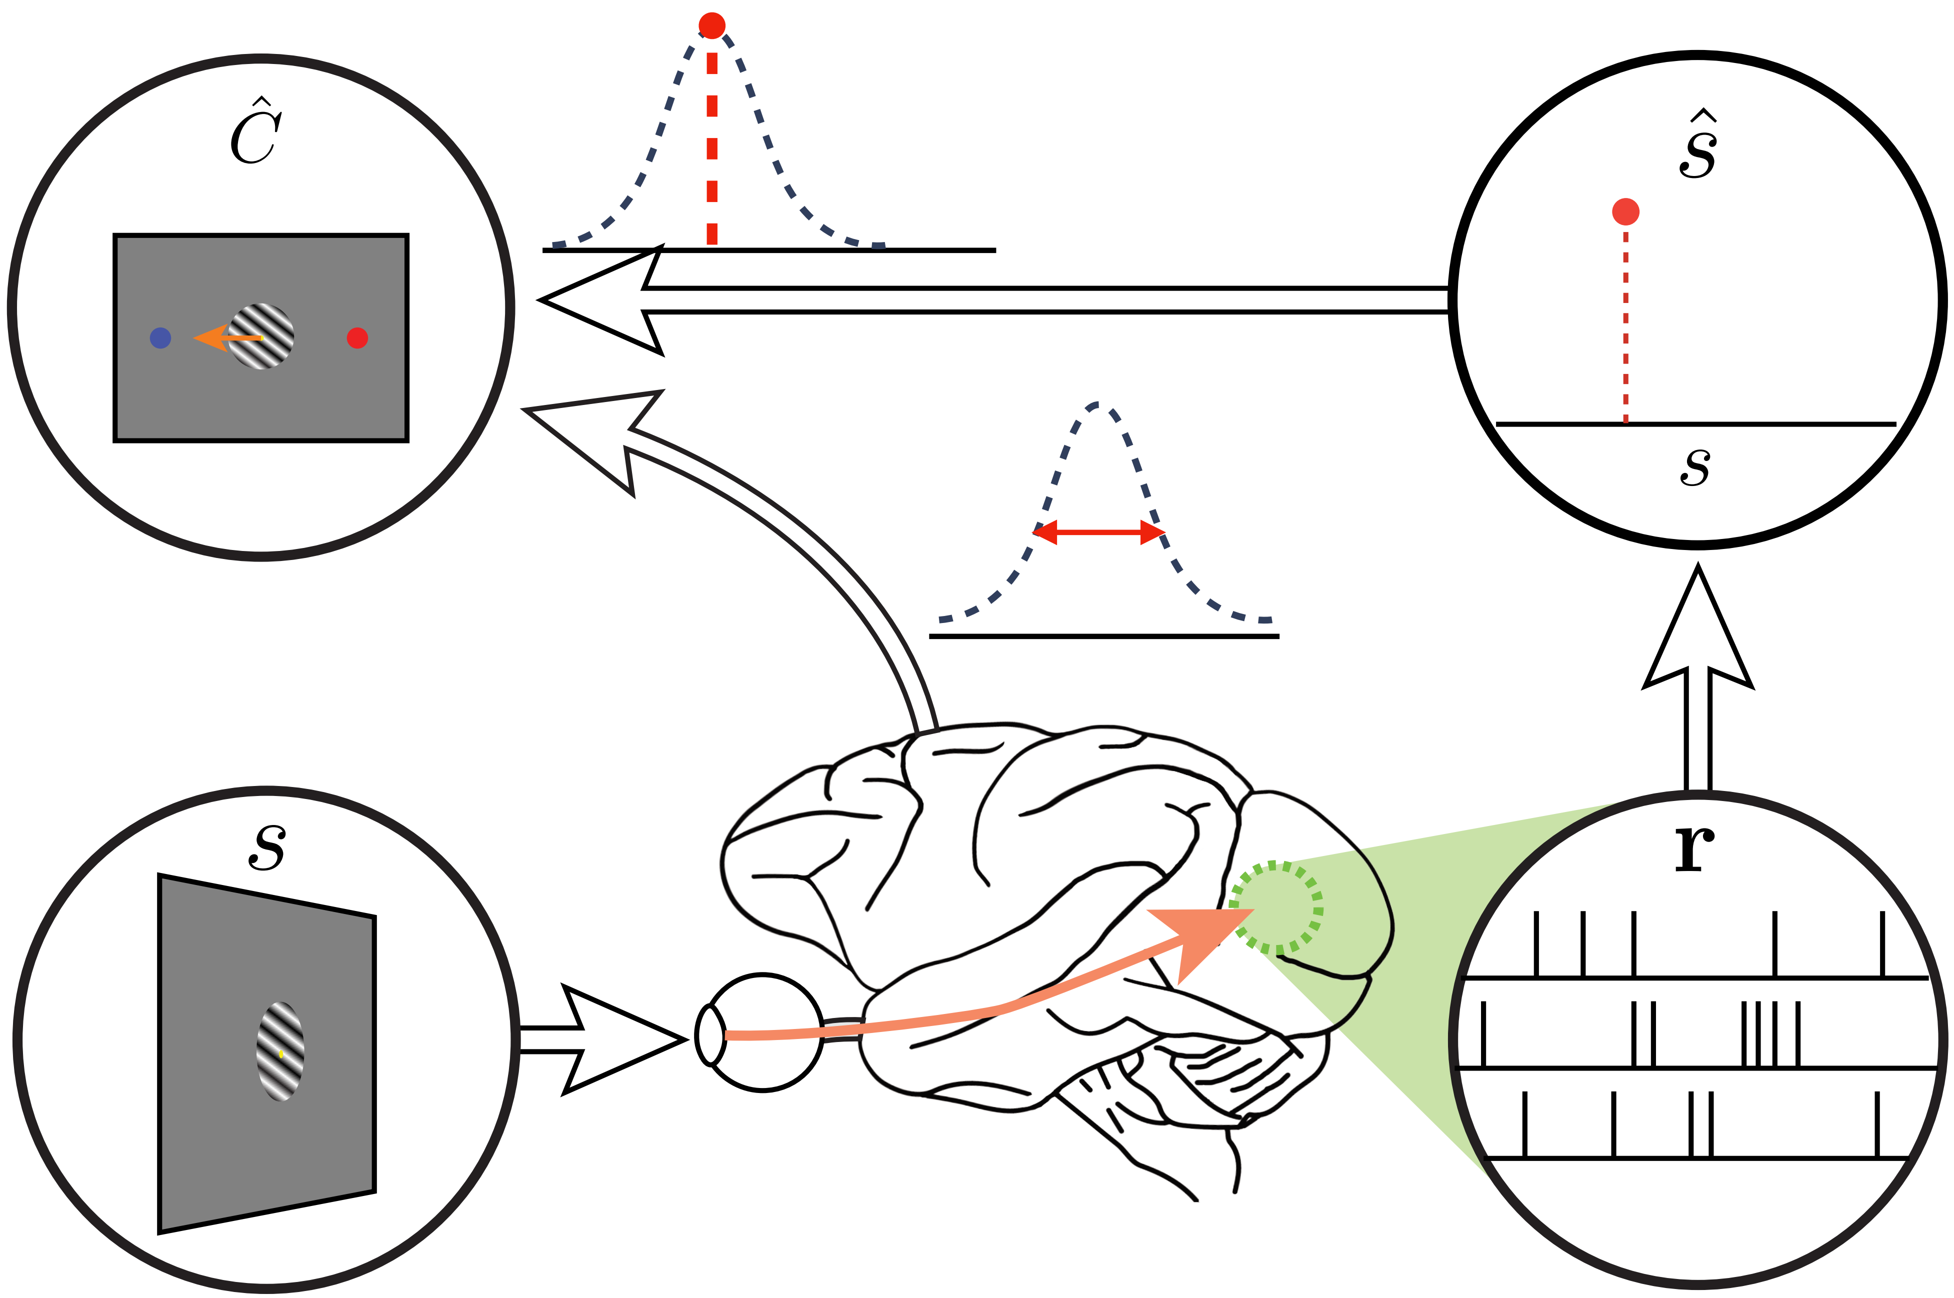


***V1 population only encodes the best point estimate***

# Testing two hypotheses

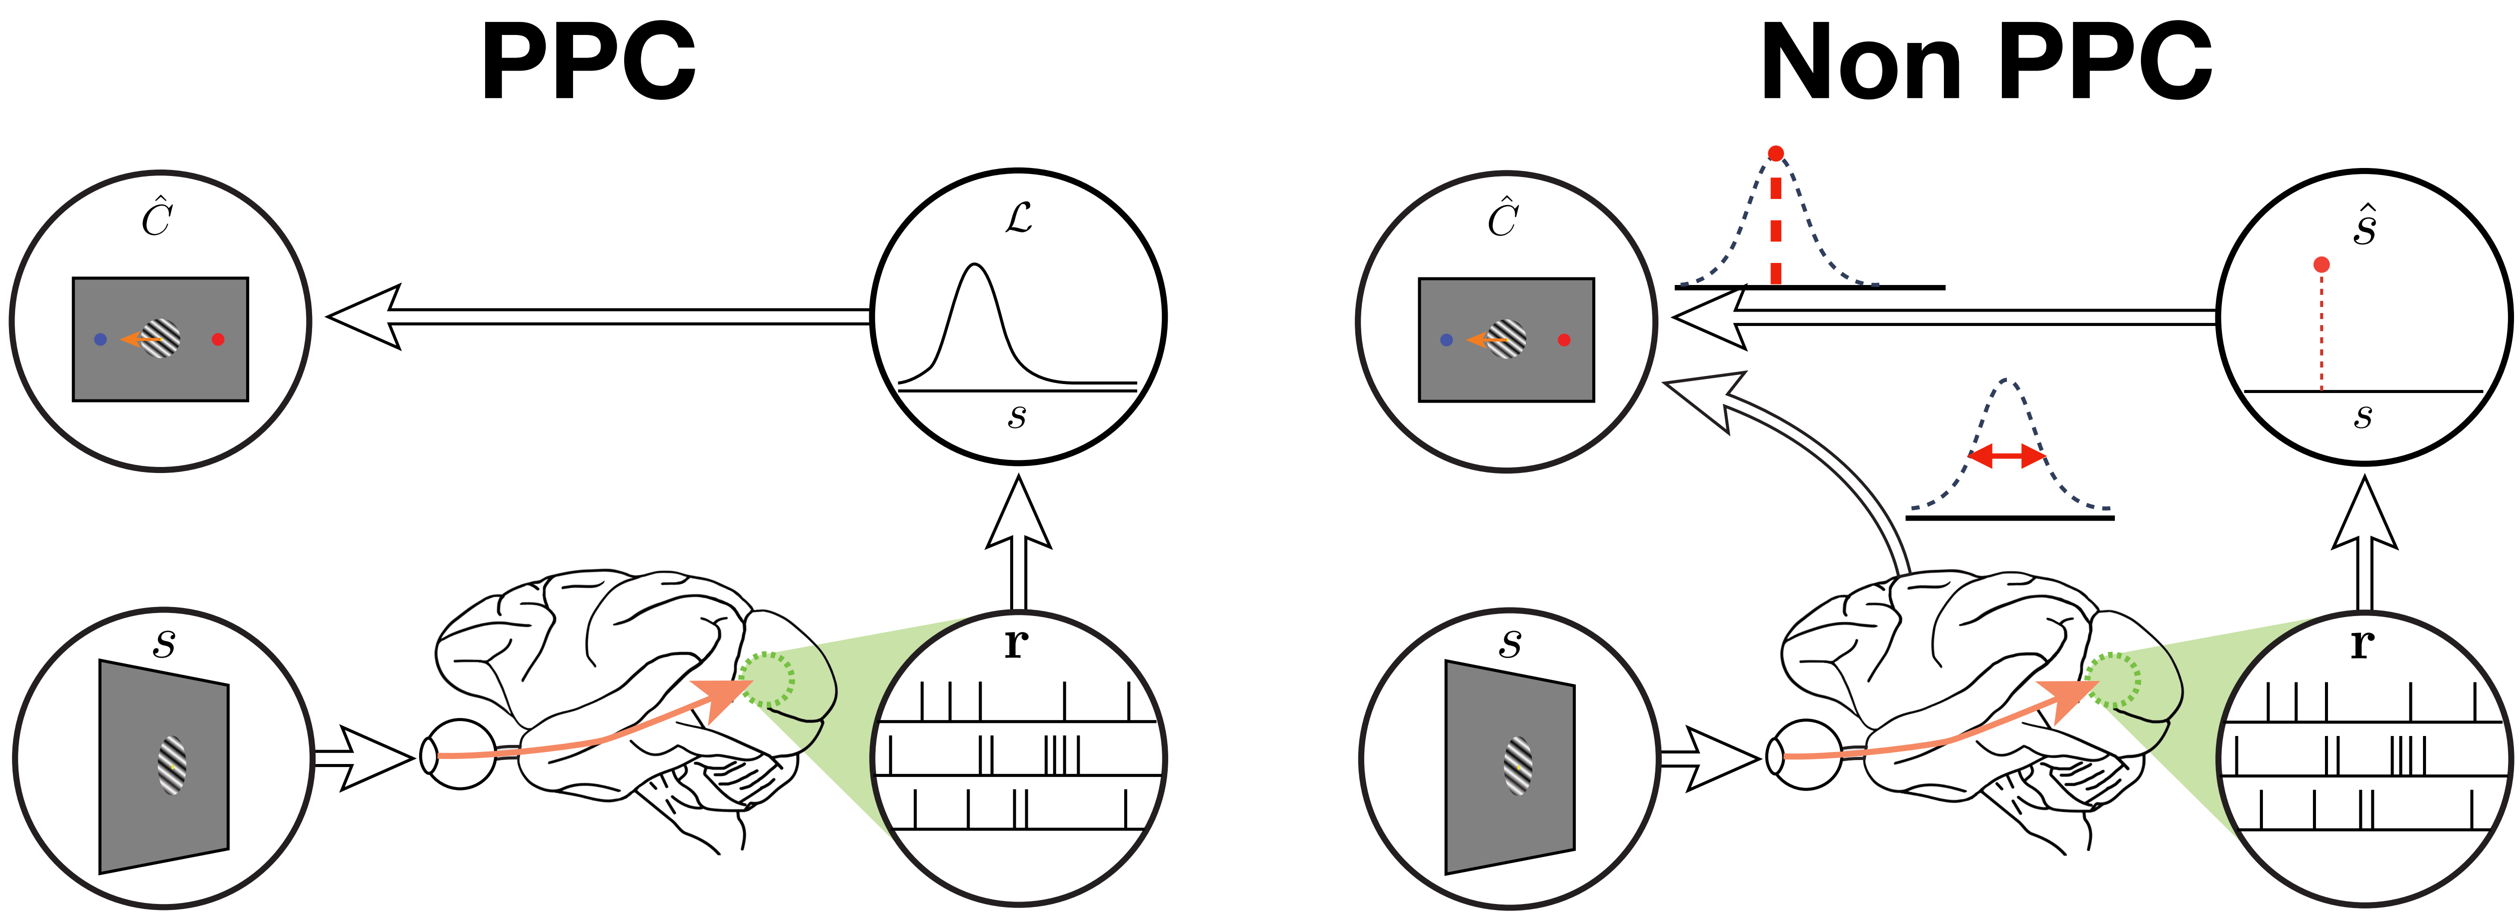


**Figure 1**

***PPC predicts that the trial to trial fluctuations in the shape of the likelihood helps to predict the behavior***

Walker, Cotton et al., 2020 *Nature Neuro*

# Critical pieces in testing PPC

- Task that requires the **use of trial-by-trial sensory uncertainty**
- To decode trial-by-trial likelihood function, you need a **simultaneous population recording** of sensory neurons
- Good method for **decoding likelihood function** on each trial

***Previous studies fail to satisfy all these criteria***

# Approach

1. Train macaque monkeys on a task in which **optimal performance requires trial-by-trial use of uncertainty** on stimulus

2. Record from a **population of V1 neurons**

3. **Decode likelihood function over stimulus** on each trial

4. Predict the monkey's trial-by-trial decisions using **models with and without uncertainty in the shape of likelihood function**

Step 1: Train monkeys on a task that requires the use of trial-by-trial sensory uncertainty

# Classification task

**Classify** an oriented stimulus into one of two classes:
$C$=1 or $C$=2

*C*=1 or *C*=2,
that is the question

Optimal performance requires the use of **trial-by-trial sensory uncertainty information** (Qamar et al., 2013)

# Trial flow

- Pick a class randomly: $C$=1 or $C$=2
- Each class defines a distribution over stimulus orientation

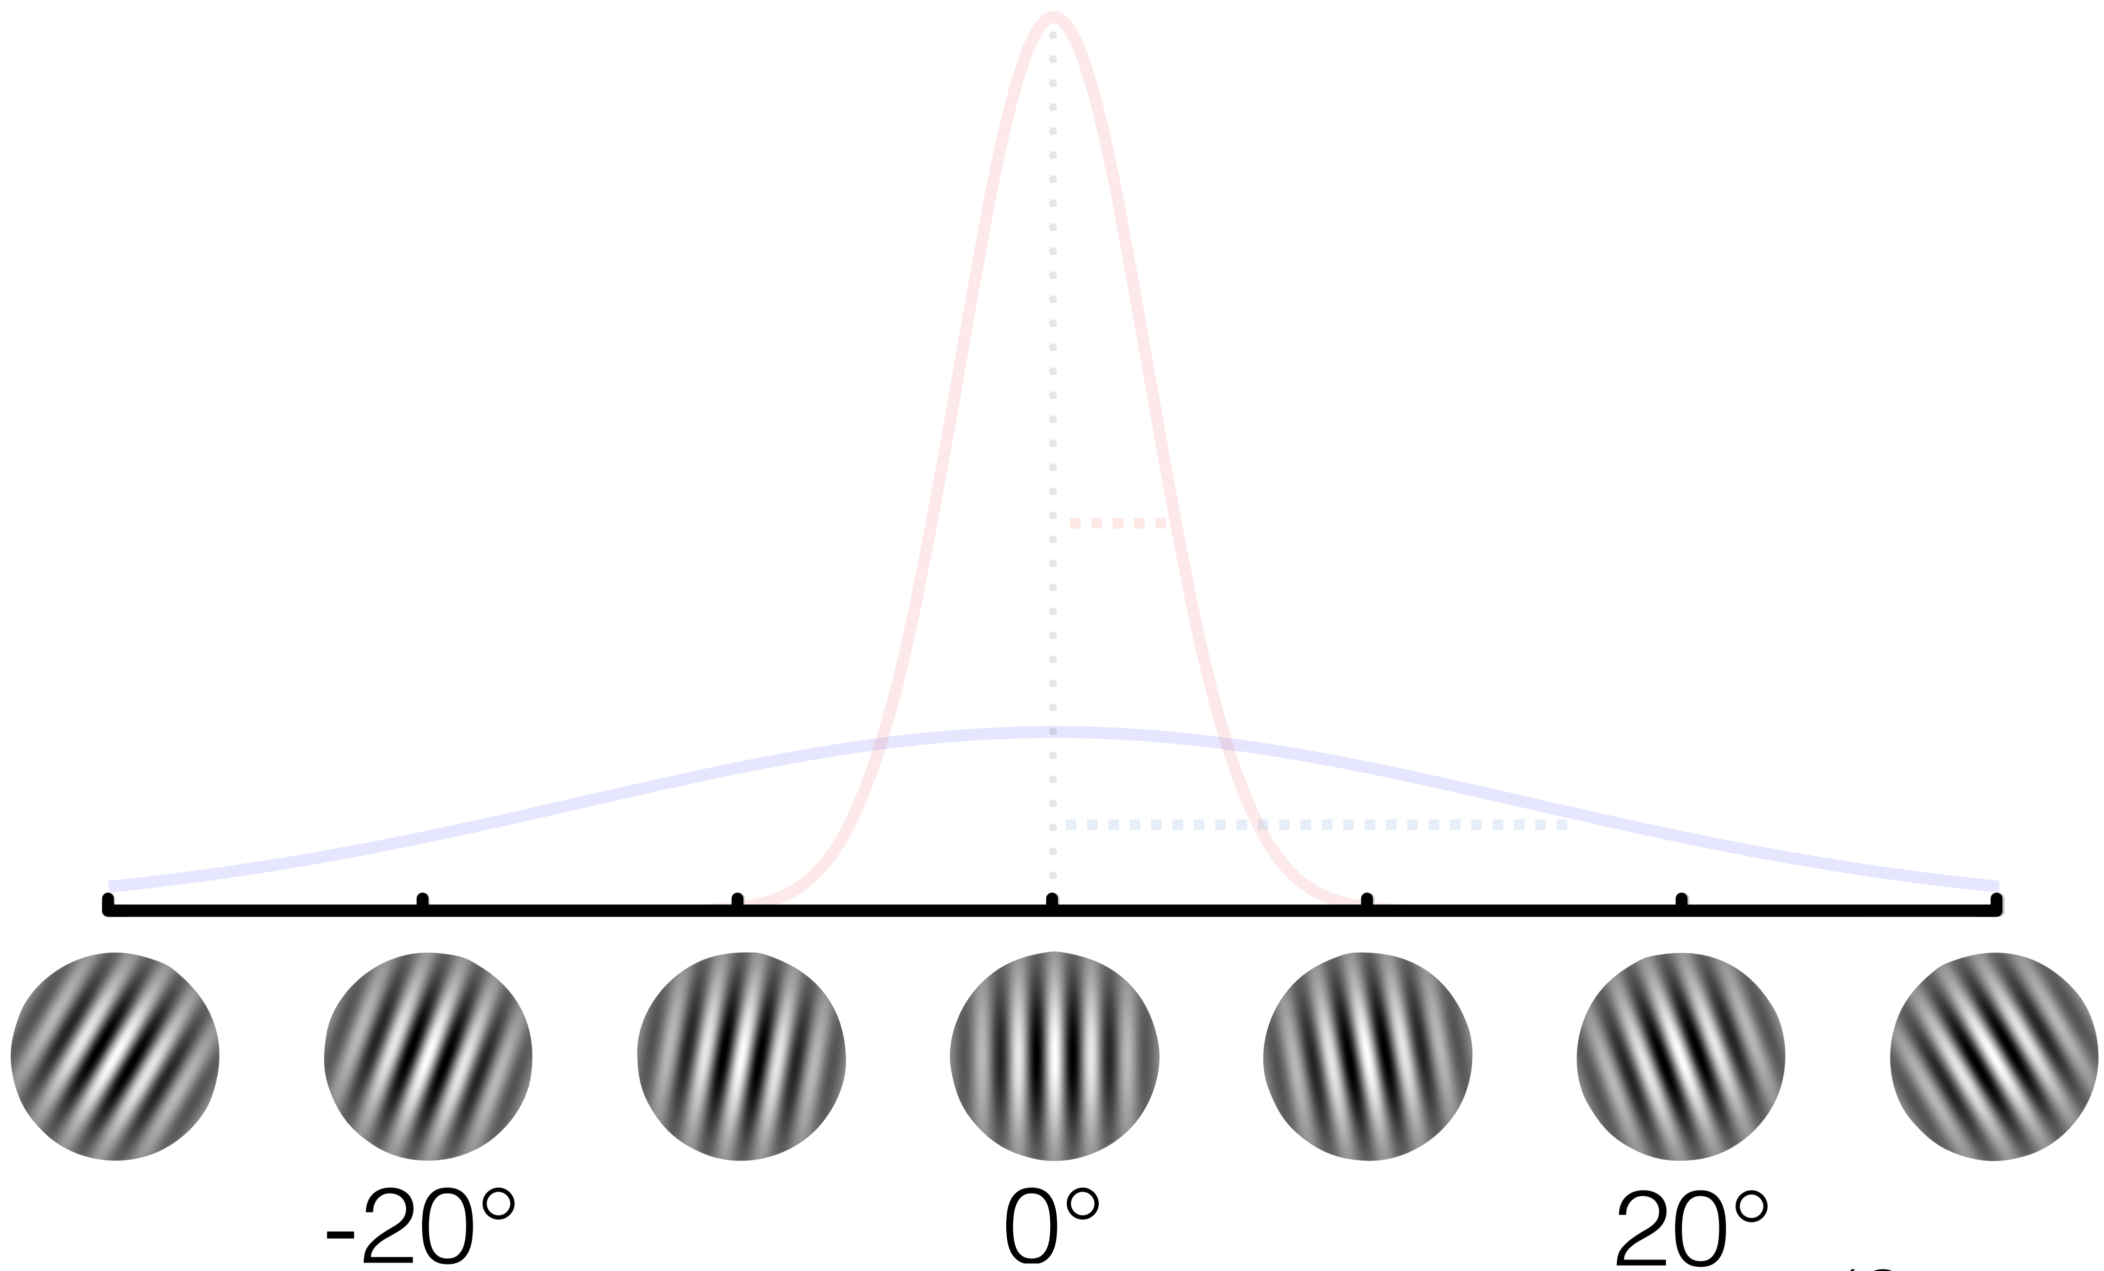


(Qamar et al., 2013)

# Trial flow

- Pick a class randomly: $C$=1 or $C$=2
- Each class defines a distribution over stimulus orientation

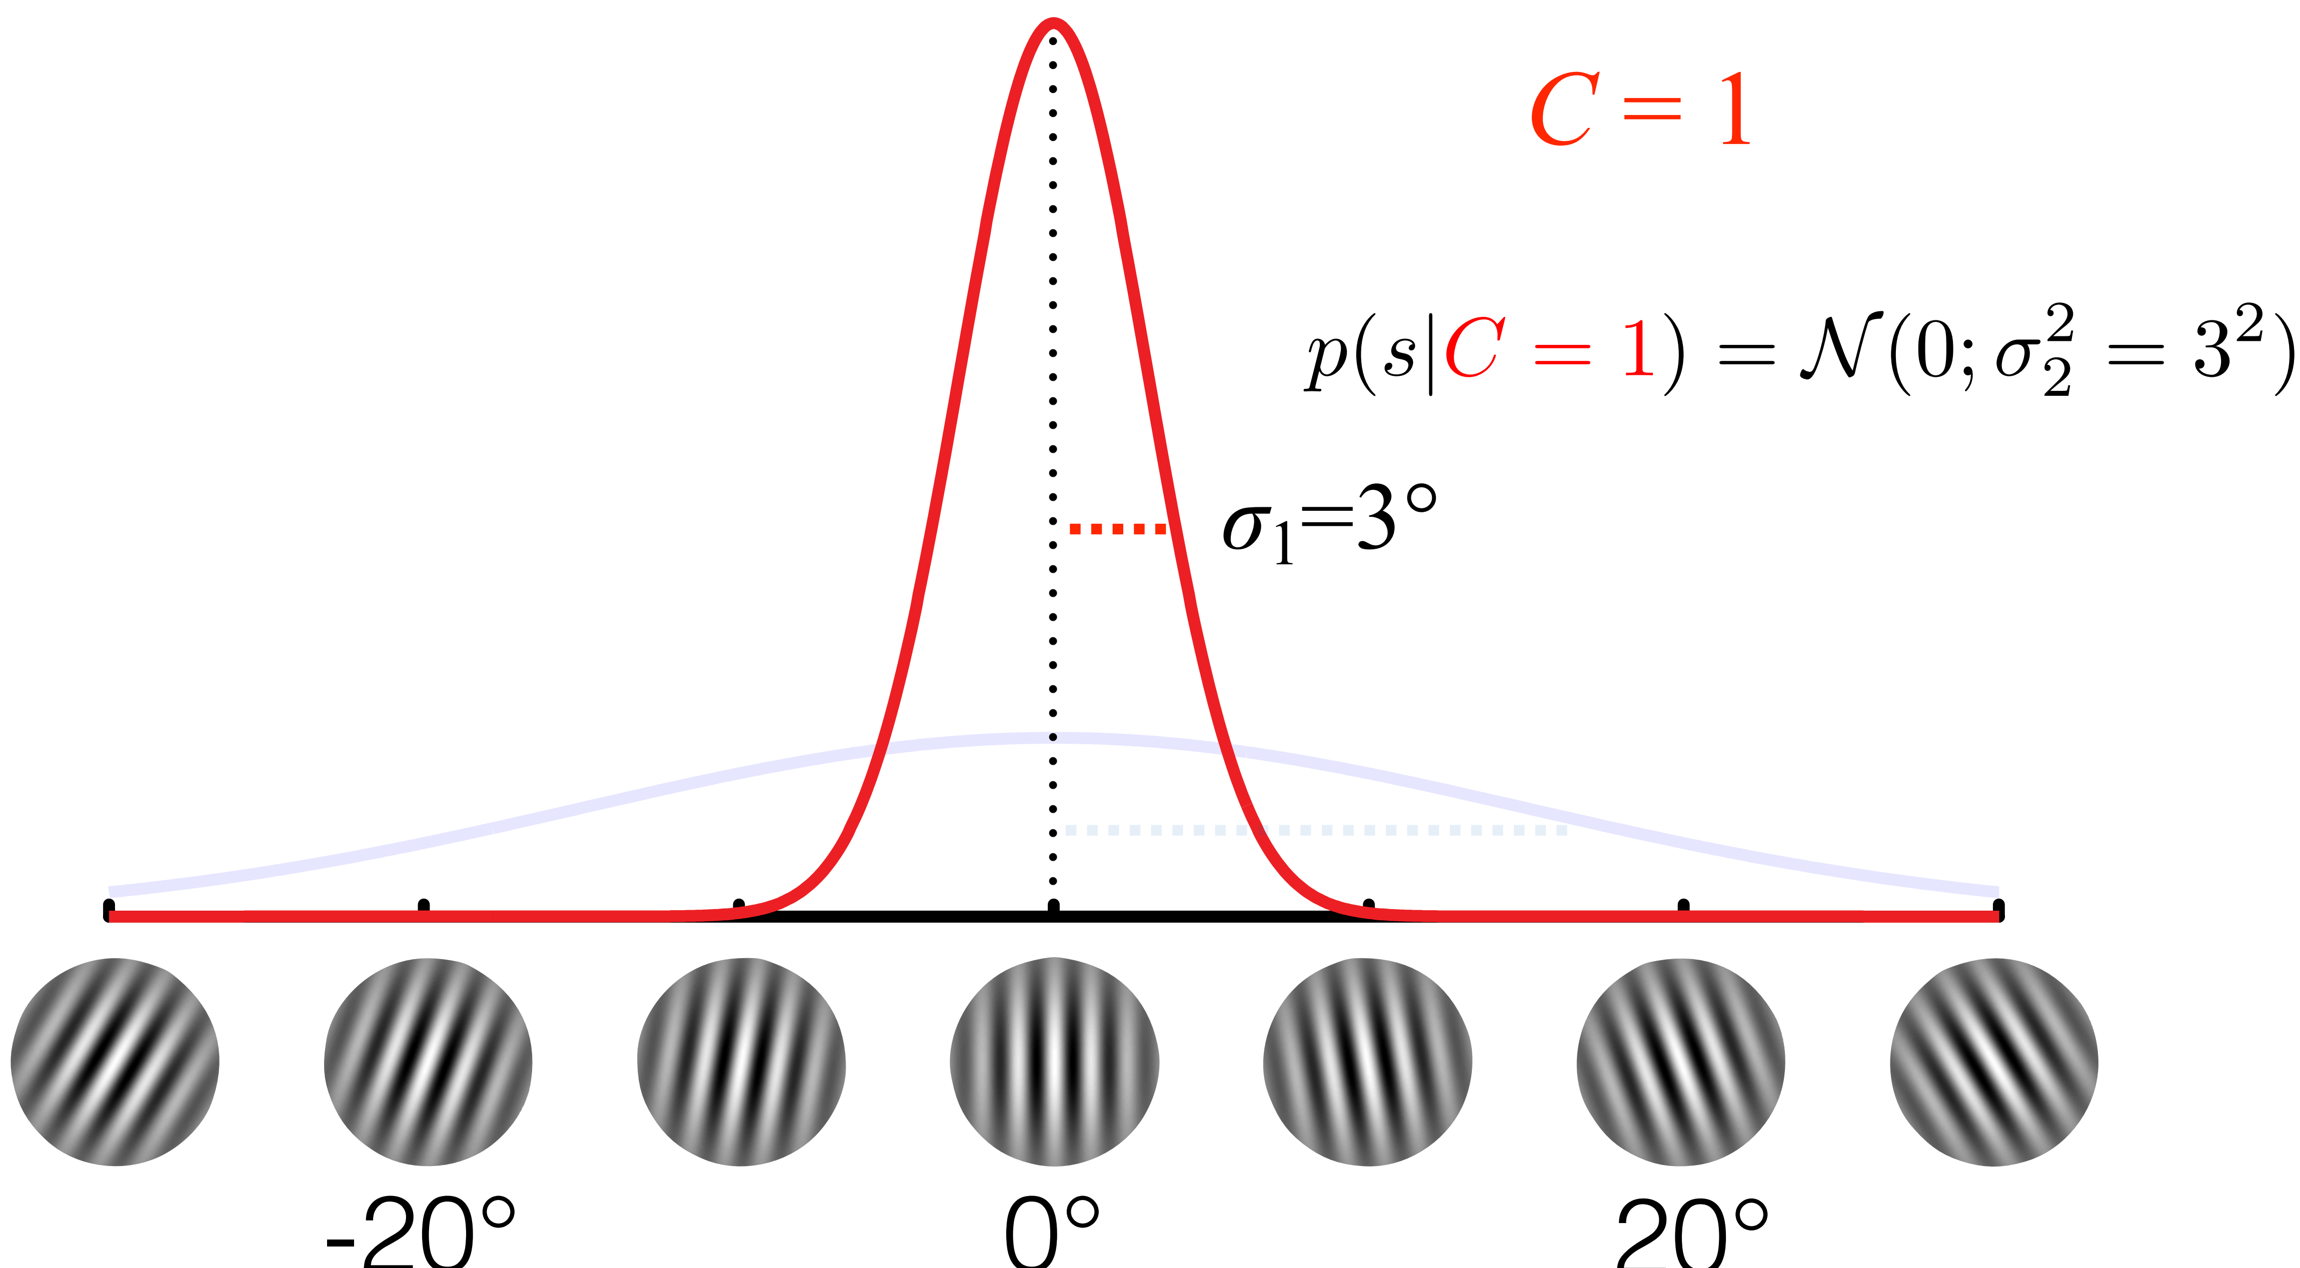


(Qamar et al., 2013)

# Trial flow

- Pick a class randomly: $C$=1 or $C$=2
- Each class defines a distribution over stimulus orientation

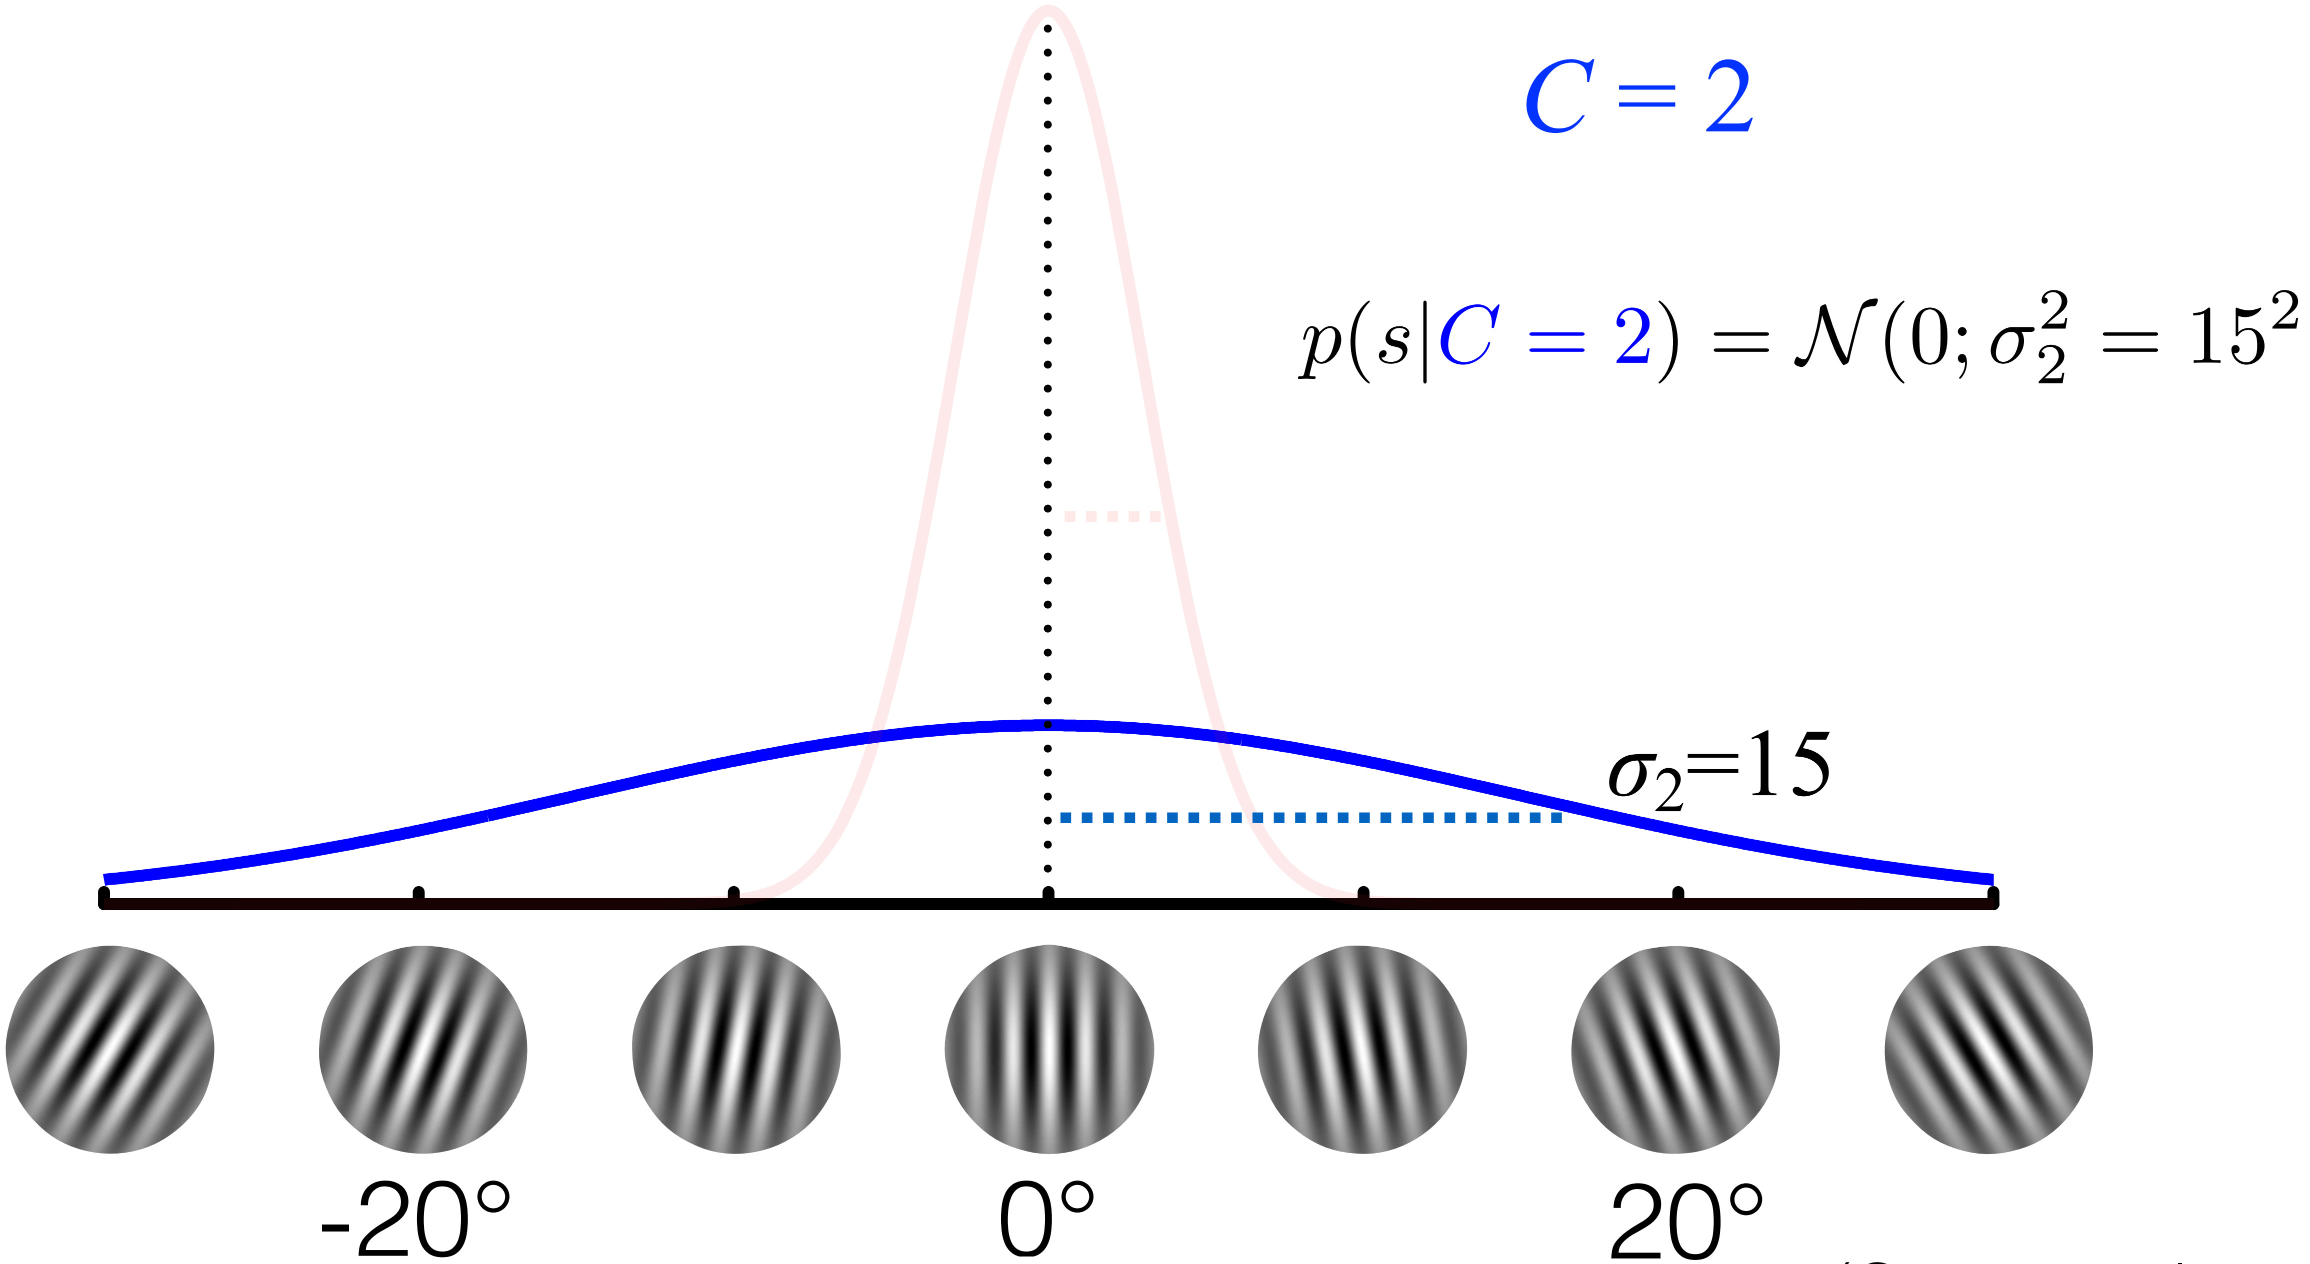


(Qamar et al., 2013)

# Trial flow

- Pick a class randomly: $C$=1 or $C$=2

- Each class defines a distribution over stimulus orientation

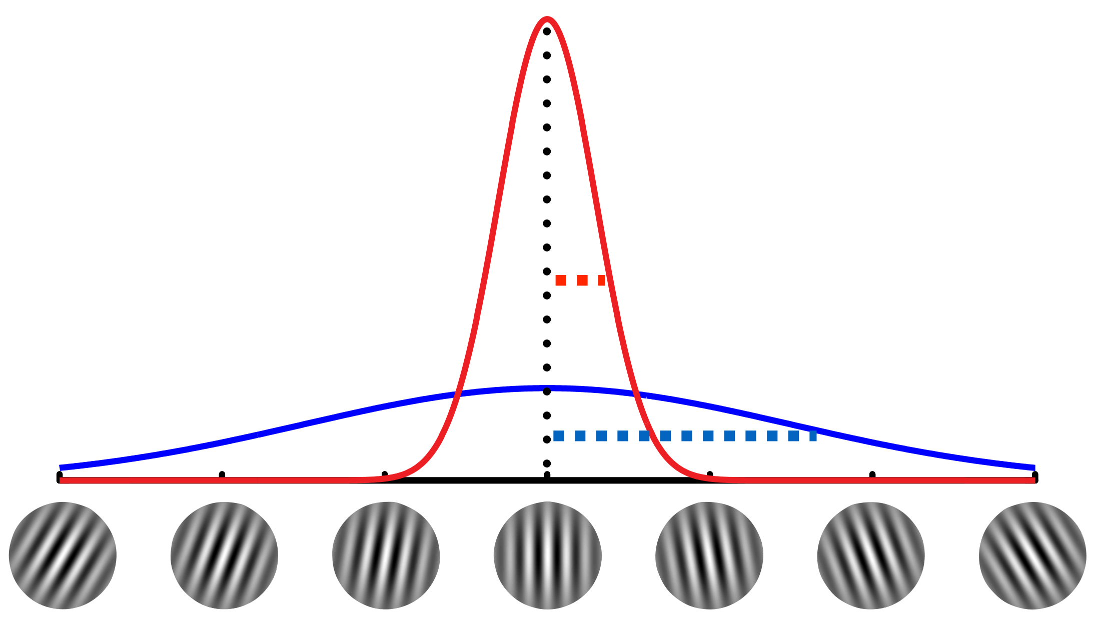

- Pick an orientation from the selected class's distribution

- Contrast is varied from trial to trial

# Optimal strategy

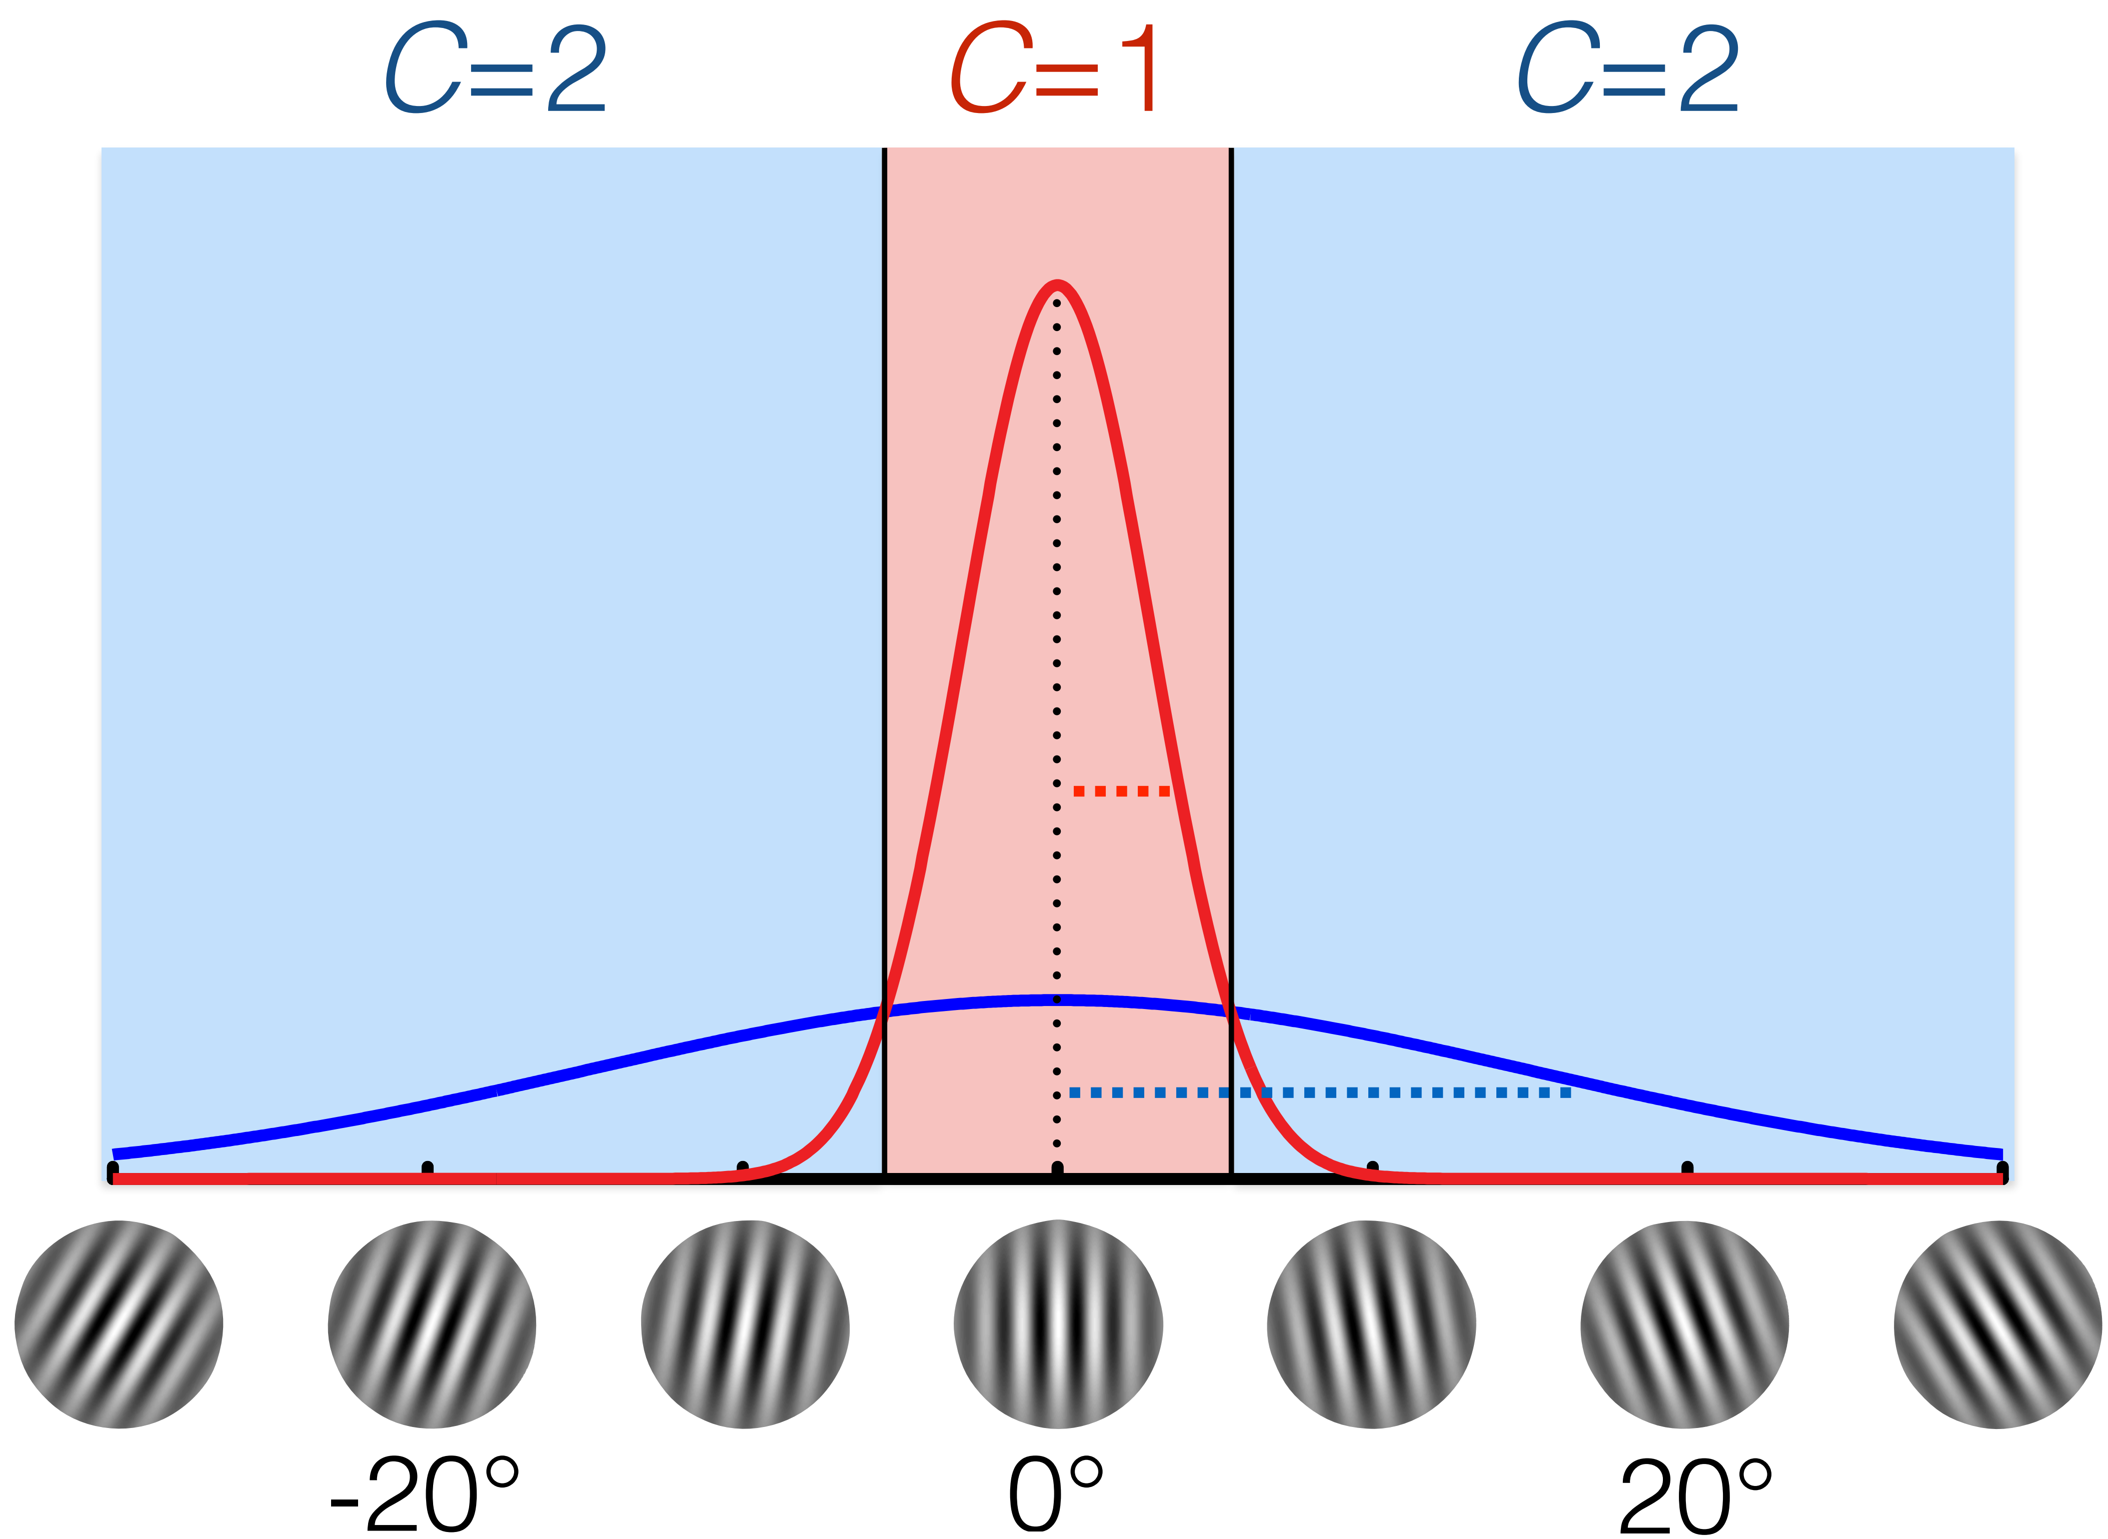

# Noisy observation

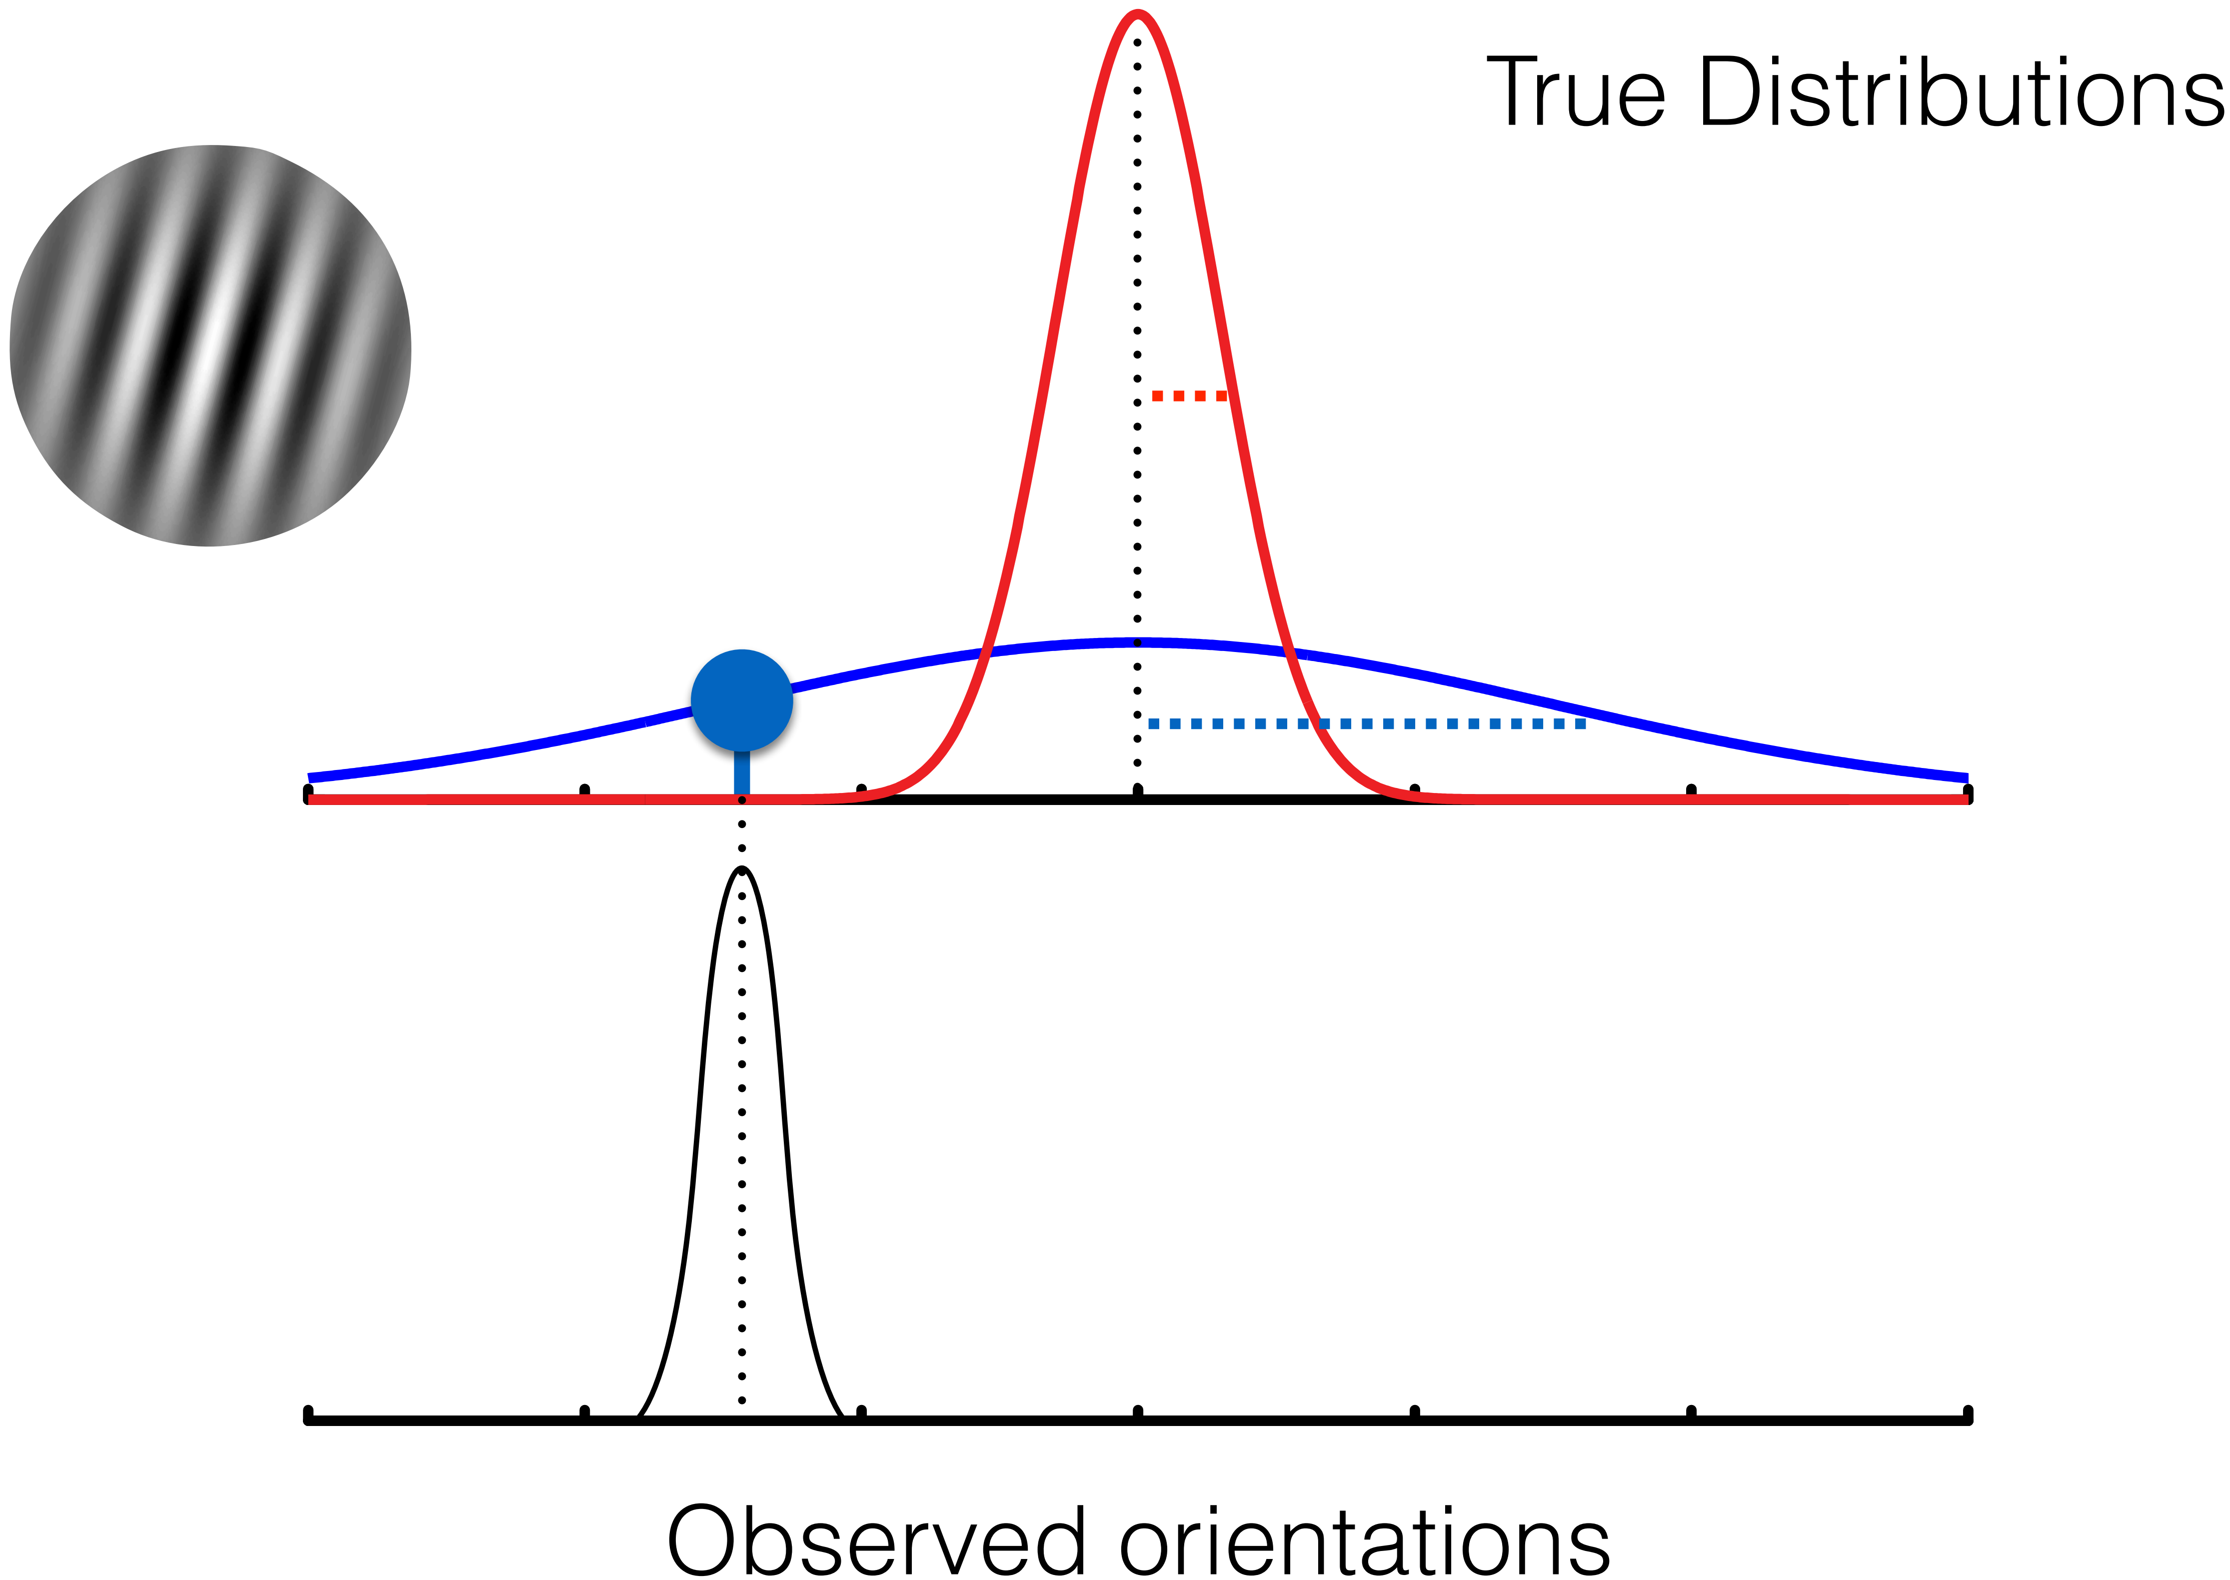

# Noisy observation

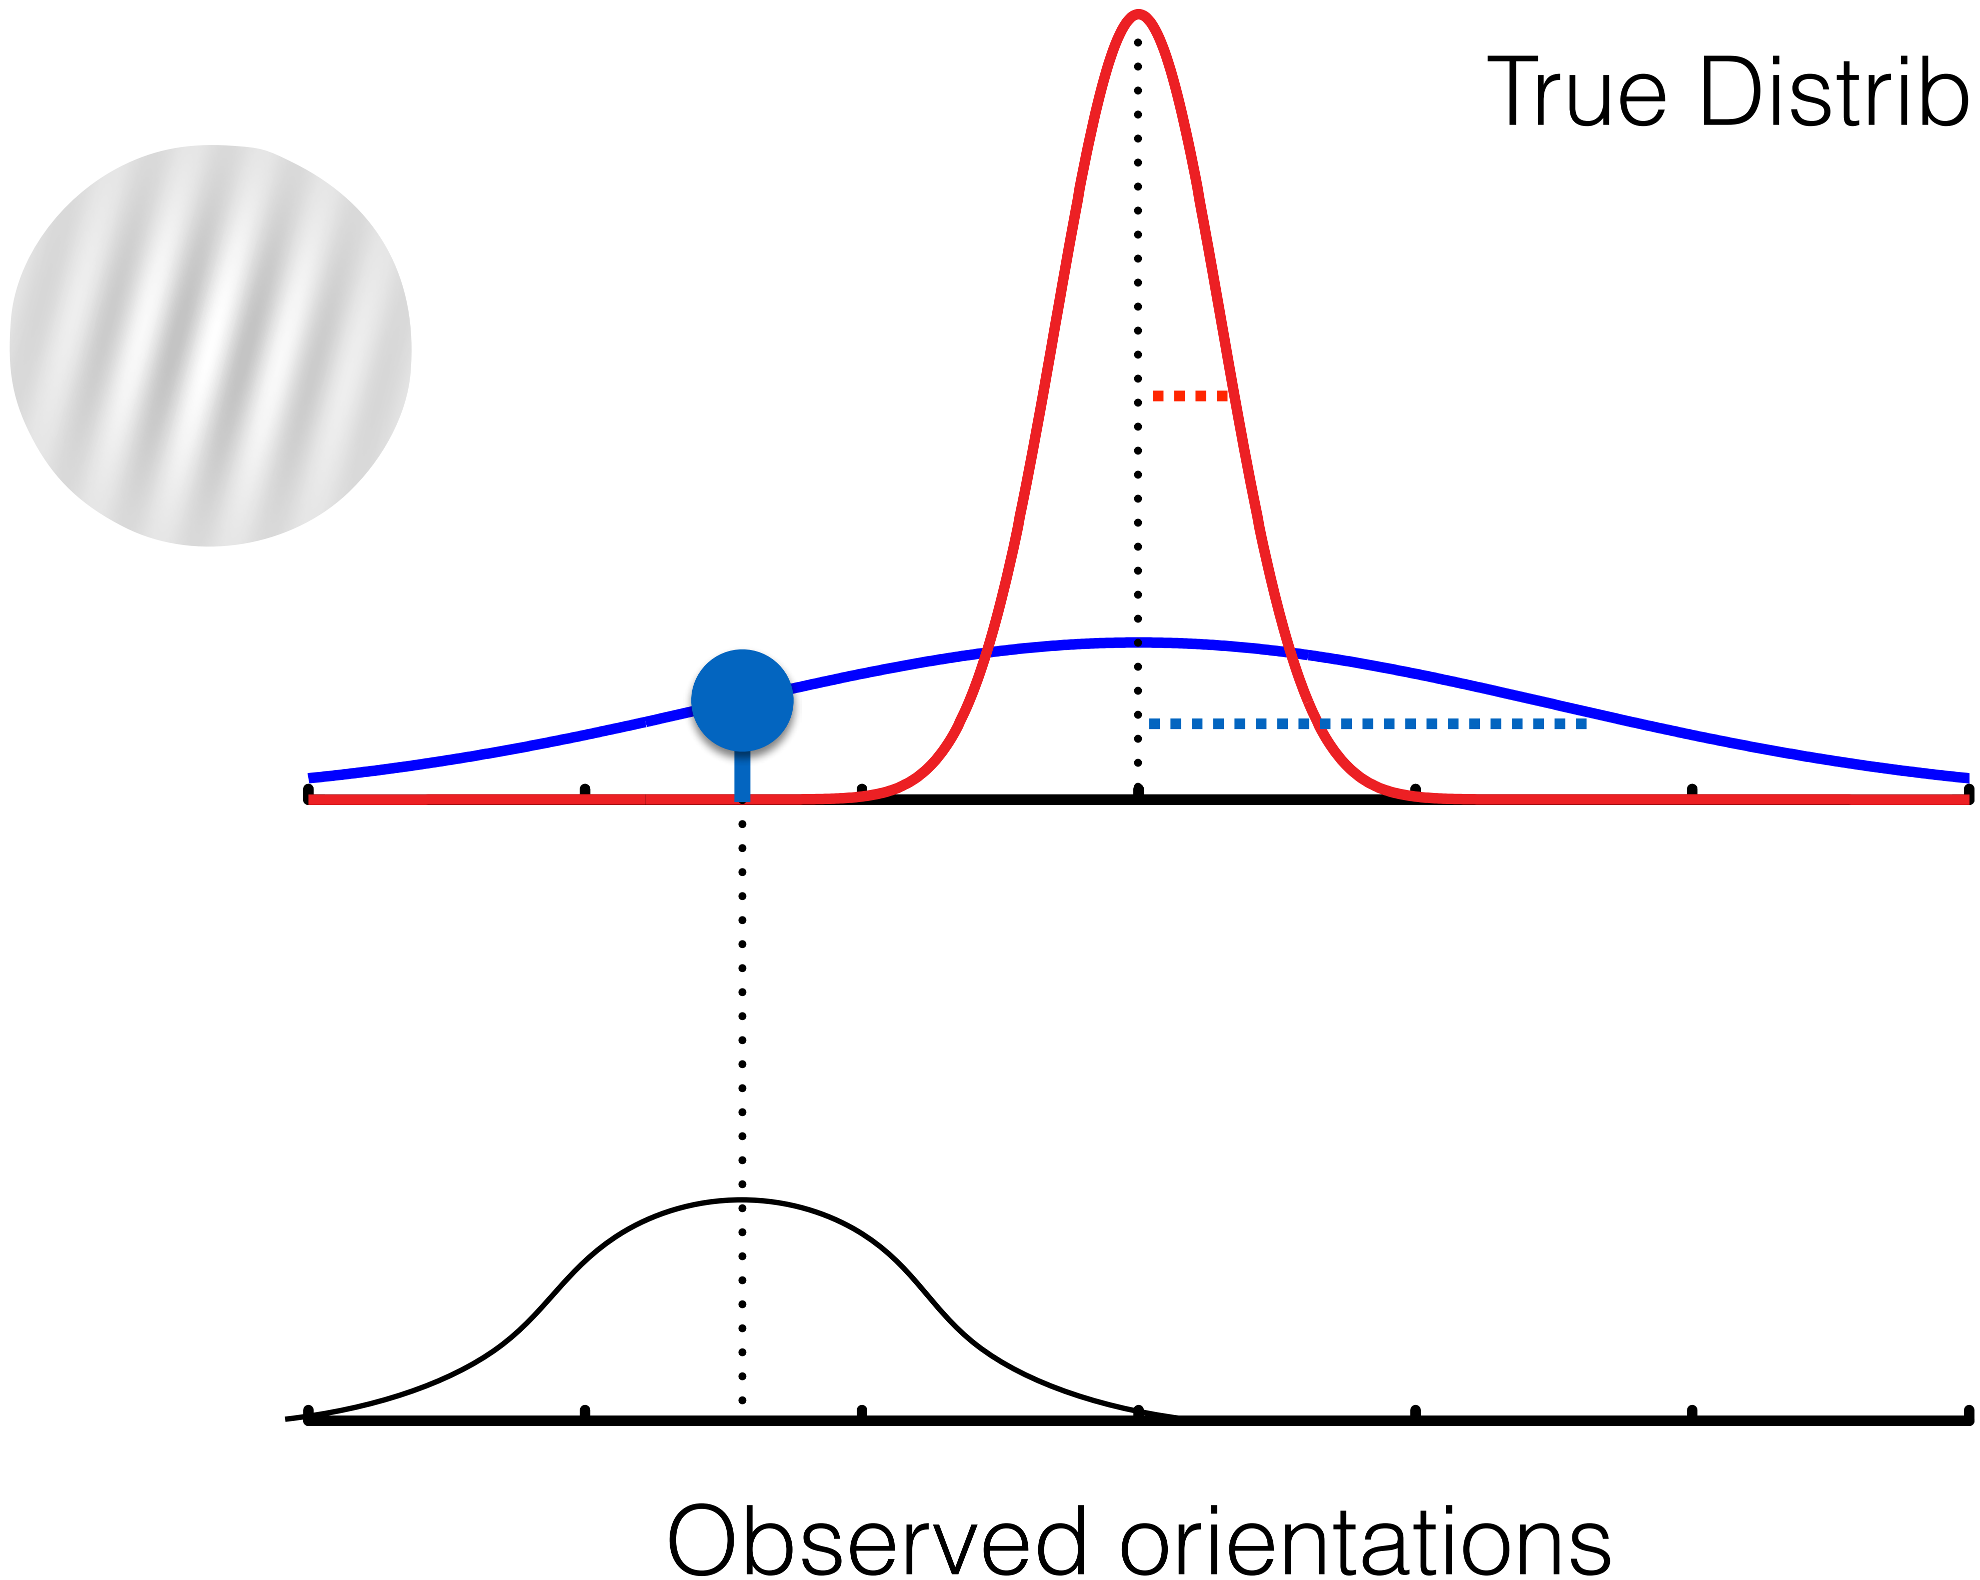

# Optimal decision boundary changes with uncertainty

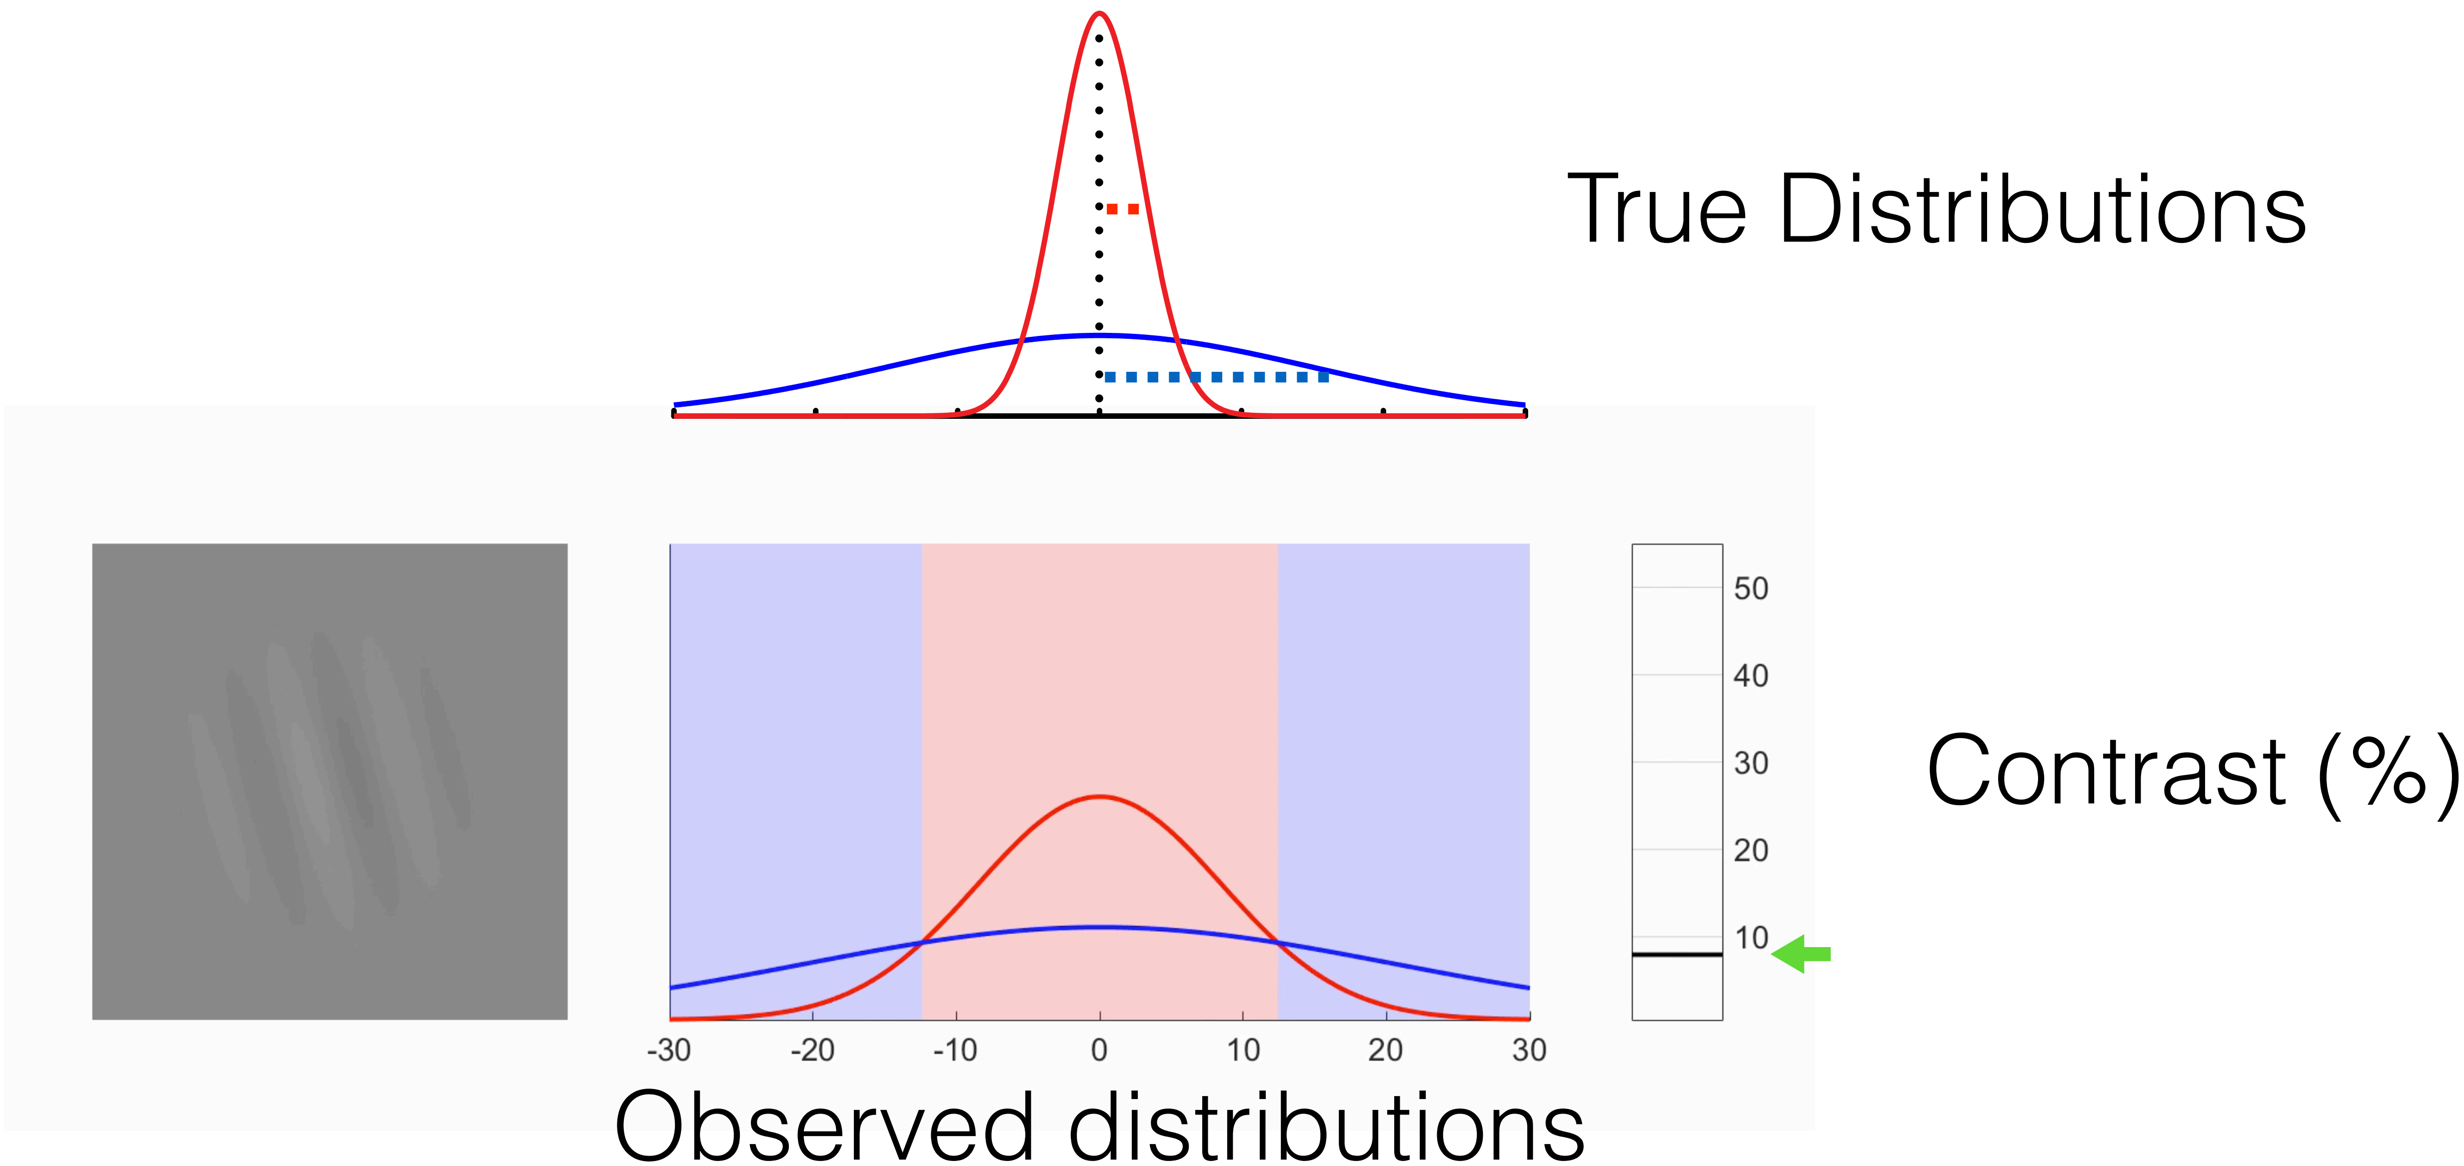


Stimulus contrast is varied on trial-by-trial basis→ optimal decision boundary changes on each trial according to the uncertainty

Step 2: Record activities of a population of V1 as they perform the task

# Recording from V1 population

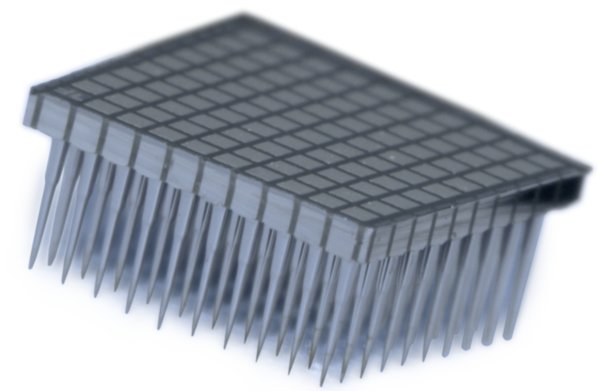

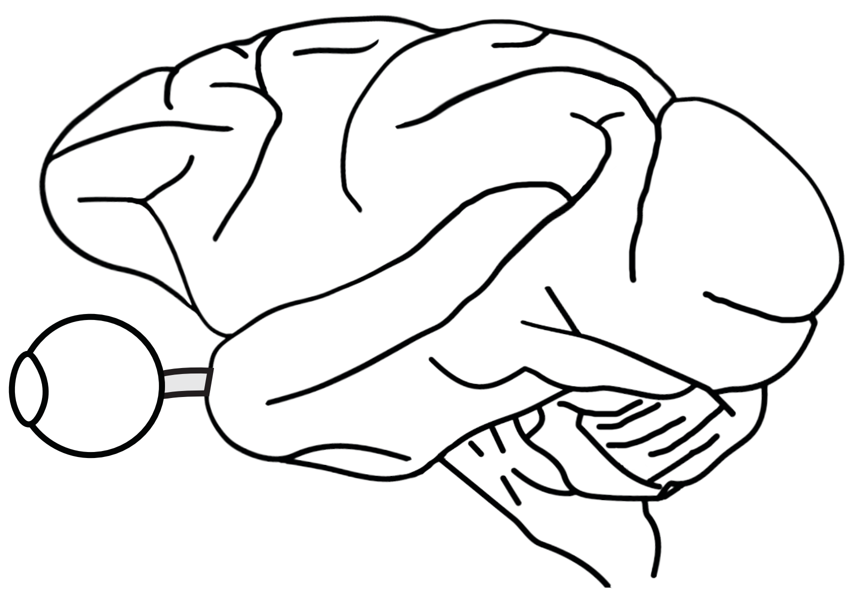

10 X 10 Multielectrode
(Utah) array
Up to 96 multi-units
recording

# Recording from V1 population

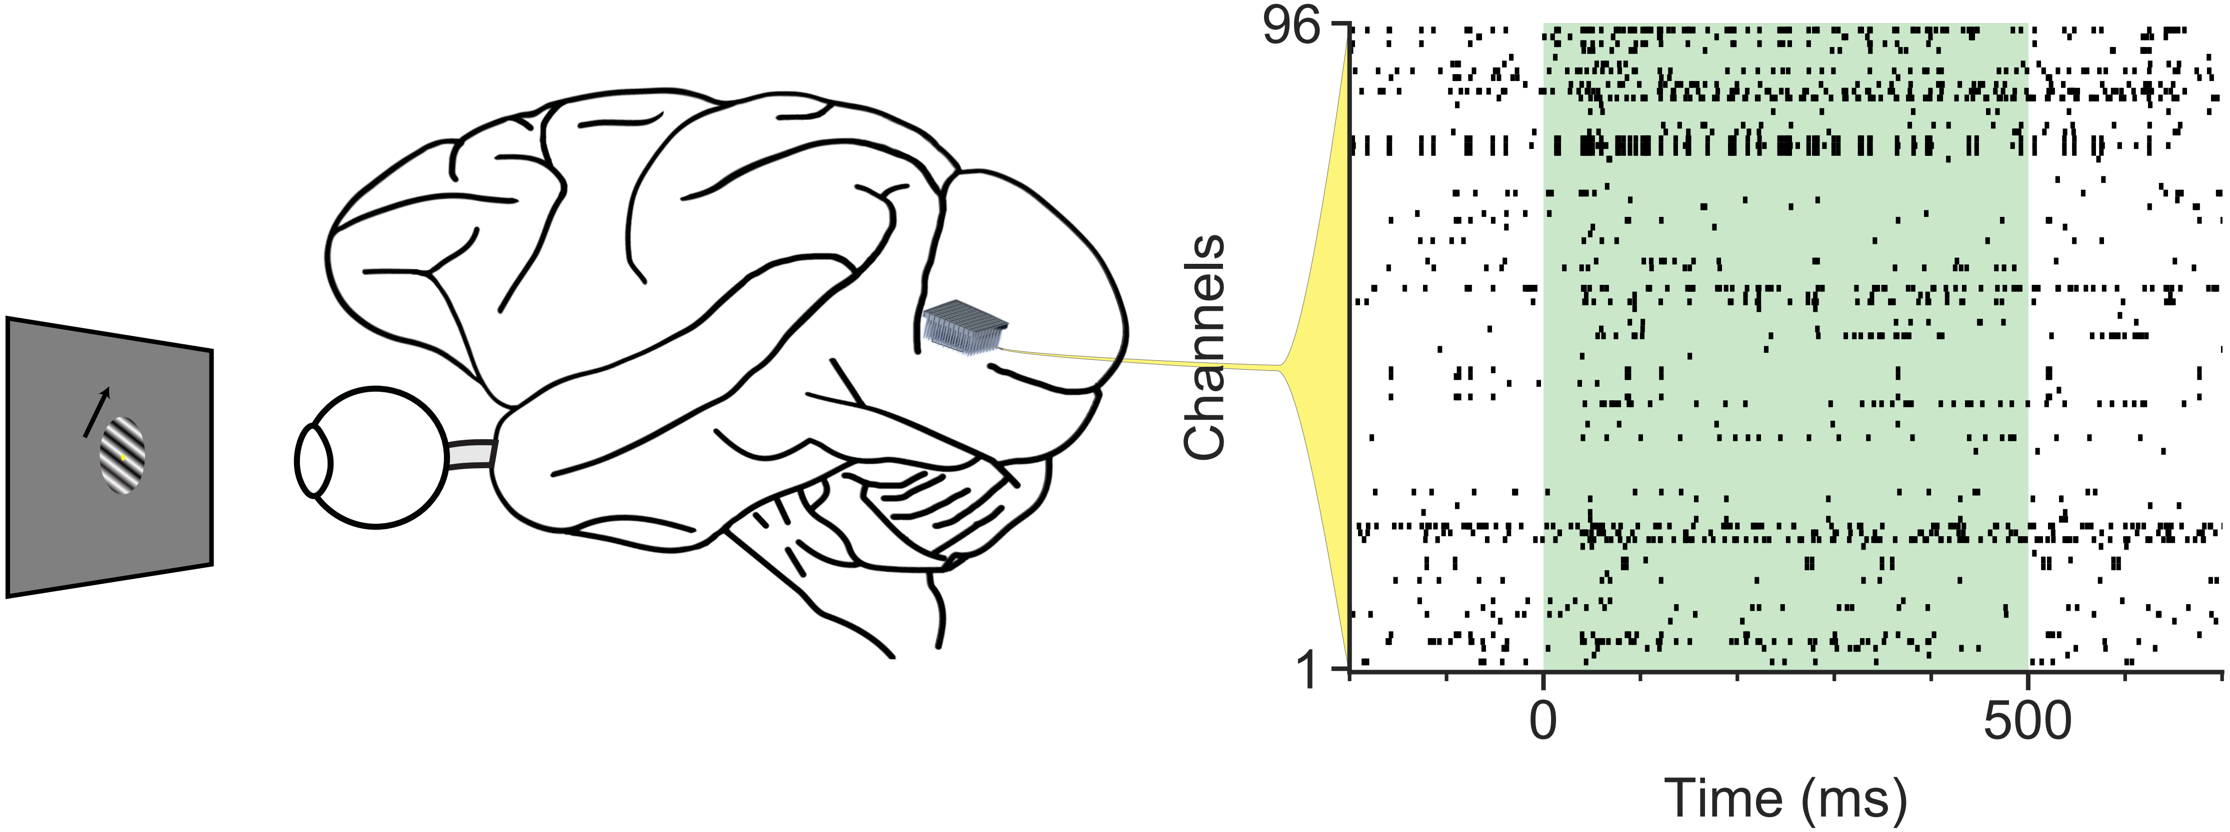


**Figure 3a**

Walker, Cotton et al., 2020 *Nature Neuro*

# Recording from V1 population

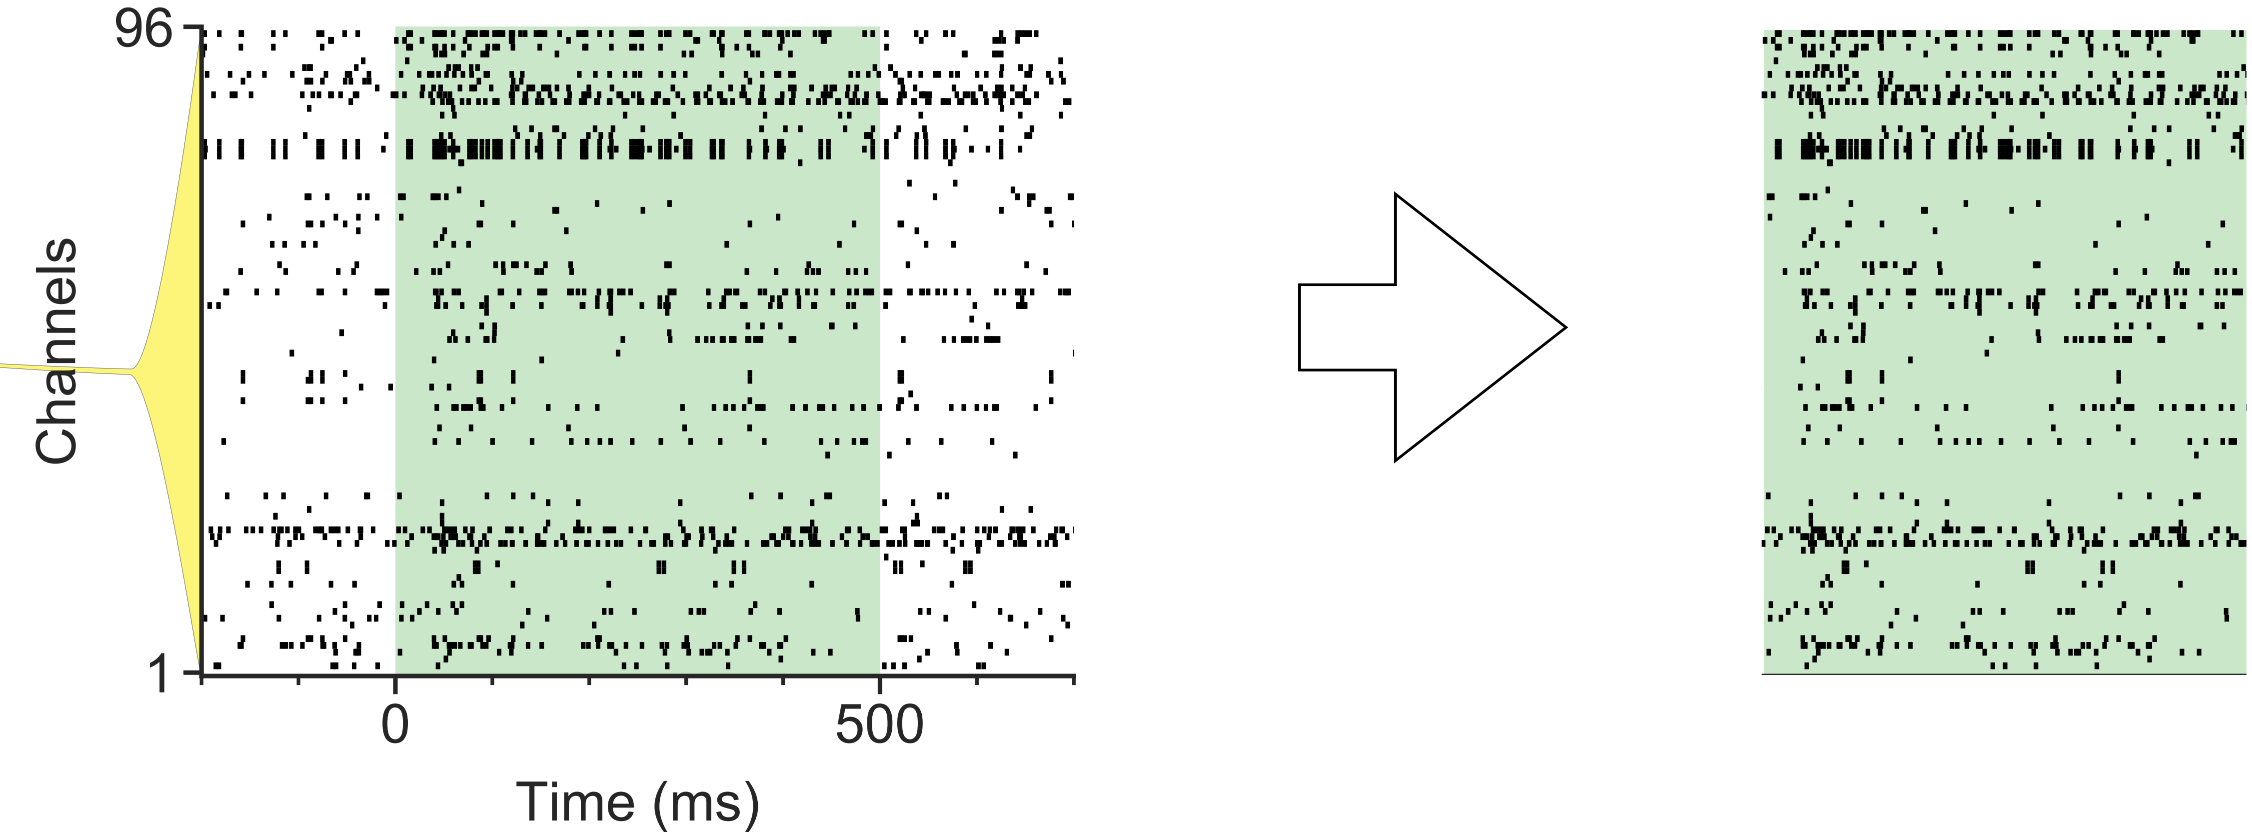


Walker, Cotton et al., 2020 *Nature Neuro*

# Recording from V1 population

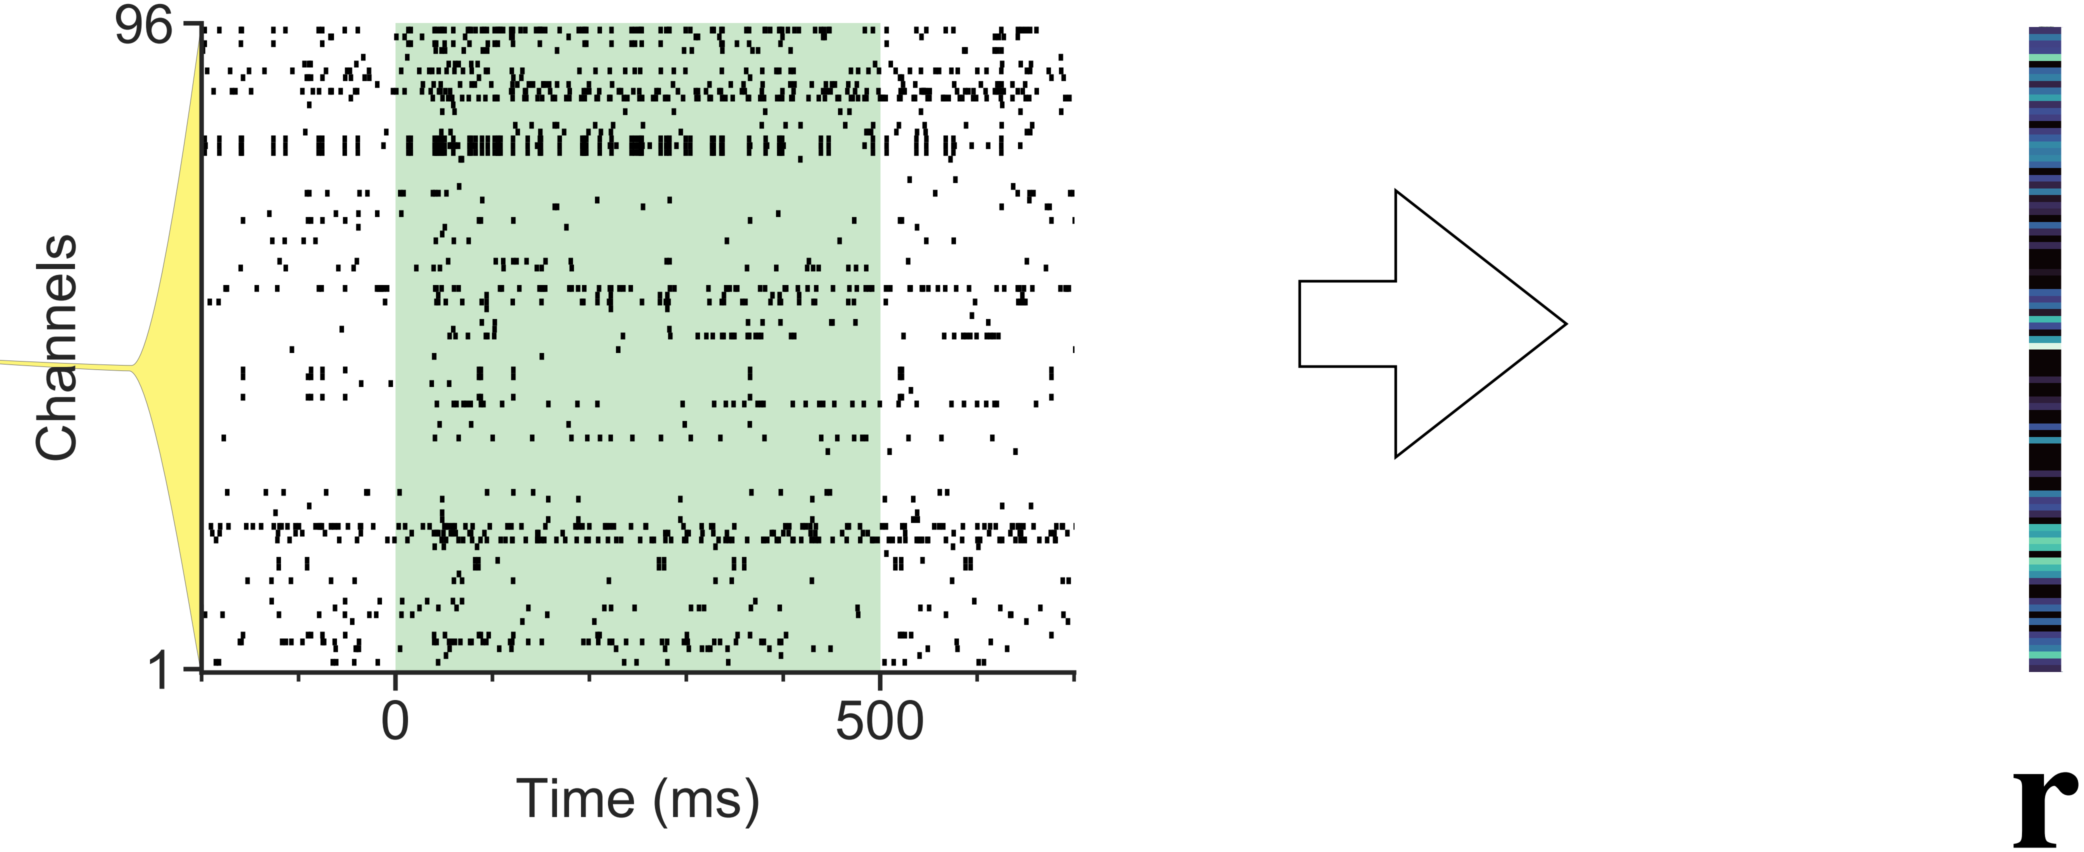


96 multi-unit activity
total spike counts

Walker, Cotton et al., 2020 *Nature Neuro*

# Step 3: Decode trial-by-trial likelihood functions

# Decoding likelihood function from population response

- Traditional method makes assumption about the **how a population of neurons would fire to each stimulus**

- This is known as the **generative model** of the population response **r:** $p(\mathbf{r} \mid s)$

- Common choice includes Tuning curve + Independent Poisson distribution

- Wrong assumptions can lead to **biases in the decoded likelihood functions**!

# DNN based likelihood decoding

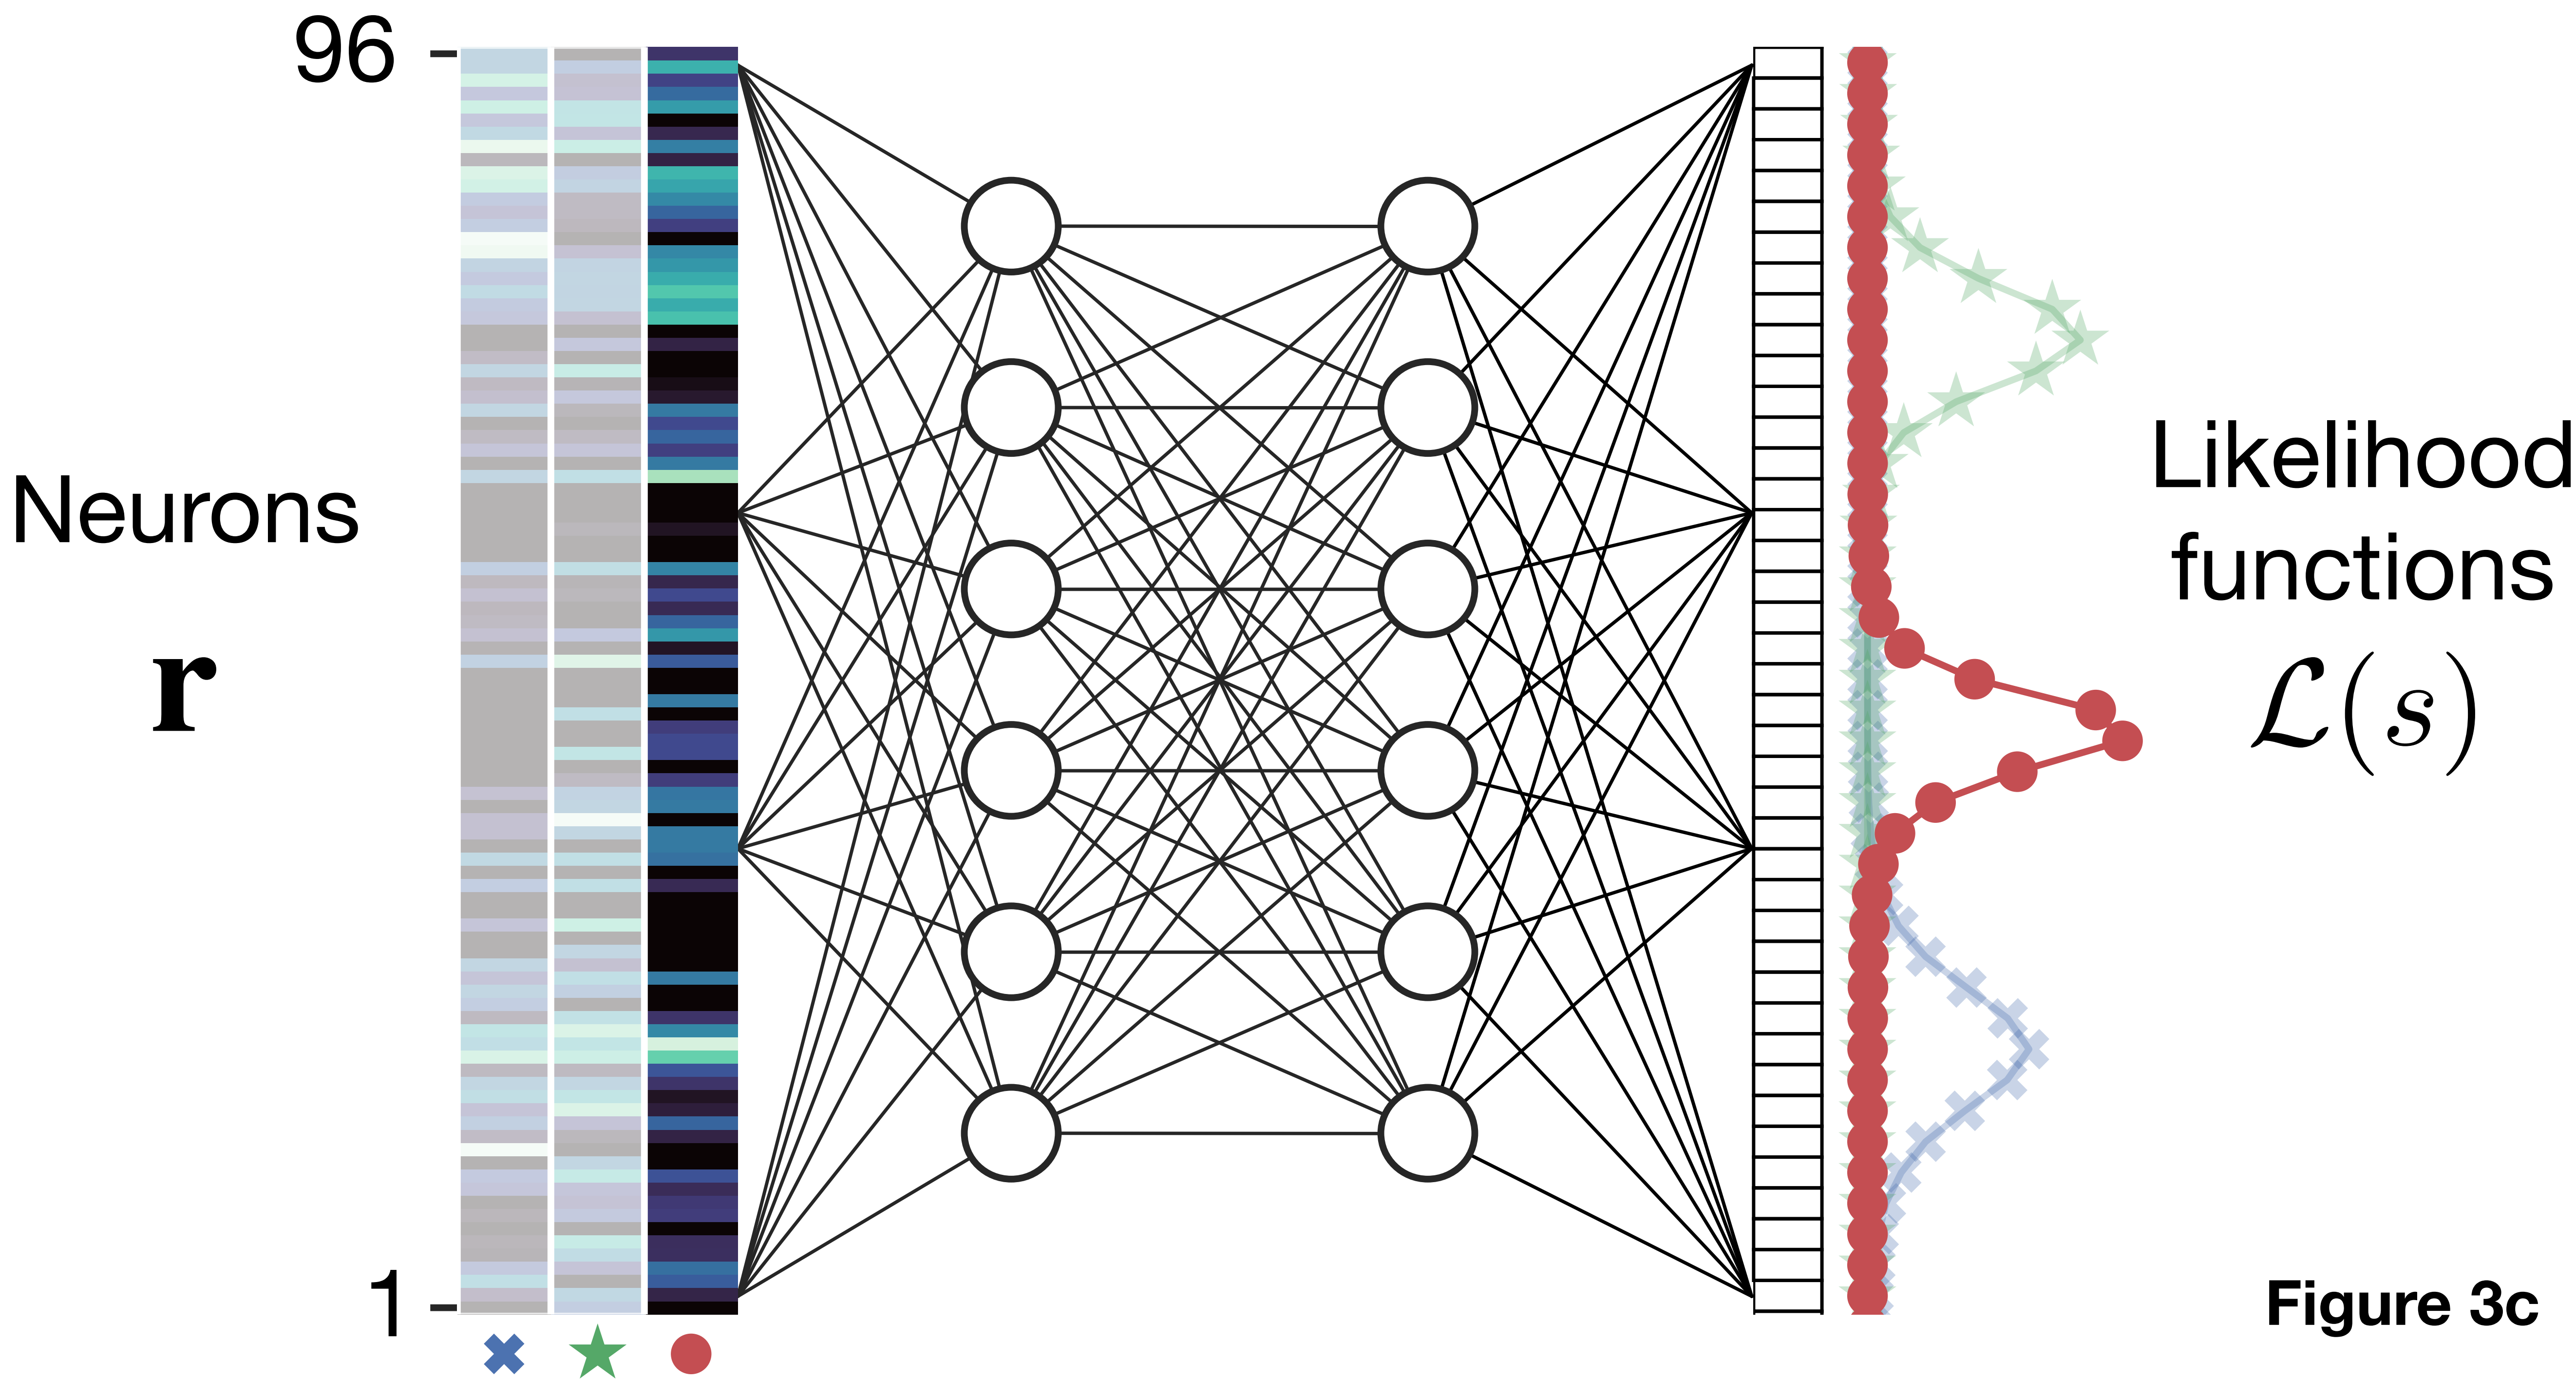


***Train a network to directly decode likelihood function from the population response r***

Walker, Cotton et al., 2020 *Nature Neuro*

# (Details on likelihood decoding) ★

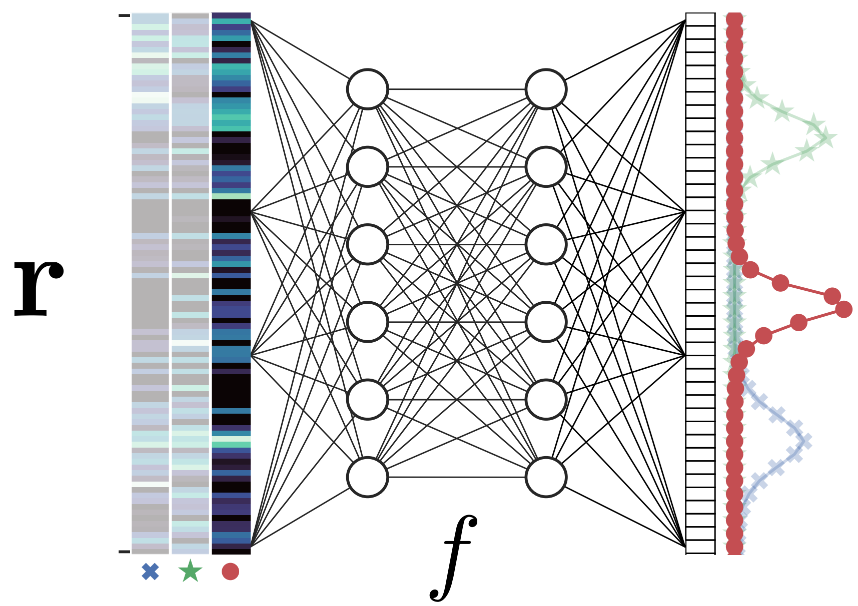


$\mathcal{L}(s) \Rightarrow \mathbf{L} = f(\mathbf{r})$

(discretized) DNN

$\boxed{\log \mathbf{L} \approx \log \mathcal{L}(s) = \log p(\mathbf{r}|s)}$ + $\boxed{\log \mathbf{p_s} \approx \log p(s)}$

**Likelihood function** **Prior**

**softmax**

**target:** $s$ $\boxed{\log p(s|\mathbf{r})}$ **Posterior**

**Cross-entropy loss**

*Could learn the likelihood function up to a multiplicative constant for each value of* $\mathbf{r}$

# Step 4: Model comparison

# Full Likelihood Model (PPC)

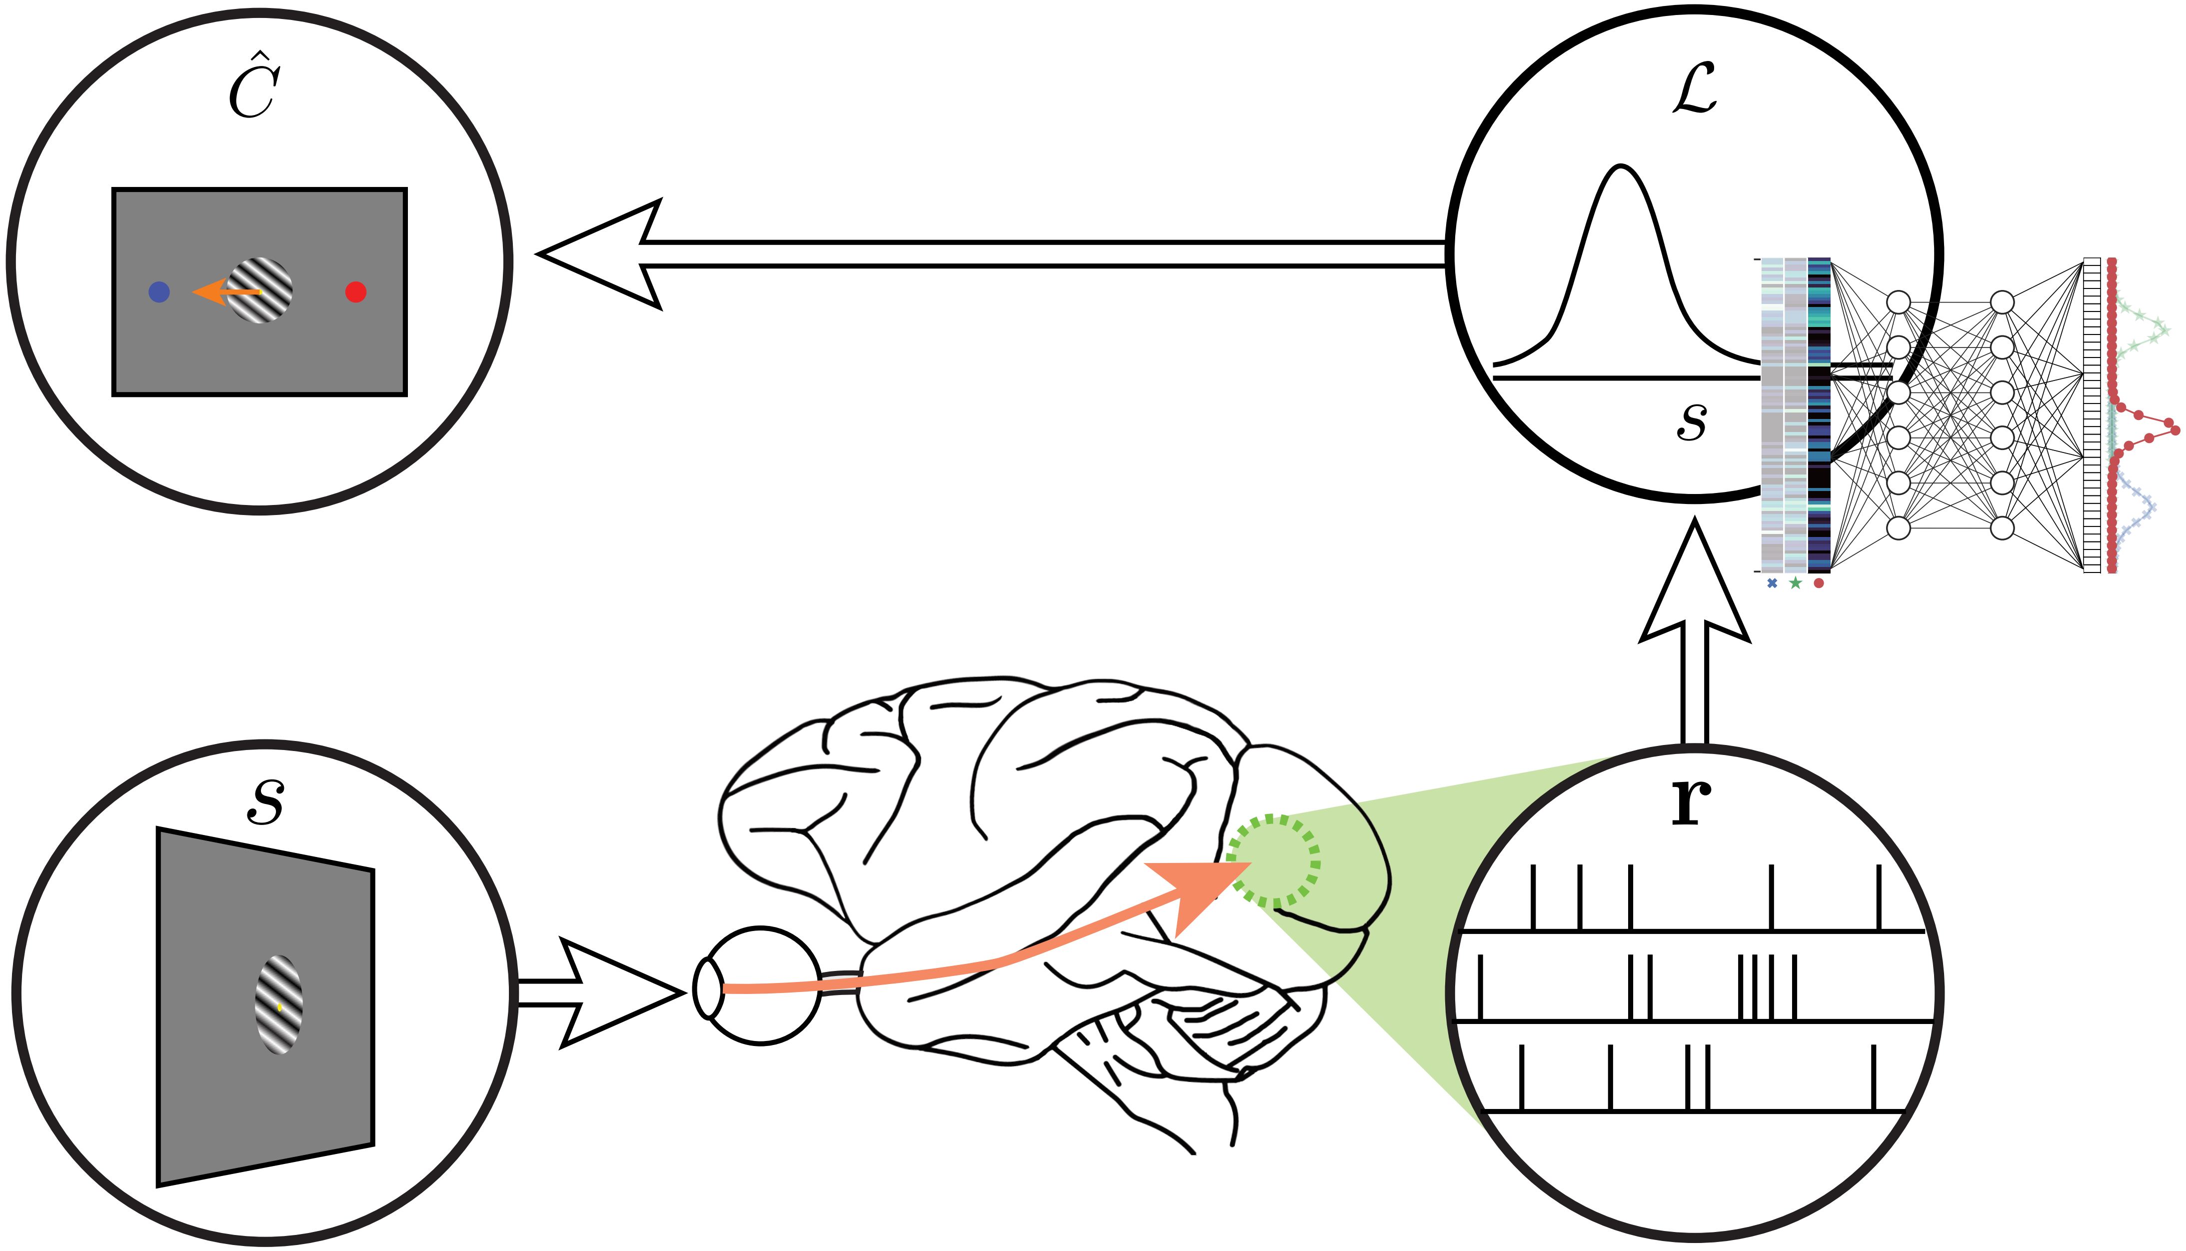


***Both the center and the shape of likelihood function changes from trial to trial***

Walker, Cotton et al., 2020 *Nature Neuro*

# Alternative model (non-PPC)

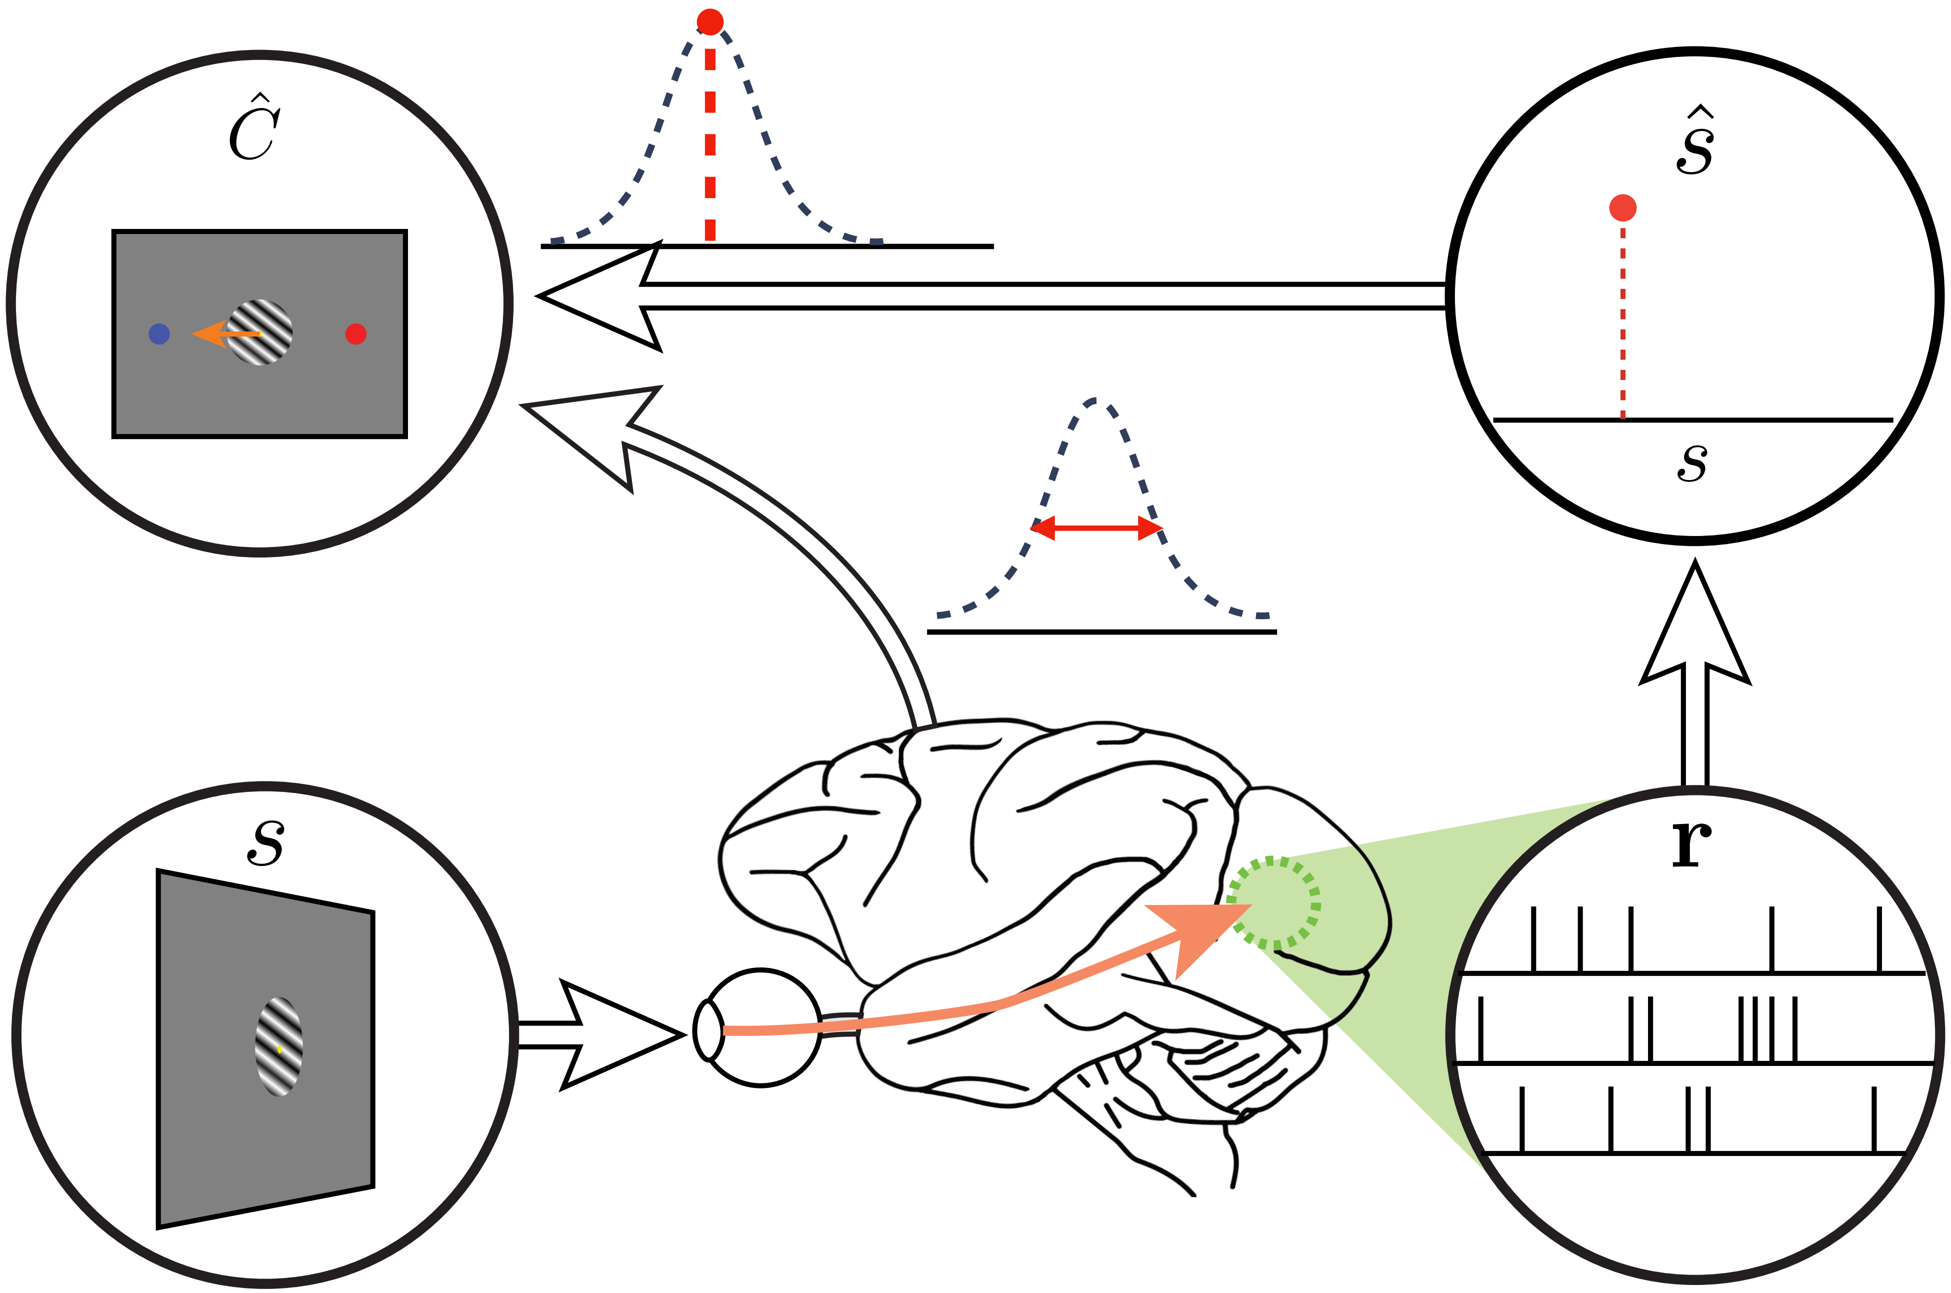


Walker, Cotton et al., 2020 *Nature Neuro*

# Alternative model (non-PPC)

- V1 population only provides **trial-by-trial best estimate** of the stimulus
- Uncertainty information is **NOT** represented by the shape of the likelihood function decoded from V1
- When contrast is fixed, the **uncertainty is expected to stay the same** across trials

*Fixed-Uncertainty model*

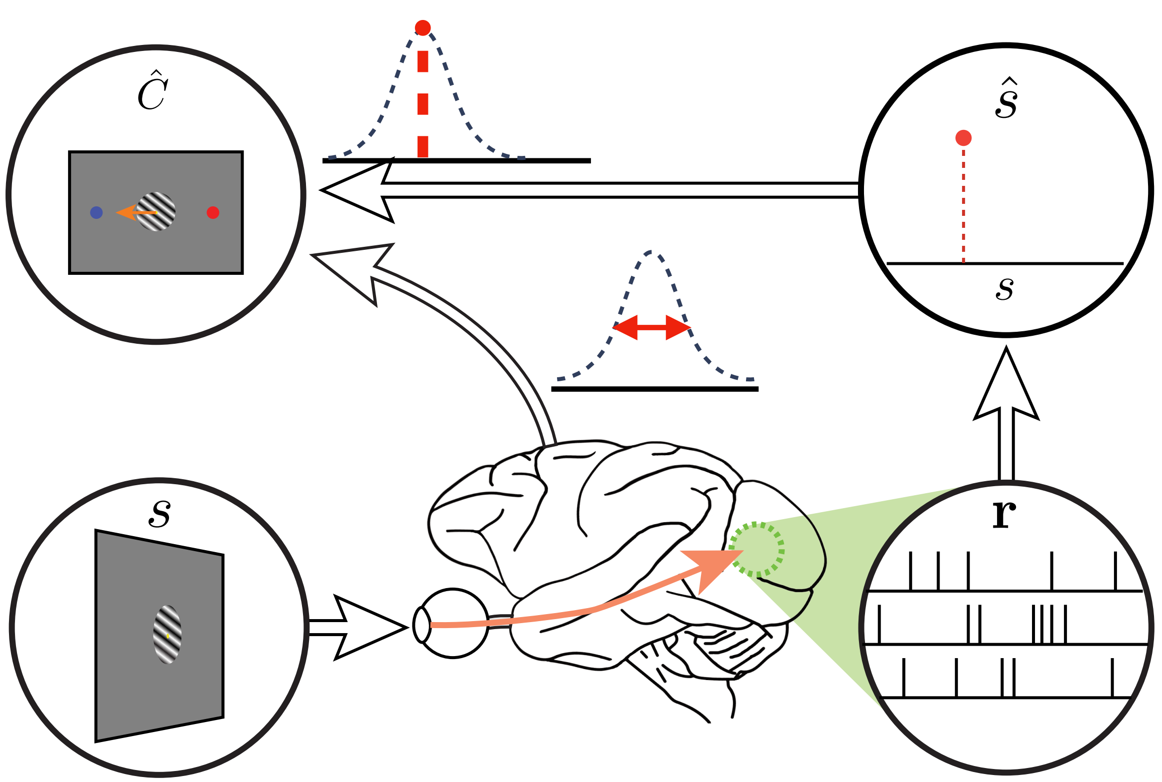

# (Fixed Uncertainty Model)

★

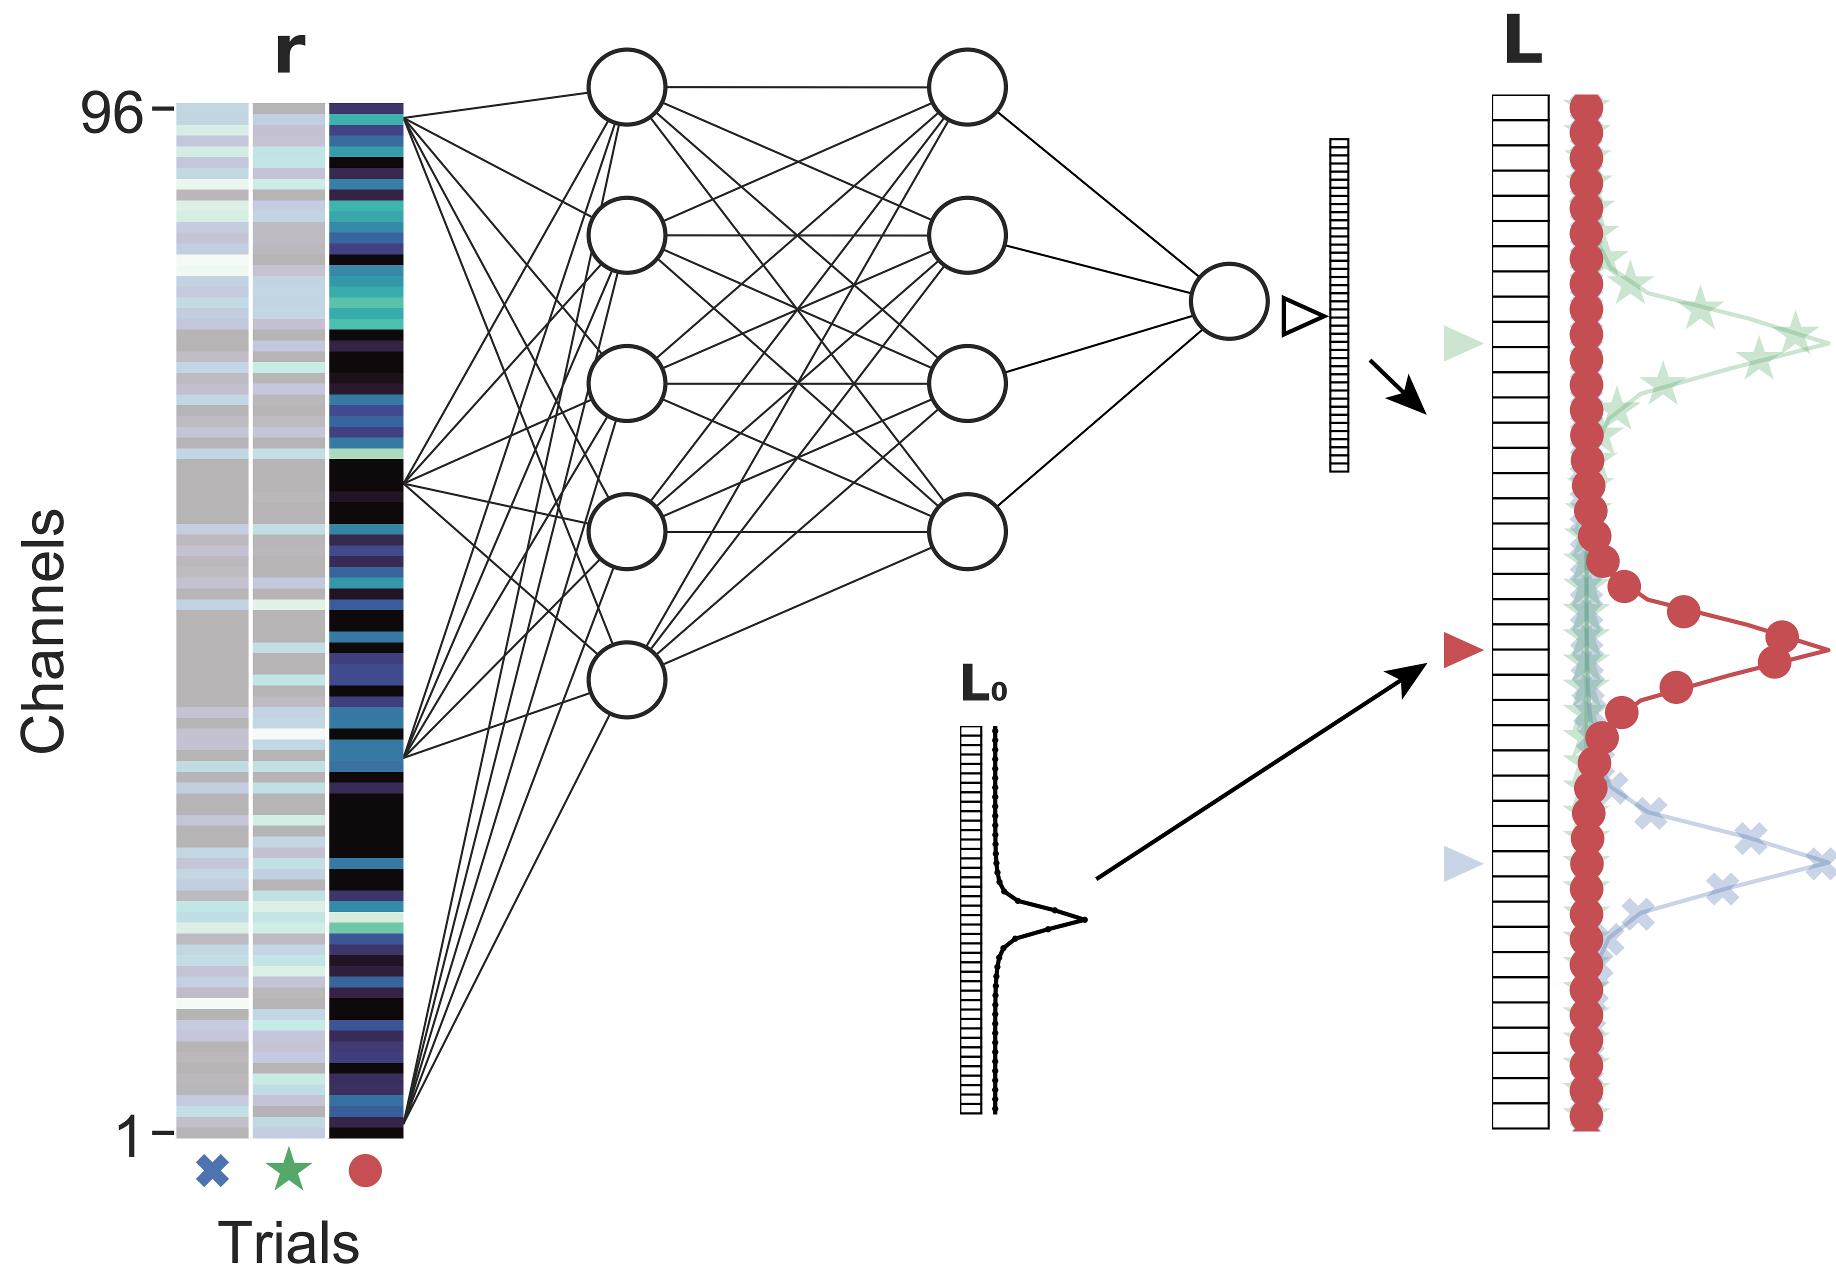


***Across trials from the same contrast, only the center of the likelihood function shifts and the shape remains the same***

Walker, Cotton et al., 2020 *Nature Neuro*

# For trials from the same contrast

## Full Likelihood Model (PPC)

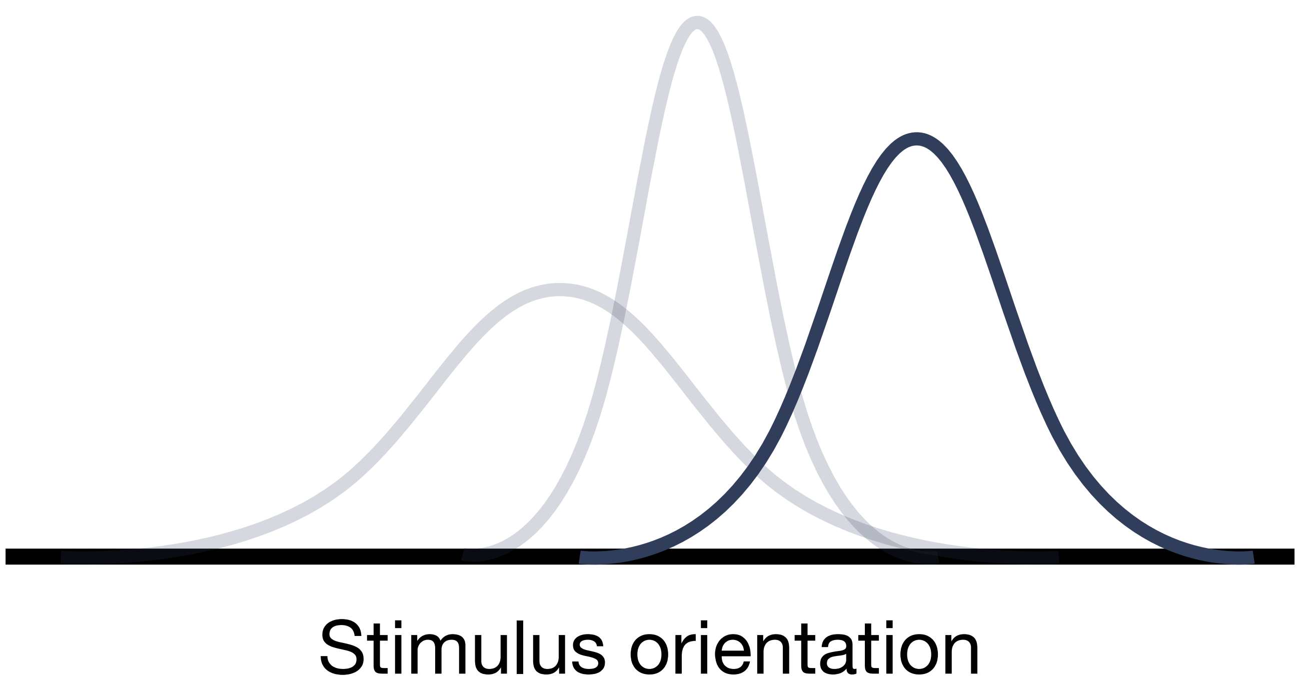


**Both *center and shape changes from trial to trial***

## Fixed Uncertainty Model (non PPC)

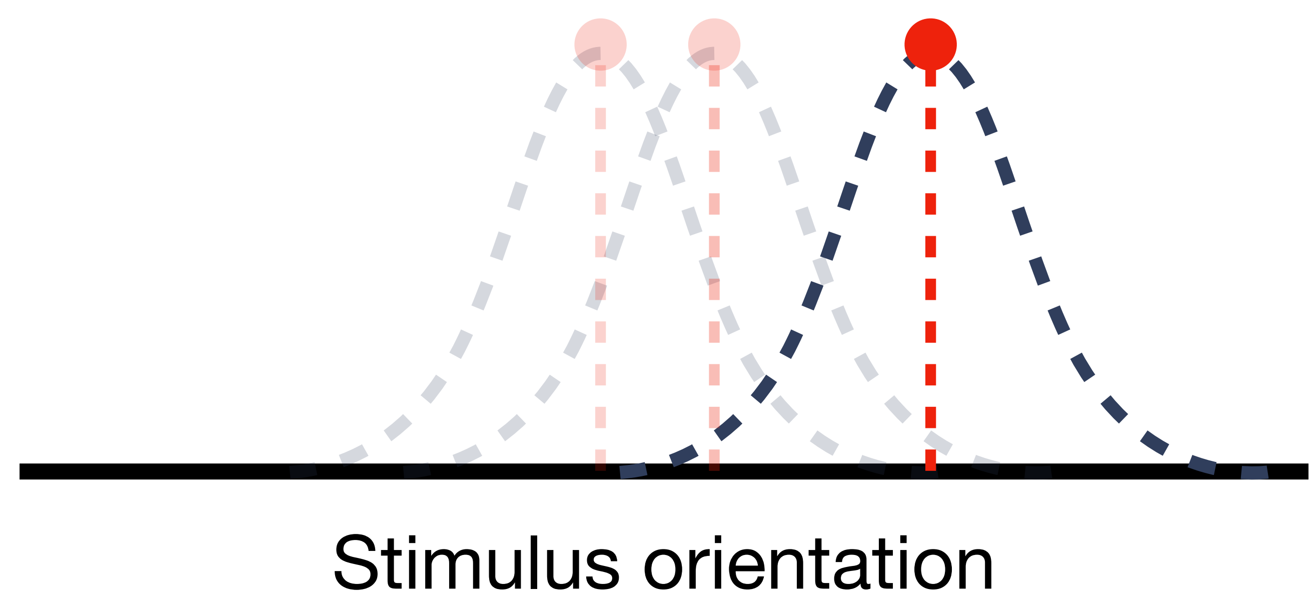


***Only center changes from trial to trial***

# Step 4: Model comparison

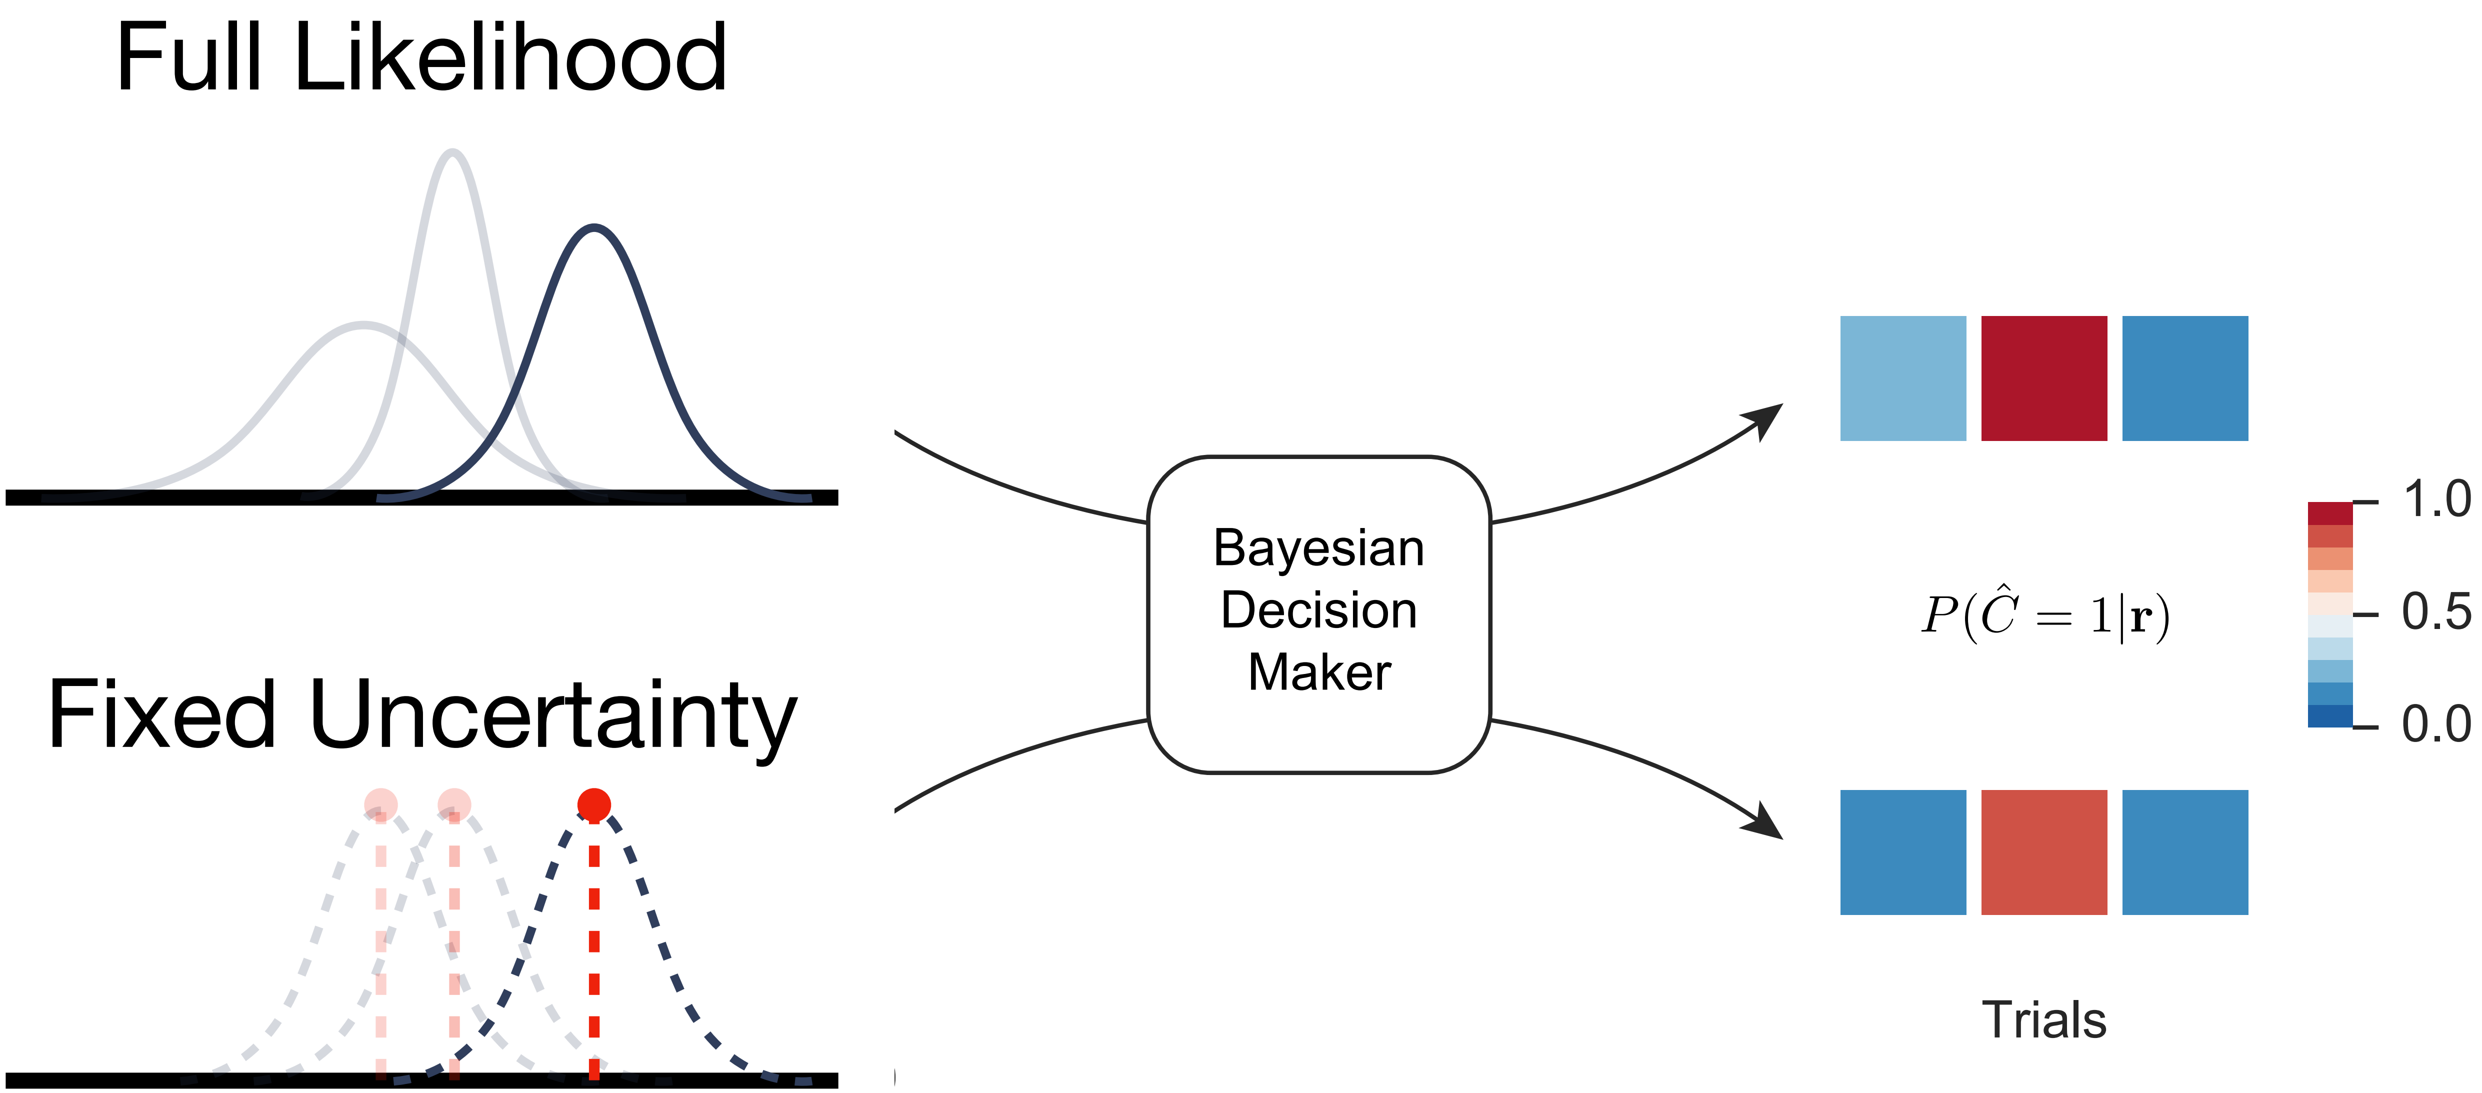


**Figure 3d**

***Compare trial-by-trial predictions of monkey's choices by the two models***

Walker, Cotton et al., 2020 *Nature Neuro*

# Model comparison results

Figure 5a

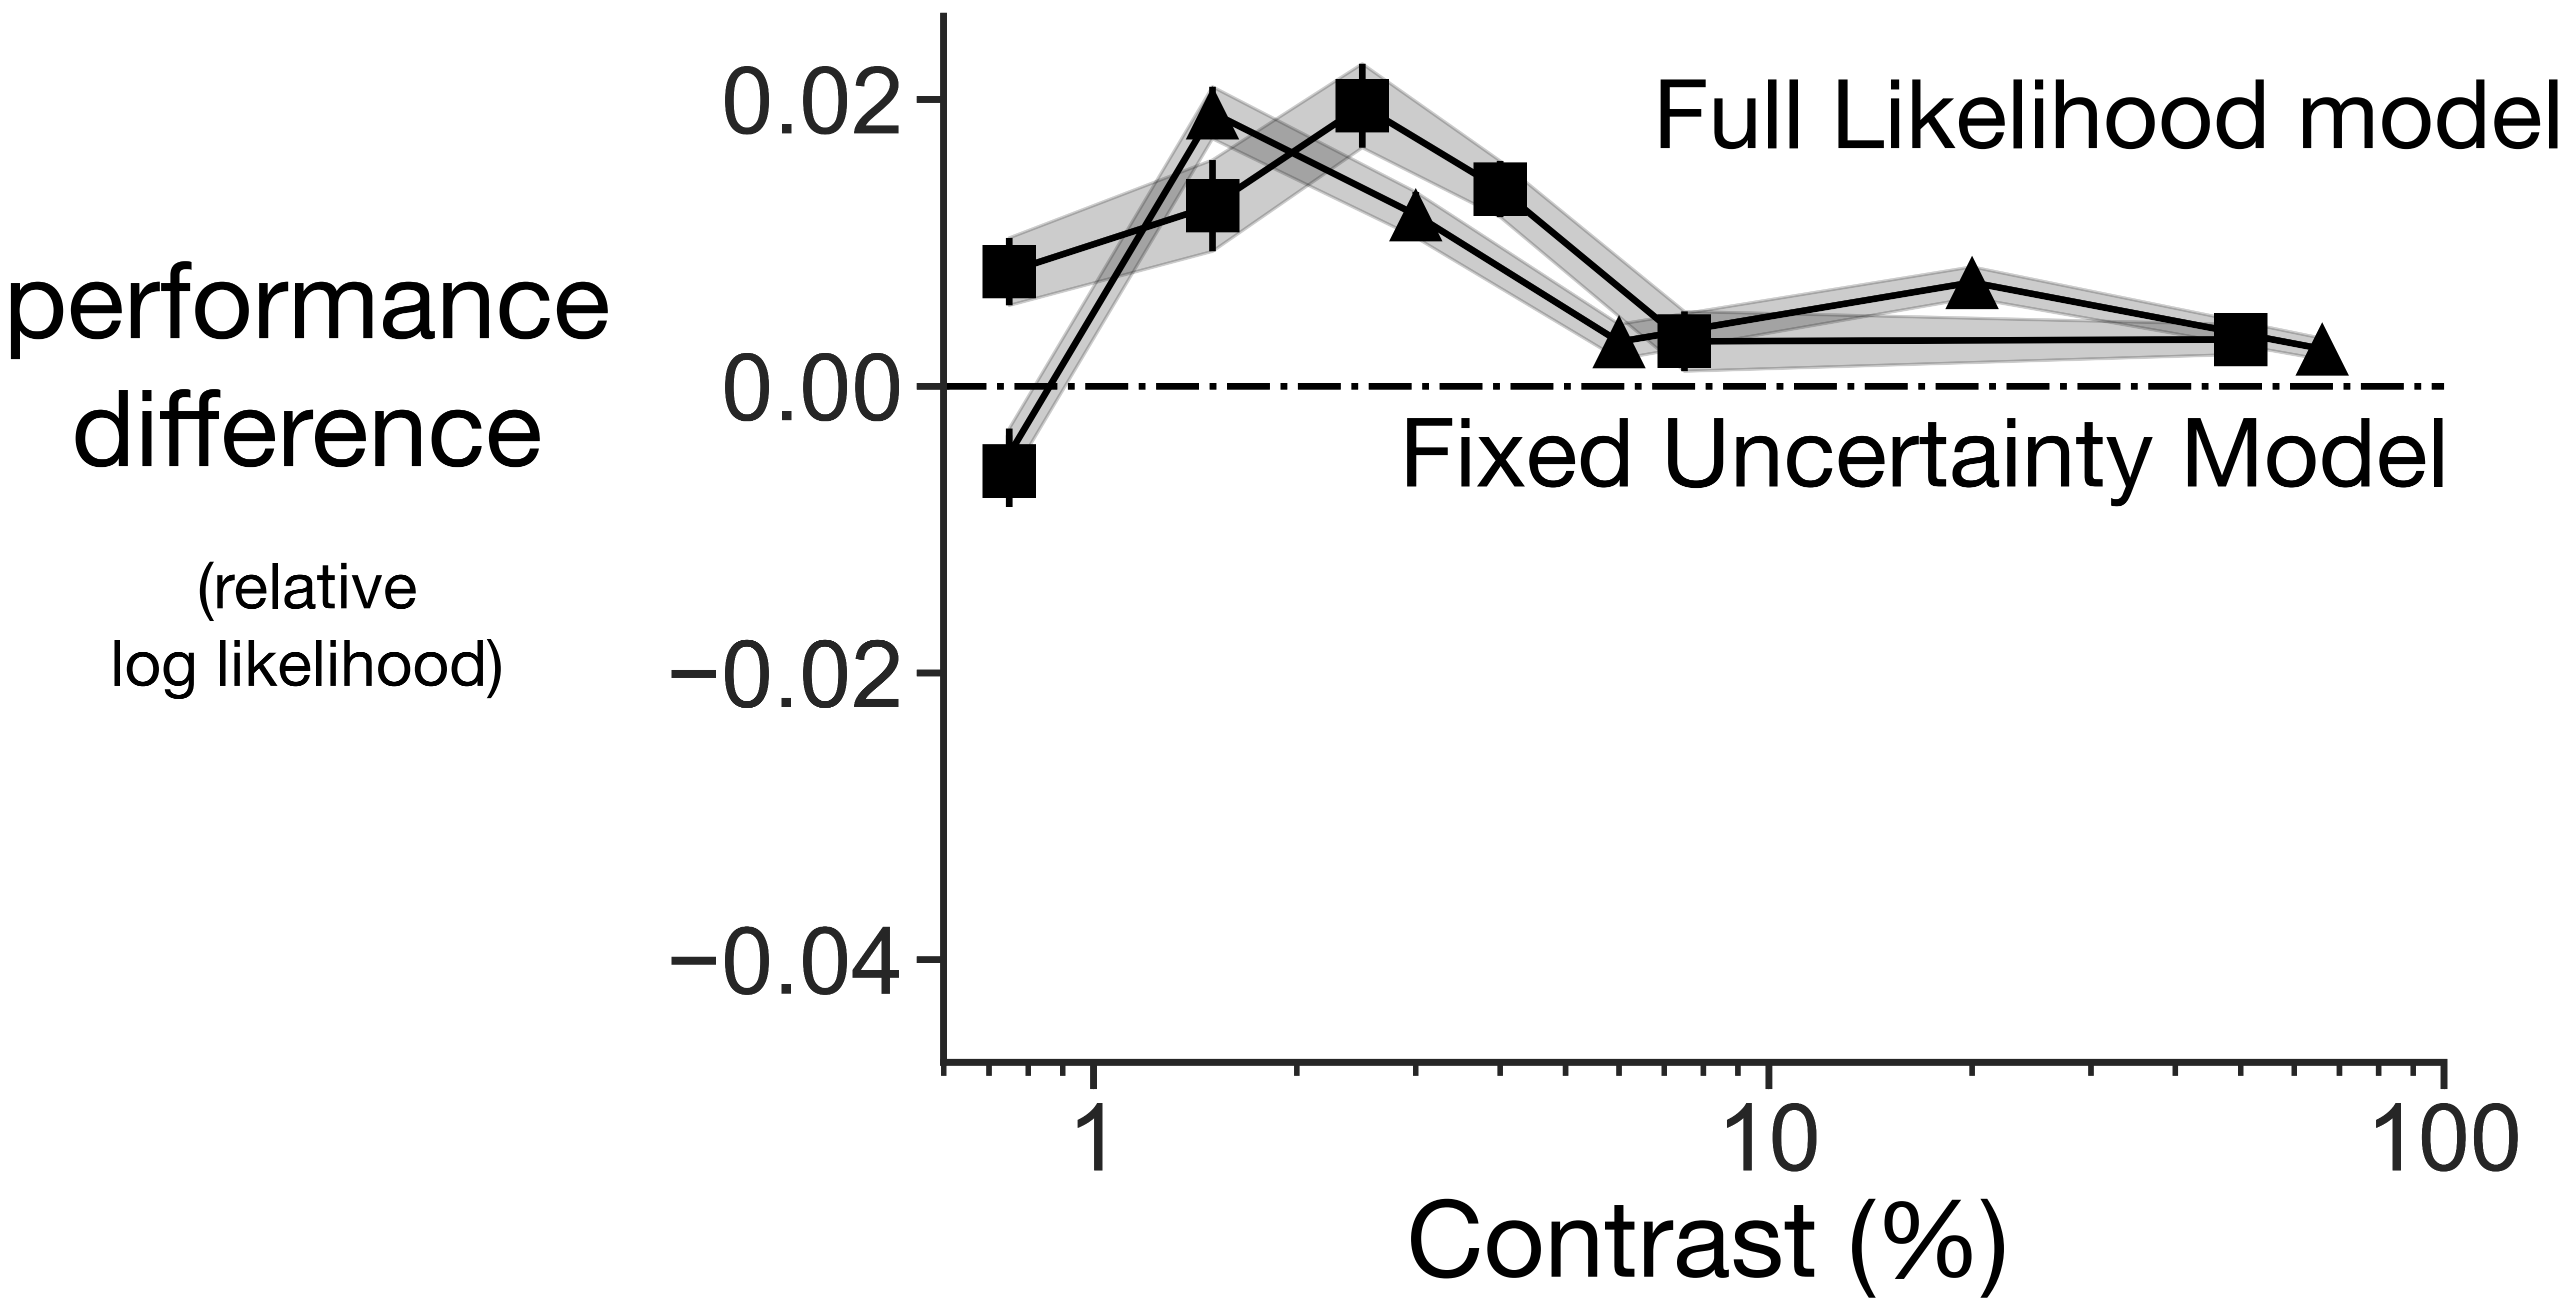


***Model using the full likelihood function predicts monkey's decisions better!***

Walker, Cotton et al., 2020 *Nature Neuro*

# Average likelihood shape varies with orientation

**Extended Figure 9c**

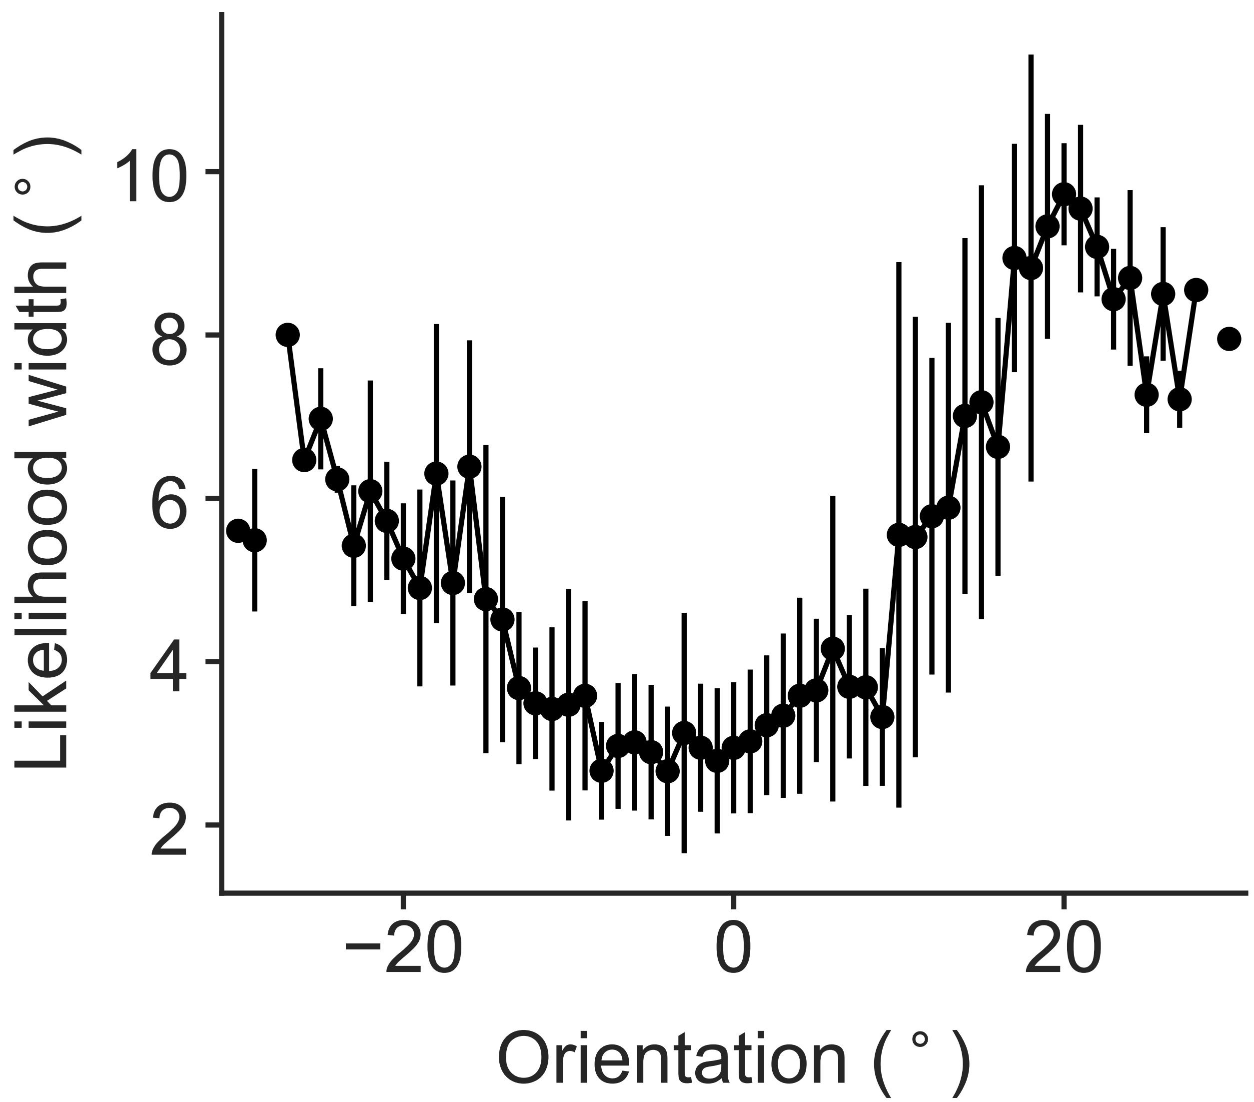


Walker, Cotton et al., 2020 *Nature Neuro*

# Shuffling of likelihood shapes conditioned on the stimulus

**Figure 5c**

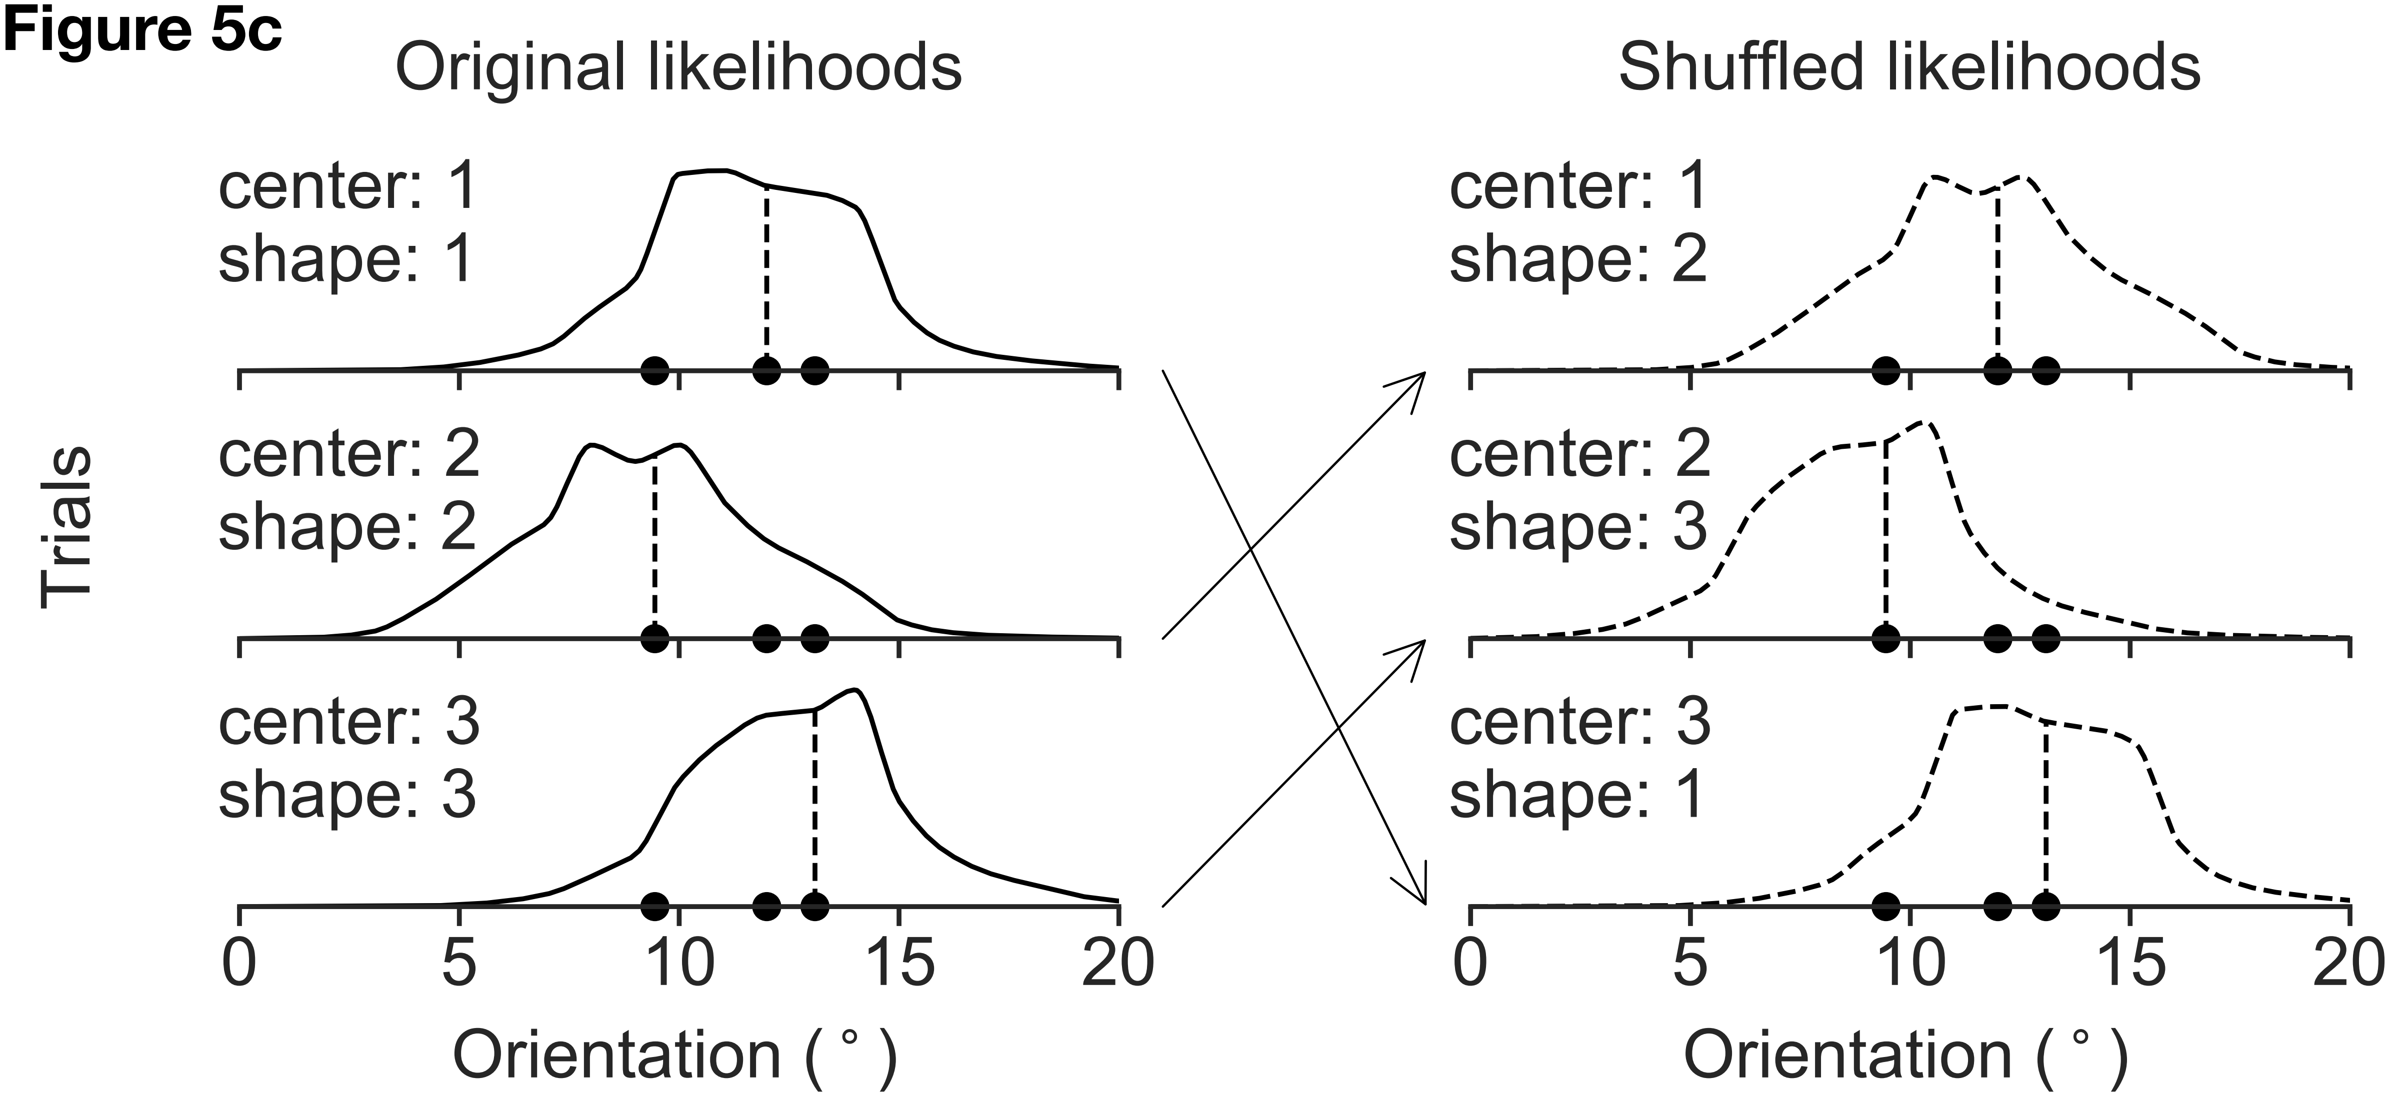


Walker, Cotton et al., 2020 *Nature Neuro*

# (Shuffling preserves stimulus-conditioned expected likelihood shape)

**Extended Figure 9c**

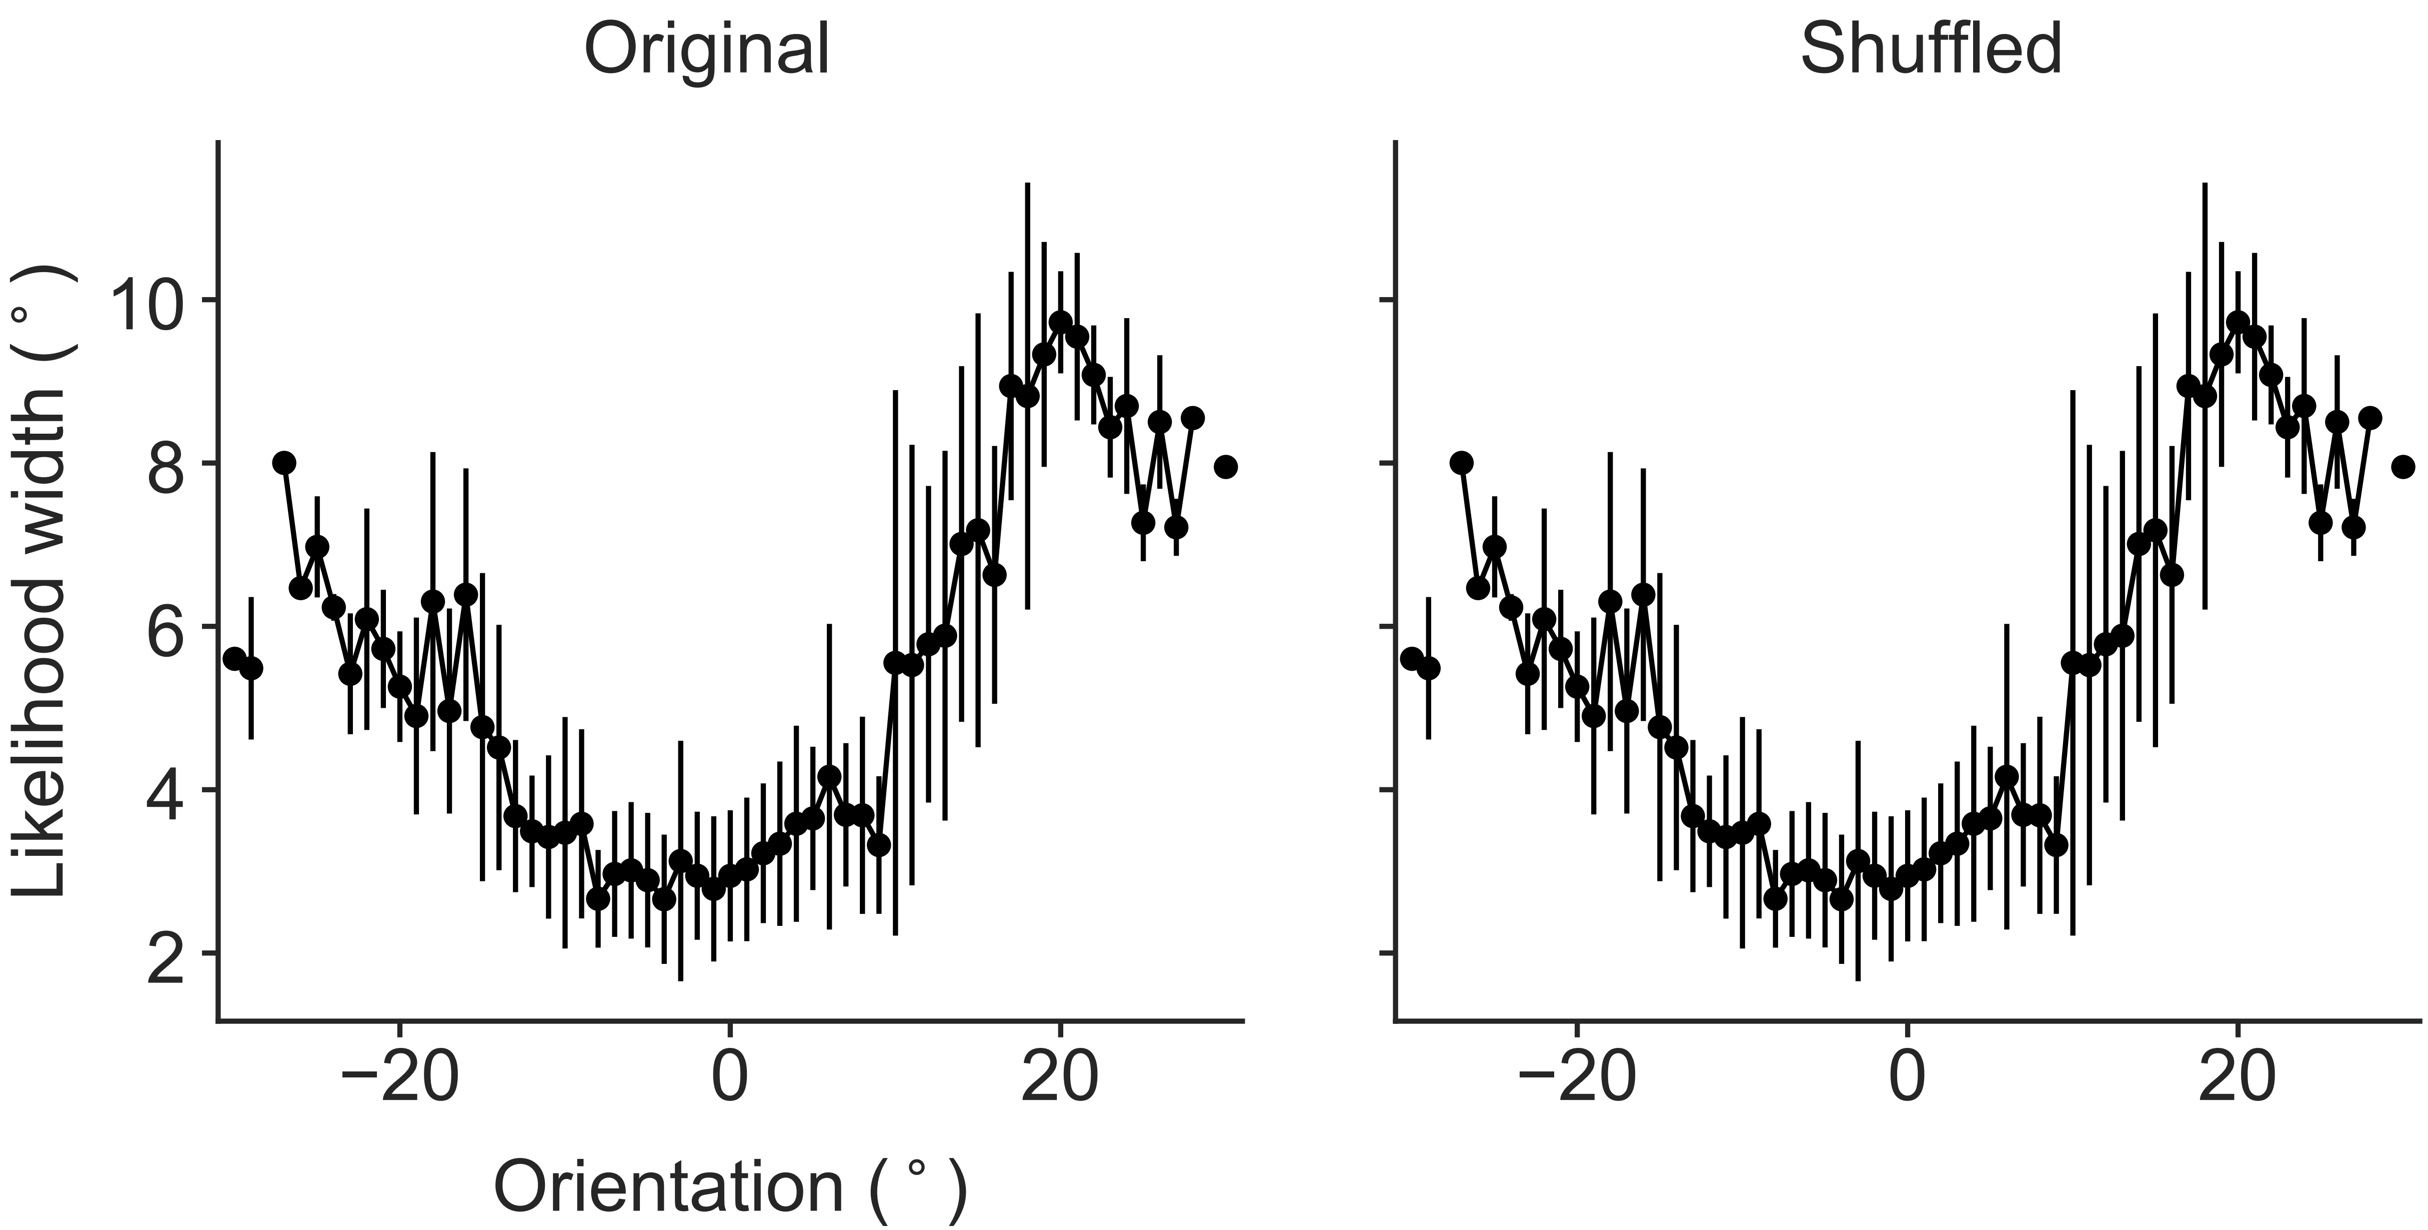


Walker, Cotton et al., 2020 *Nature Neuro*

# Shuffling shapes deteriorates performance

**Figure 5a**

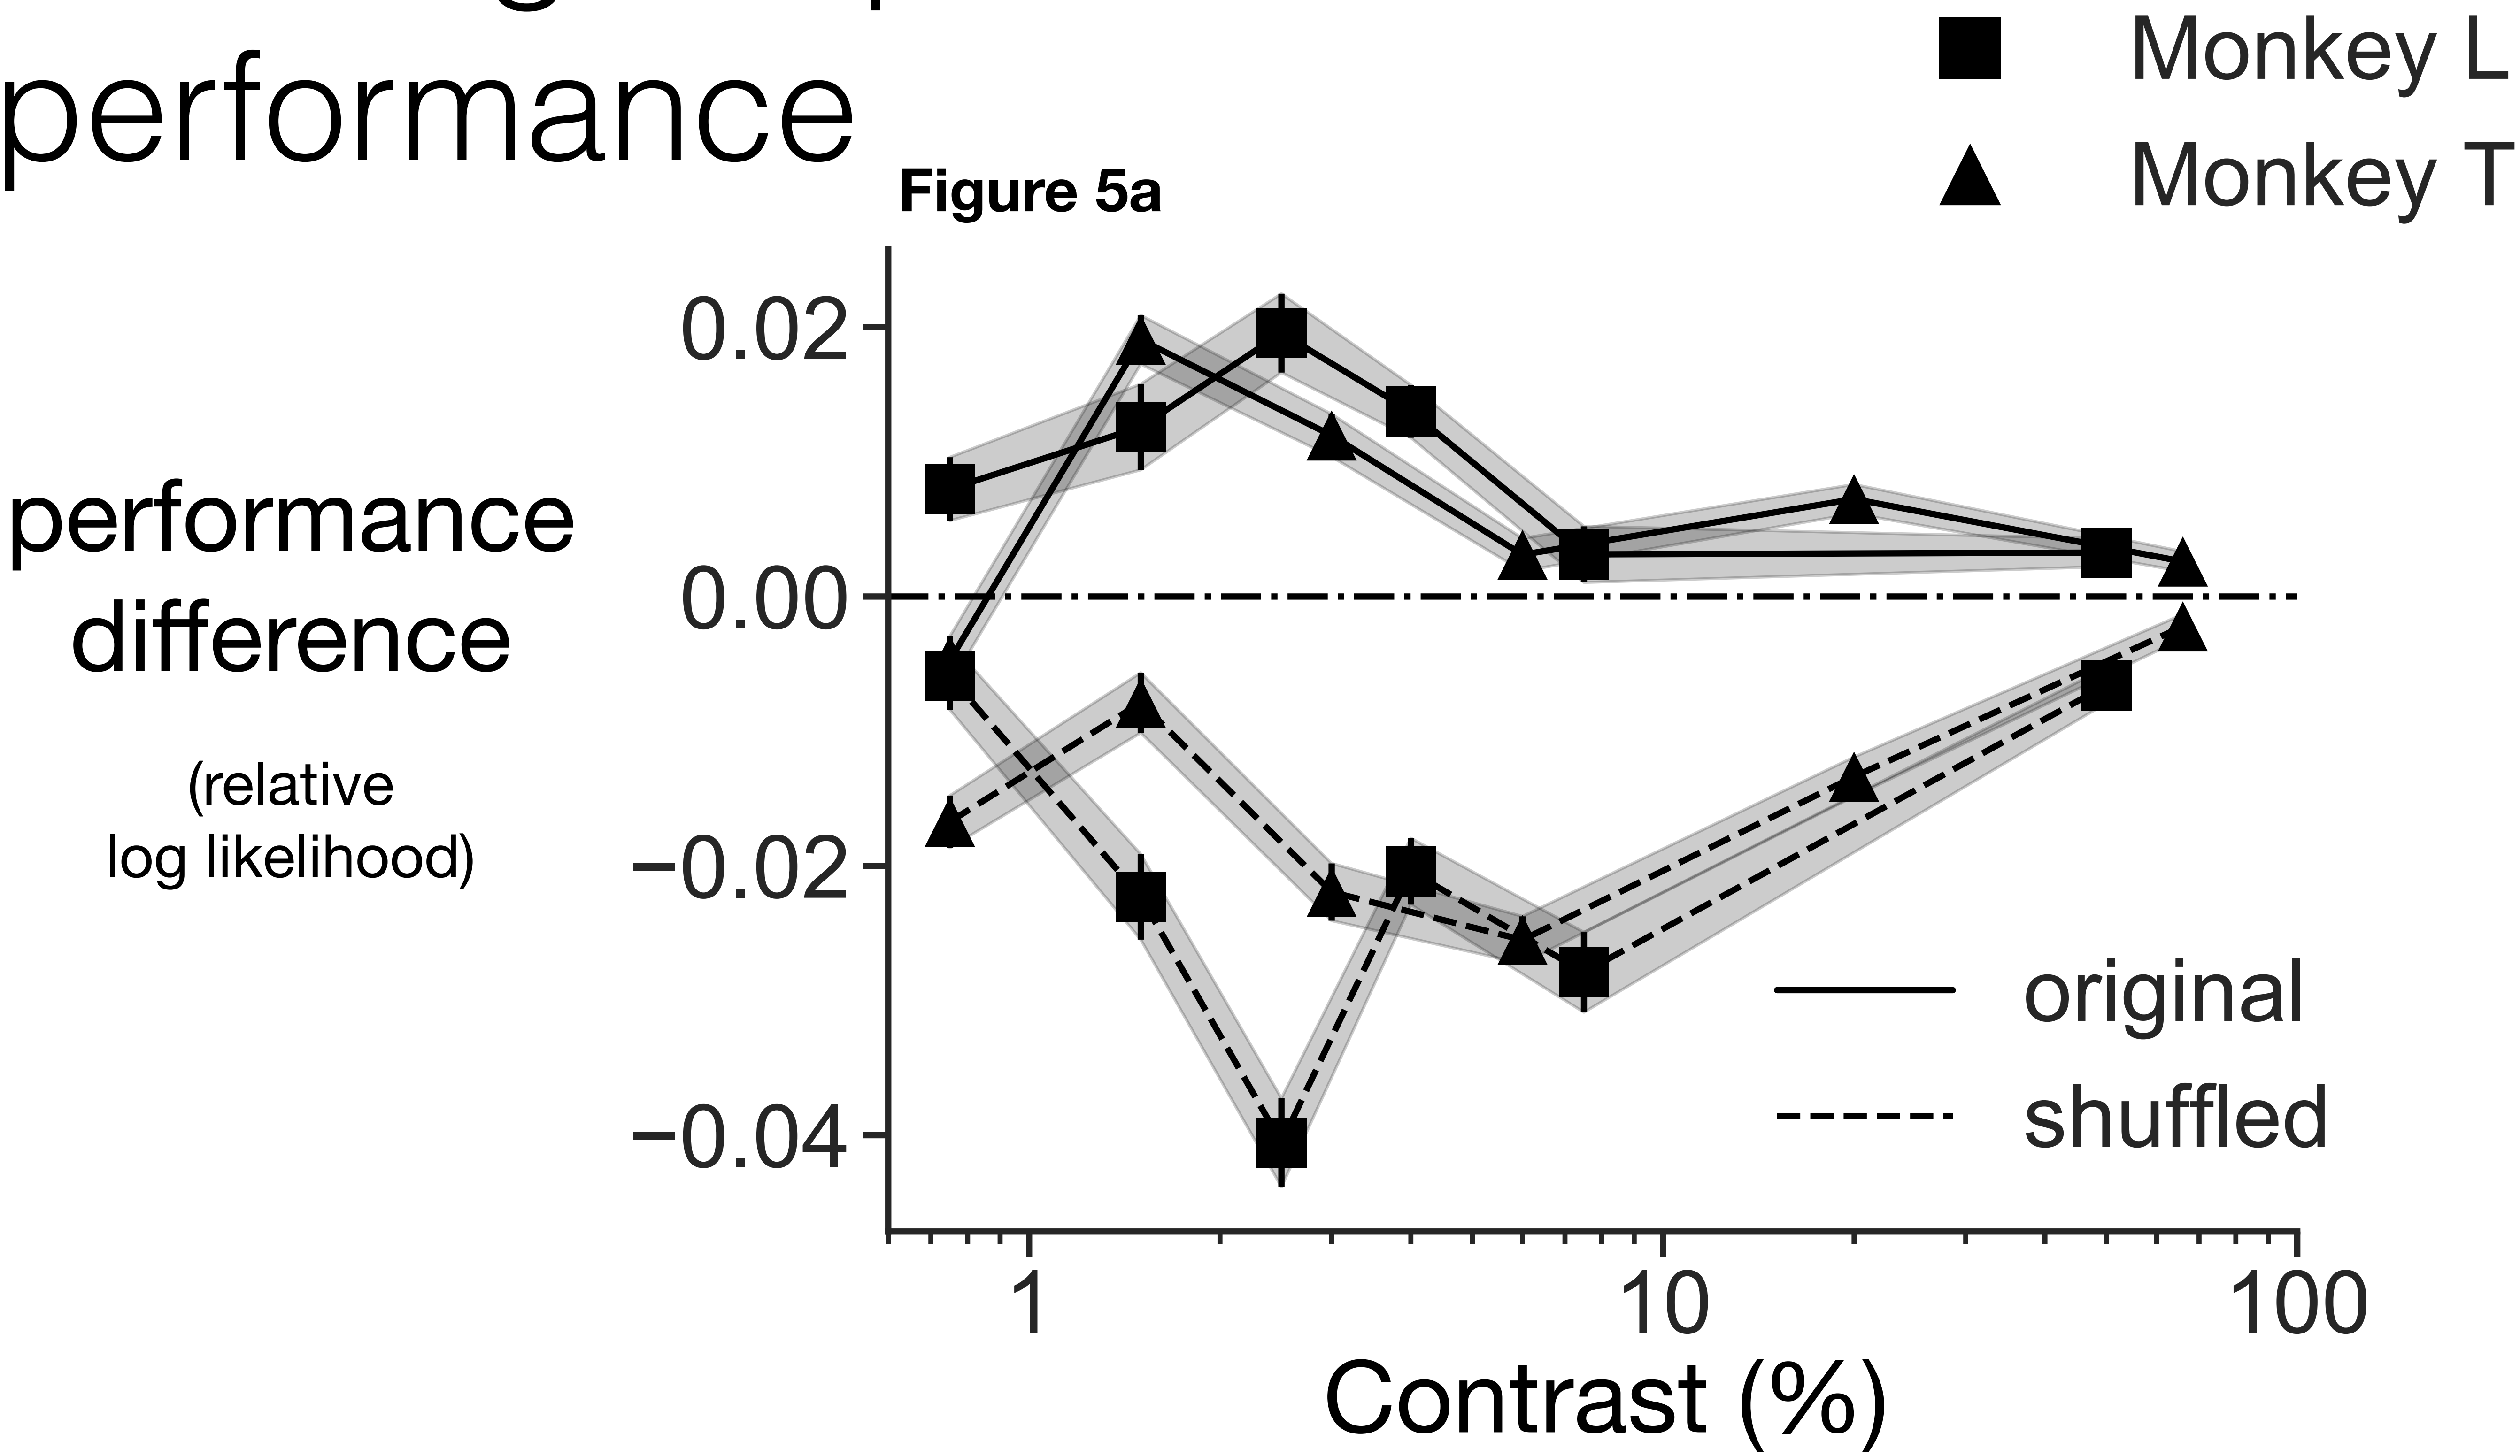


***Shuffling of likelihood function shape kills the benefit of Full-Likelihood Model, as expected***

# Summary

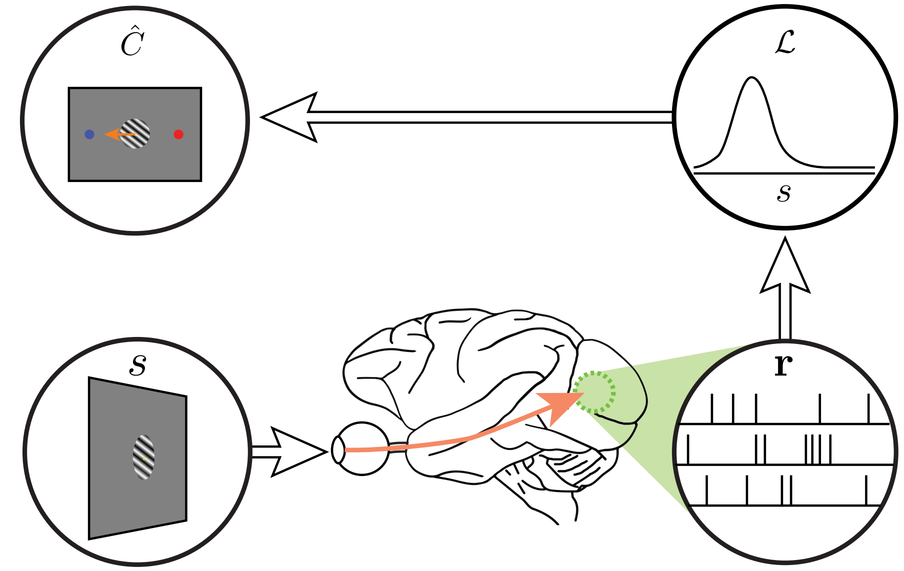


**Results:**

- Trial-by-trial fluctuations in the shape of likelihood function are informative about the monkey's trial-by-trial decisions on the task

**Significance:**

- **First population level electrophysiological evidence** in support of the hypothesis that the population of V1 neurons encode uncertainty in the form of likelihood function on trial-by-trial basis, **supporting PPC**

# Big thanks to all collaborators…

**Tolias Lab**
**Baylor College of Medicine**
**(Houston, TX)**

**Sinz Lab**
**University of Tübingen**
**(Tübingen, Germany)**

**Ma Lab**
**New York University**
**(New York, NY)**

# Bonus slides

# Recording from V1 population

**Figure 3b**

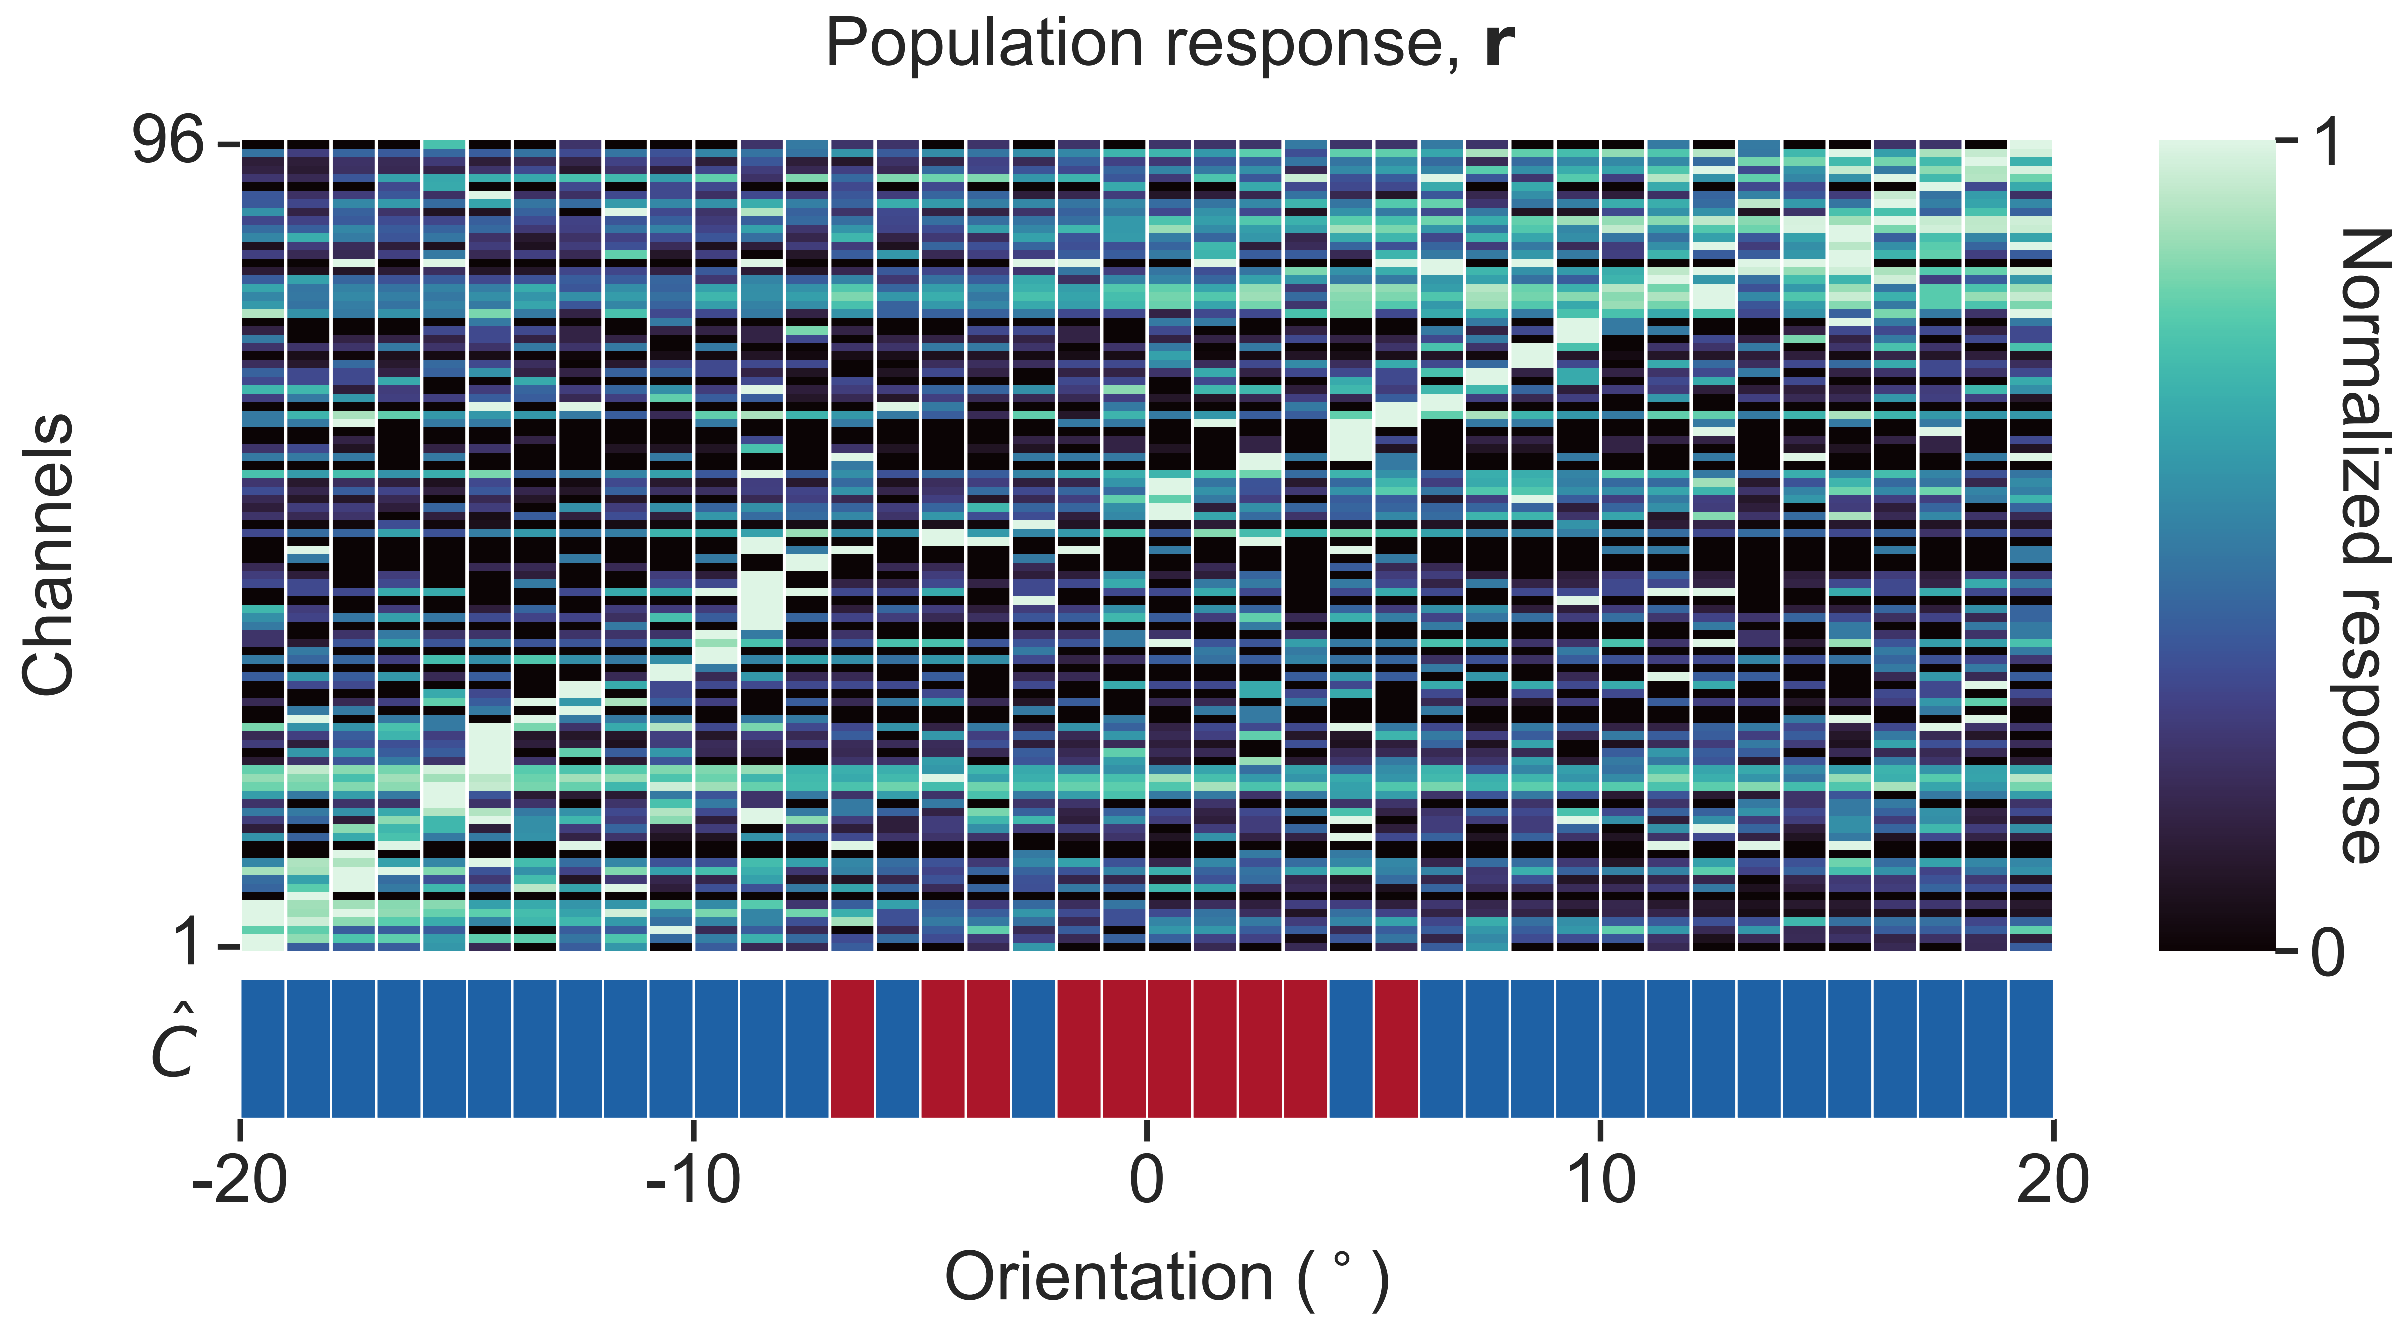


Walker, Cotton et al., 2020 *Nature Neuro*

# (Effectiveness of DNN likelihood decoder) ★

**Extended Figure 3b: Ground truth - independent Poisson distribution**

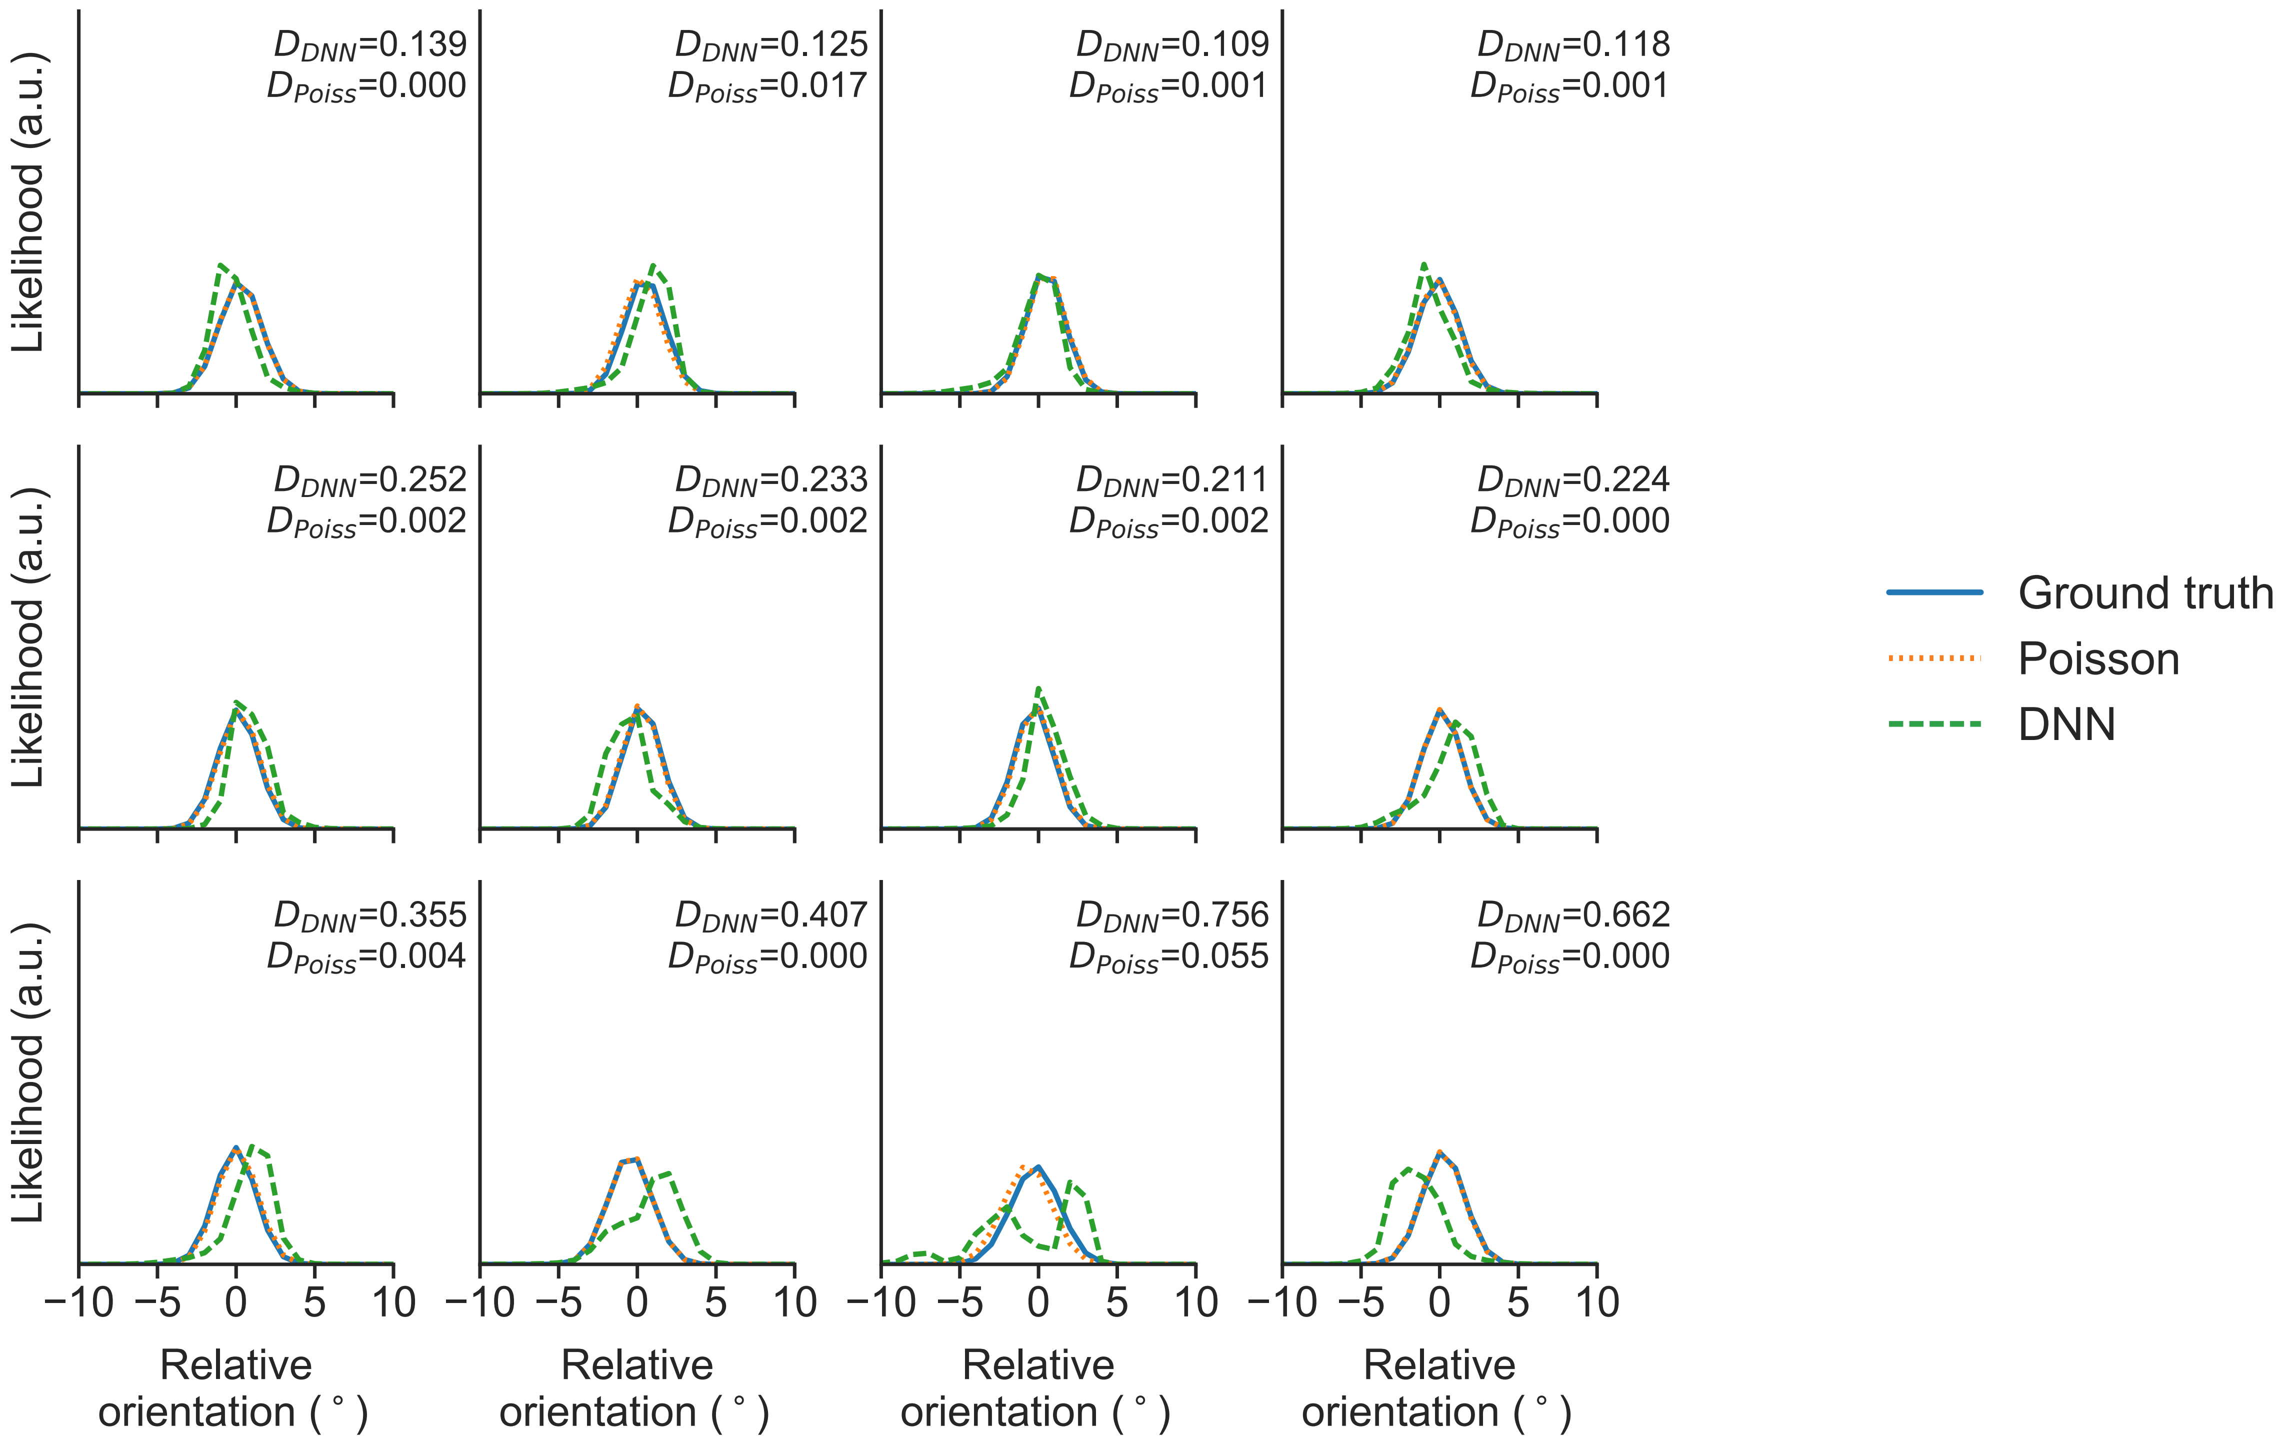

# (Effectiveness of DNN likelihood decoder)

**Extended Figure 3d: Ground truth - scaled correlated Gaussian distribution**

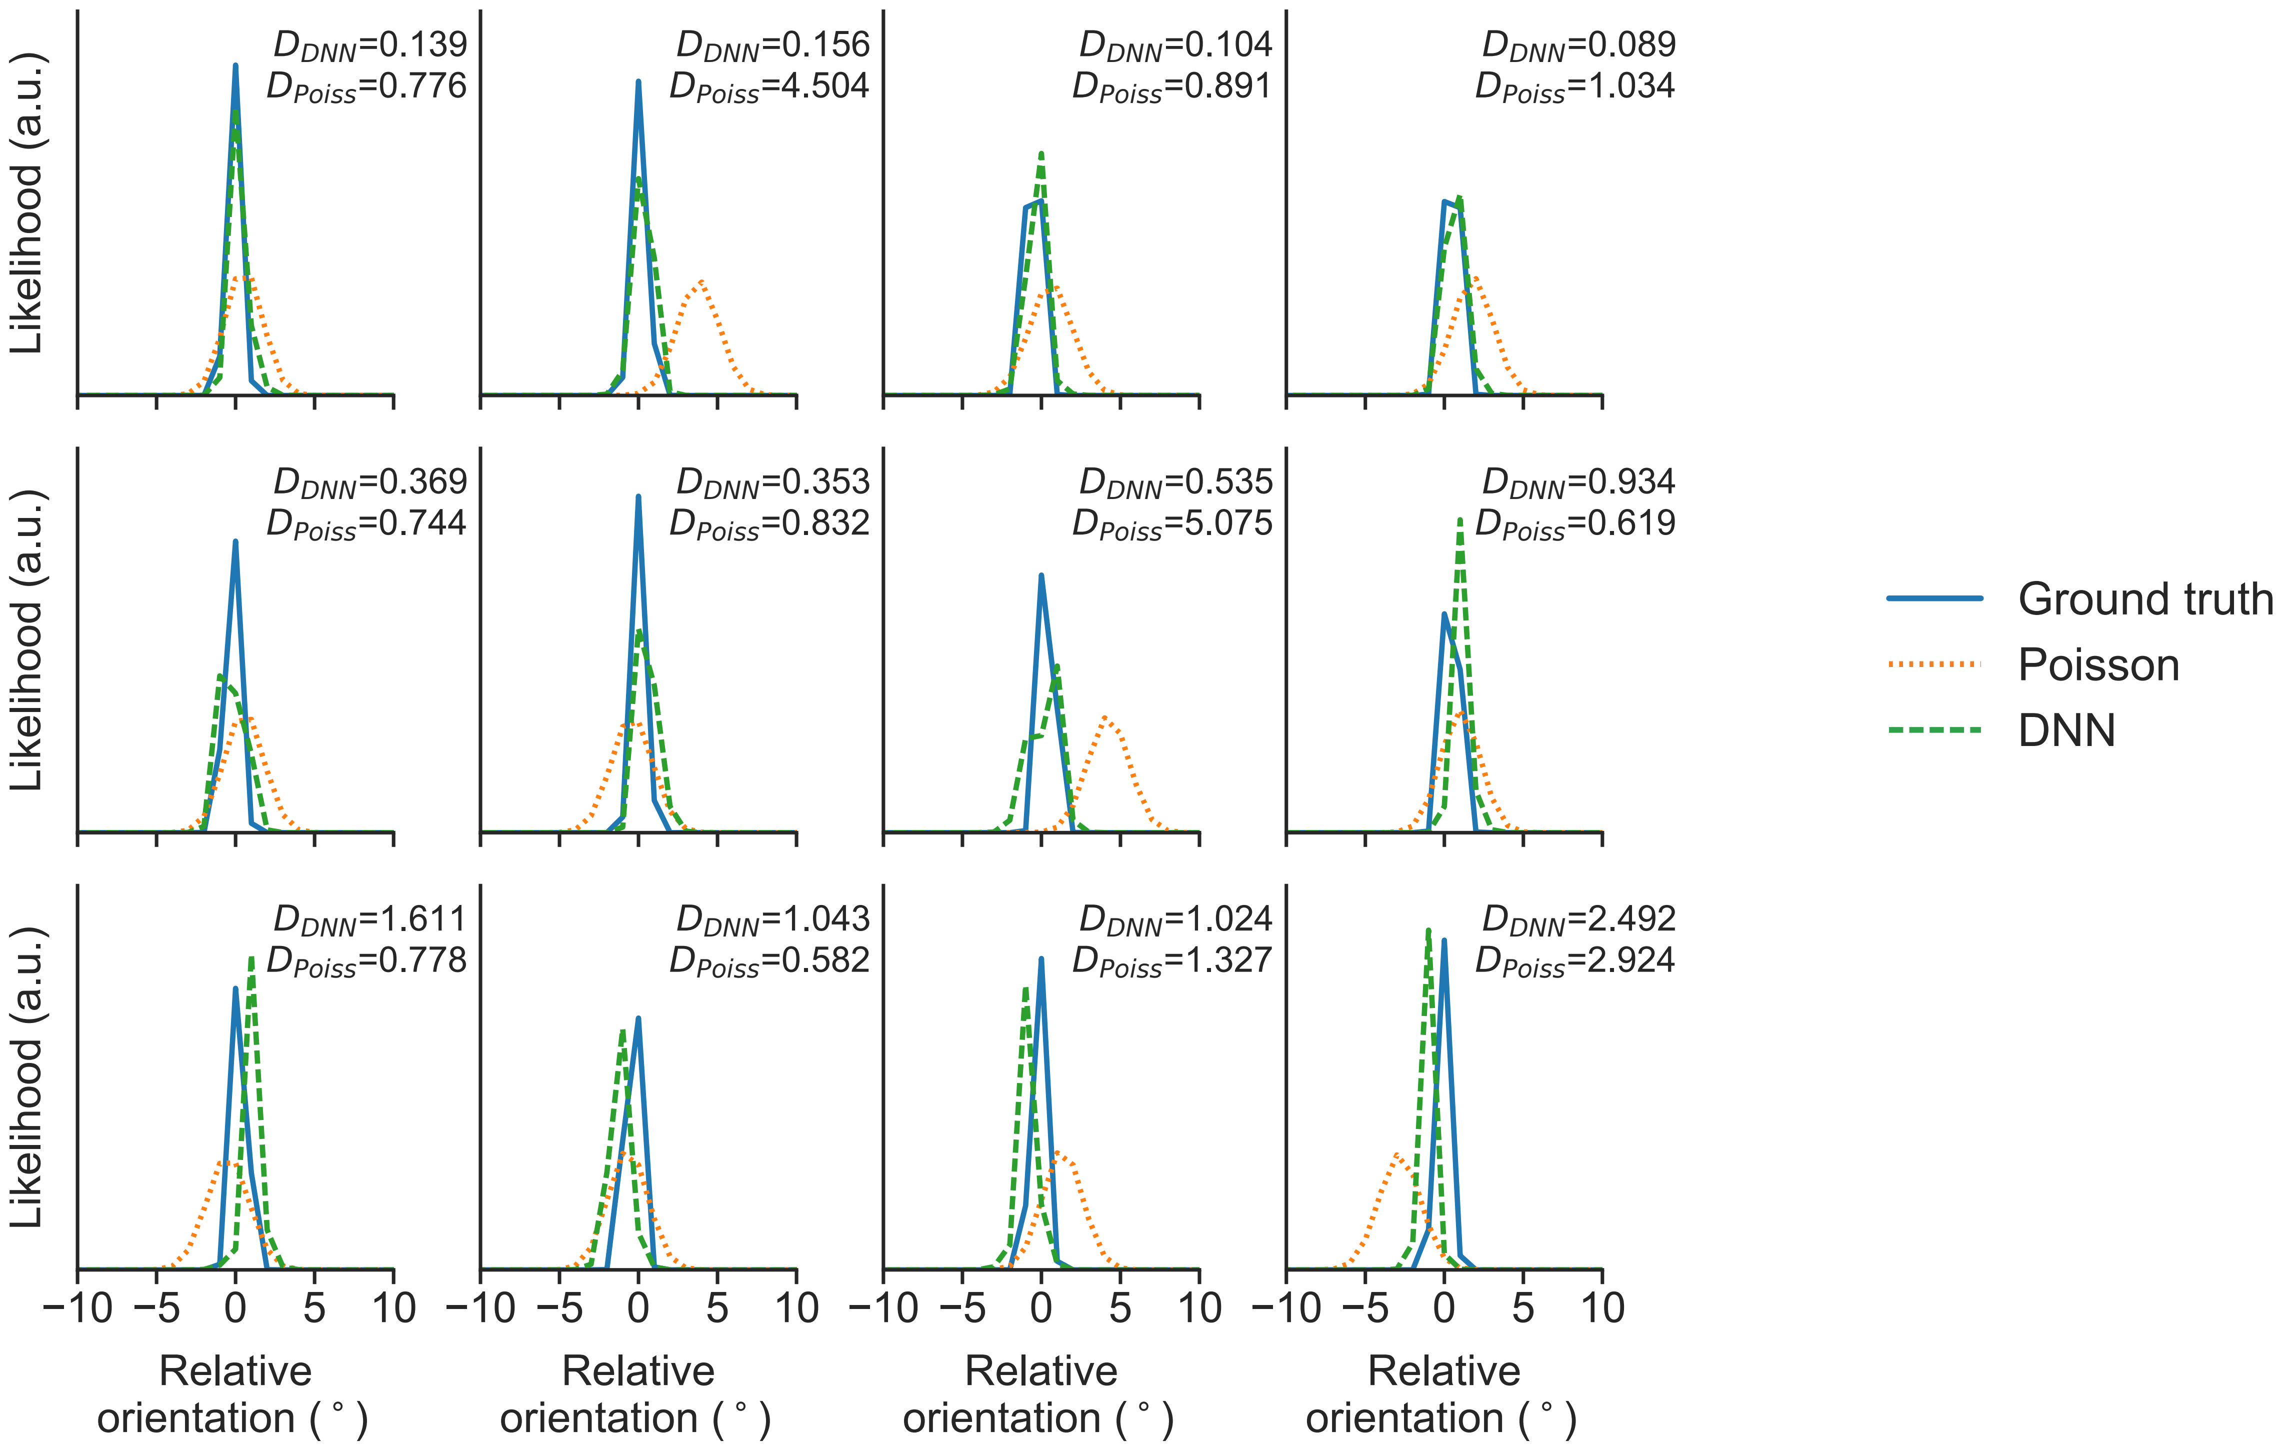

# Width of decoded likelihood varies with contrast

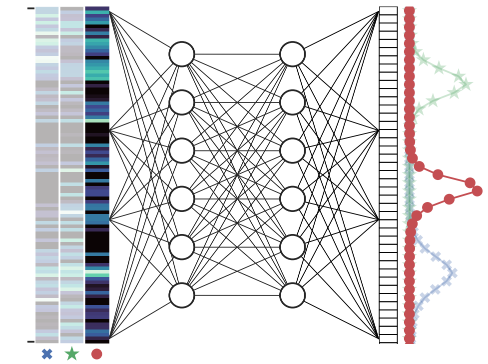

**Figure 4a—c**

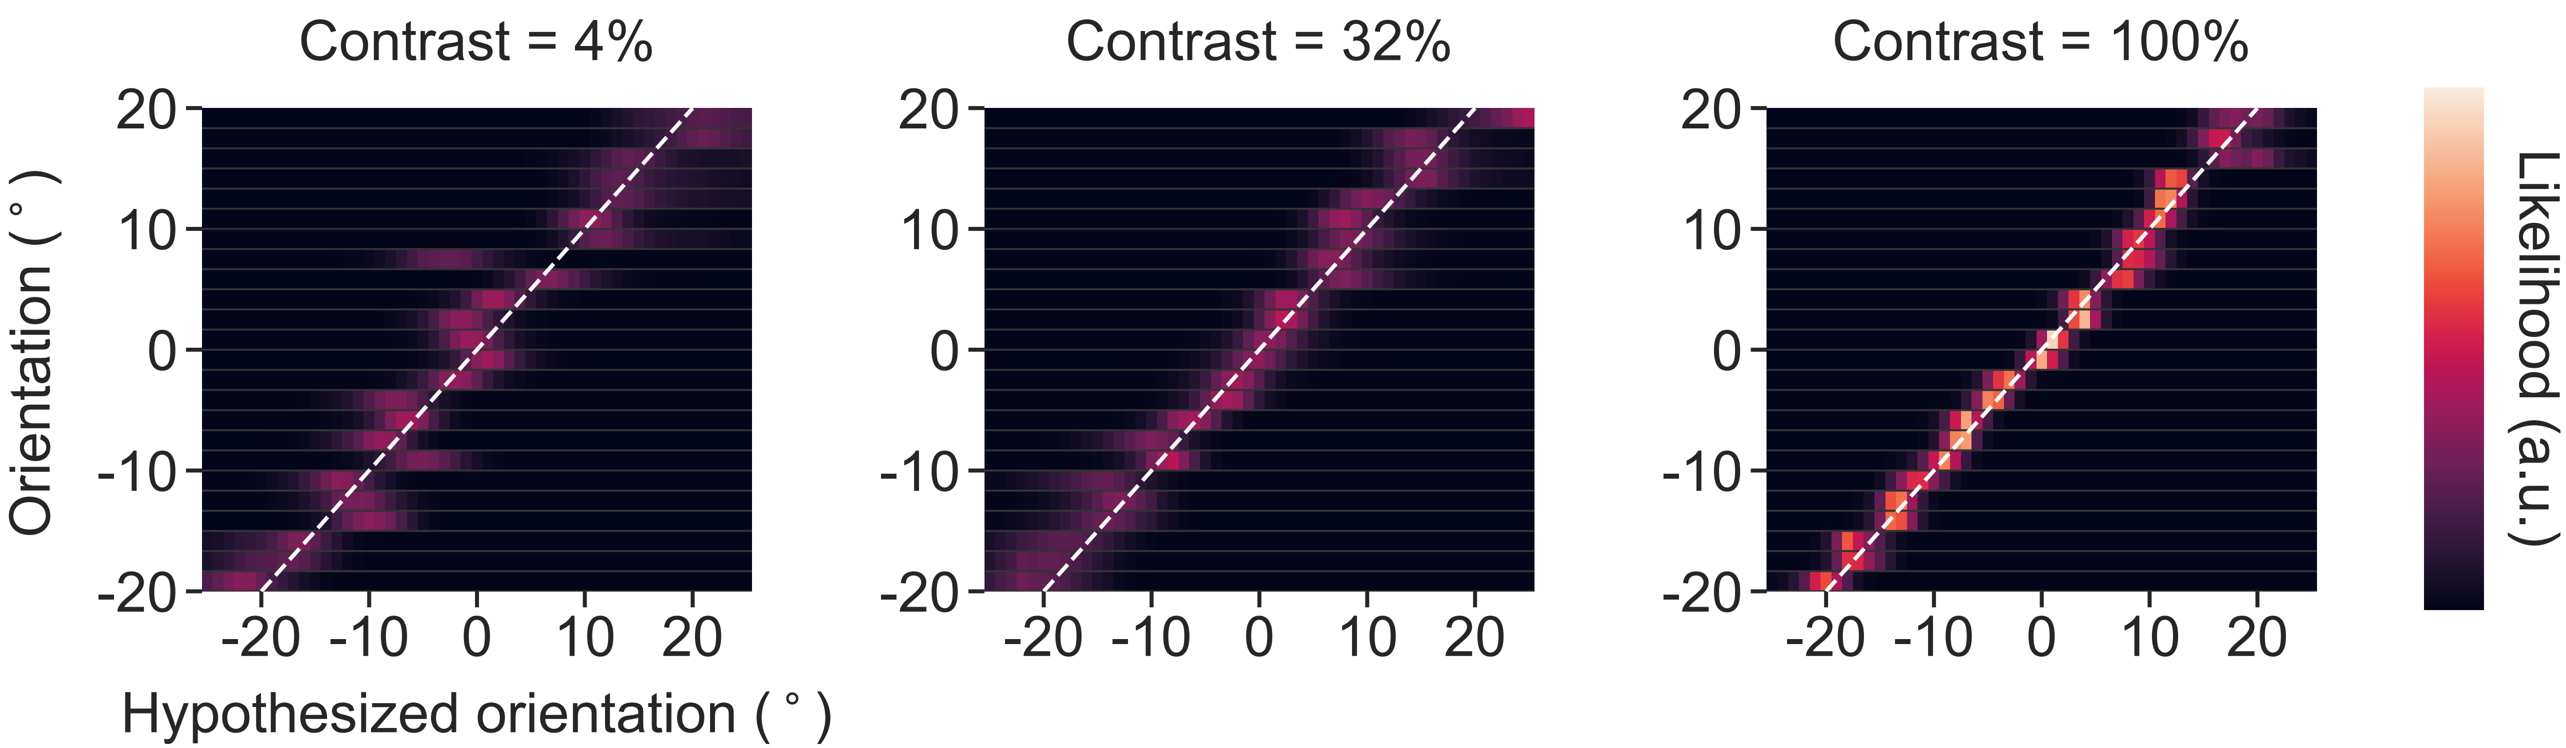


Walker, Cotton et al., 2020 *Nature Neuro*

# Width of decoded likelihood varies with contrast

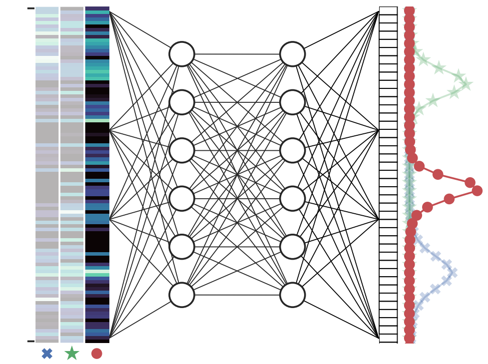

**Figure 4d**

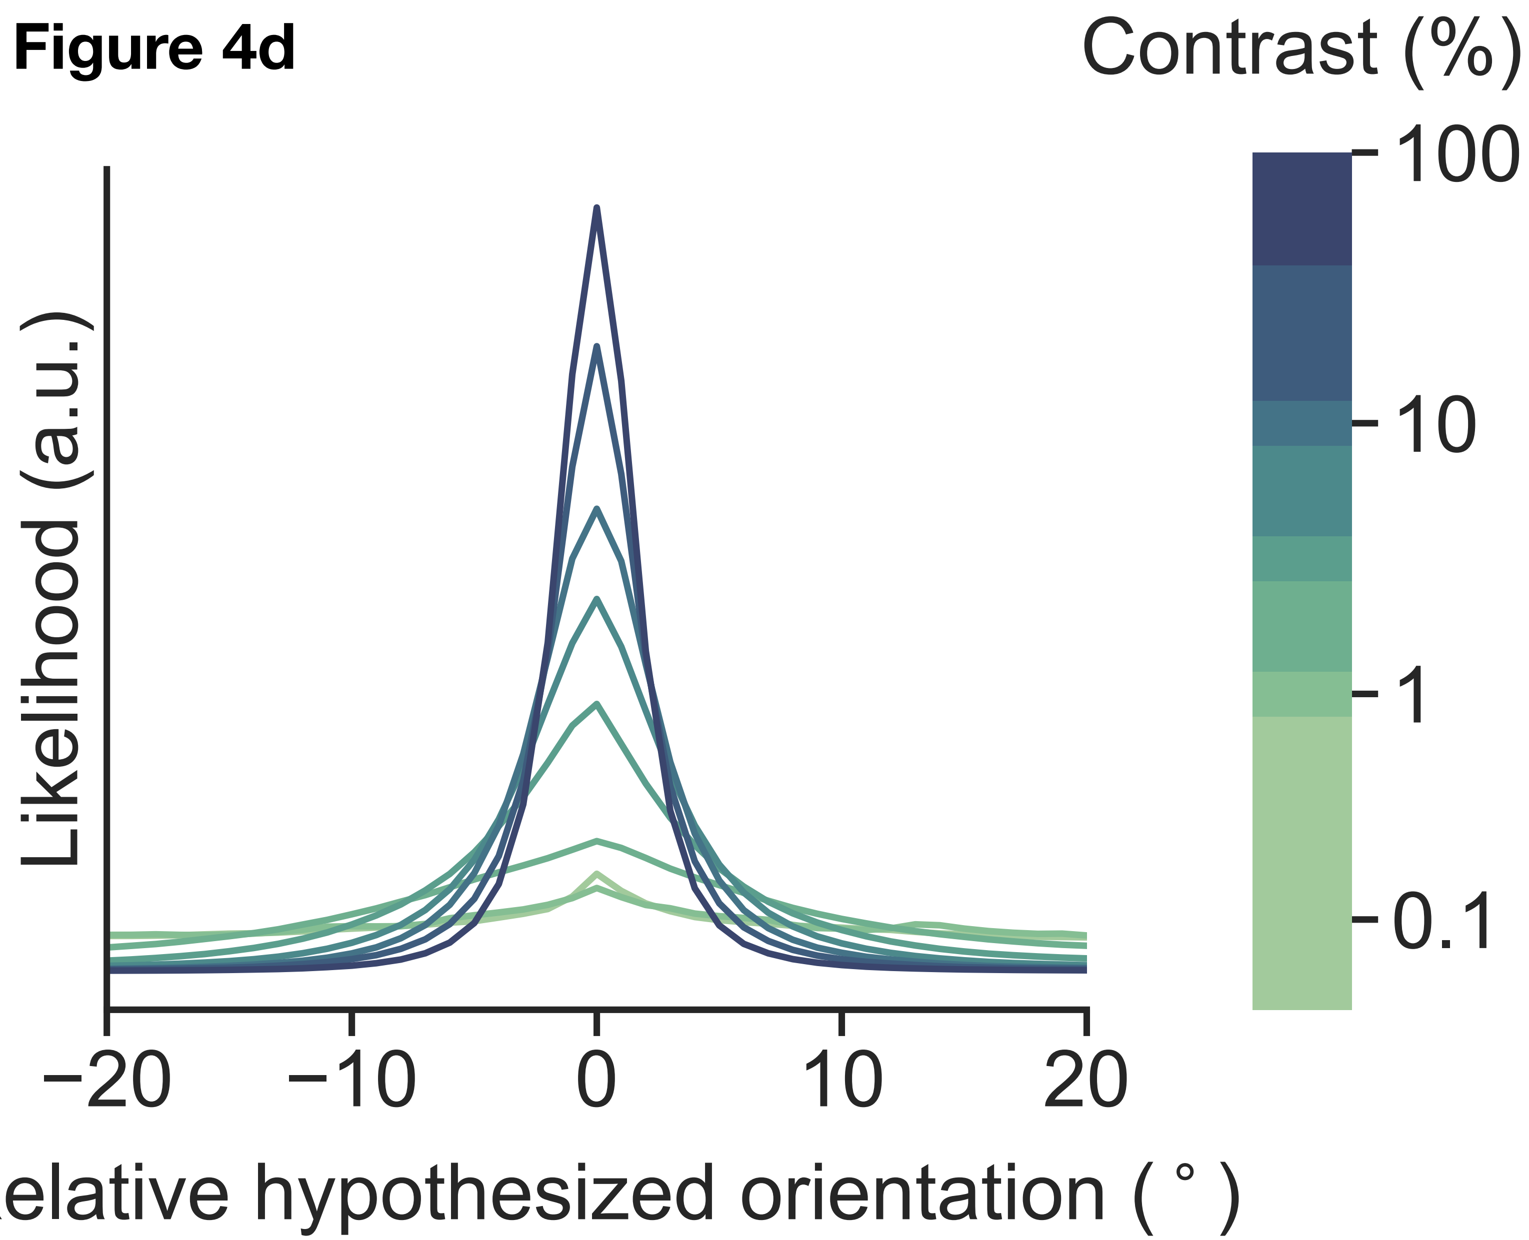


Walker, Cotton et al., 2020 *Nature Neuro*

# Which neurons contribute to different aspects of the likelihood functions?

**Figure 6**

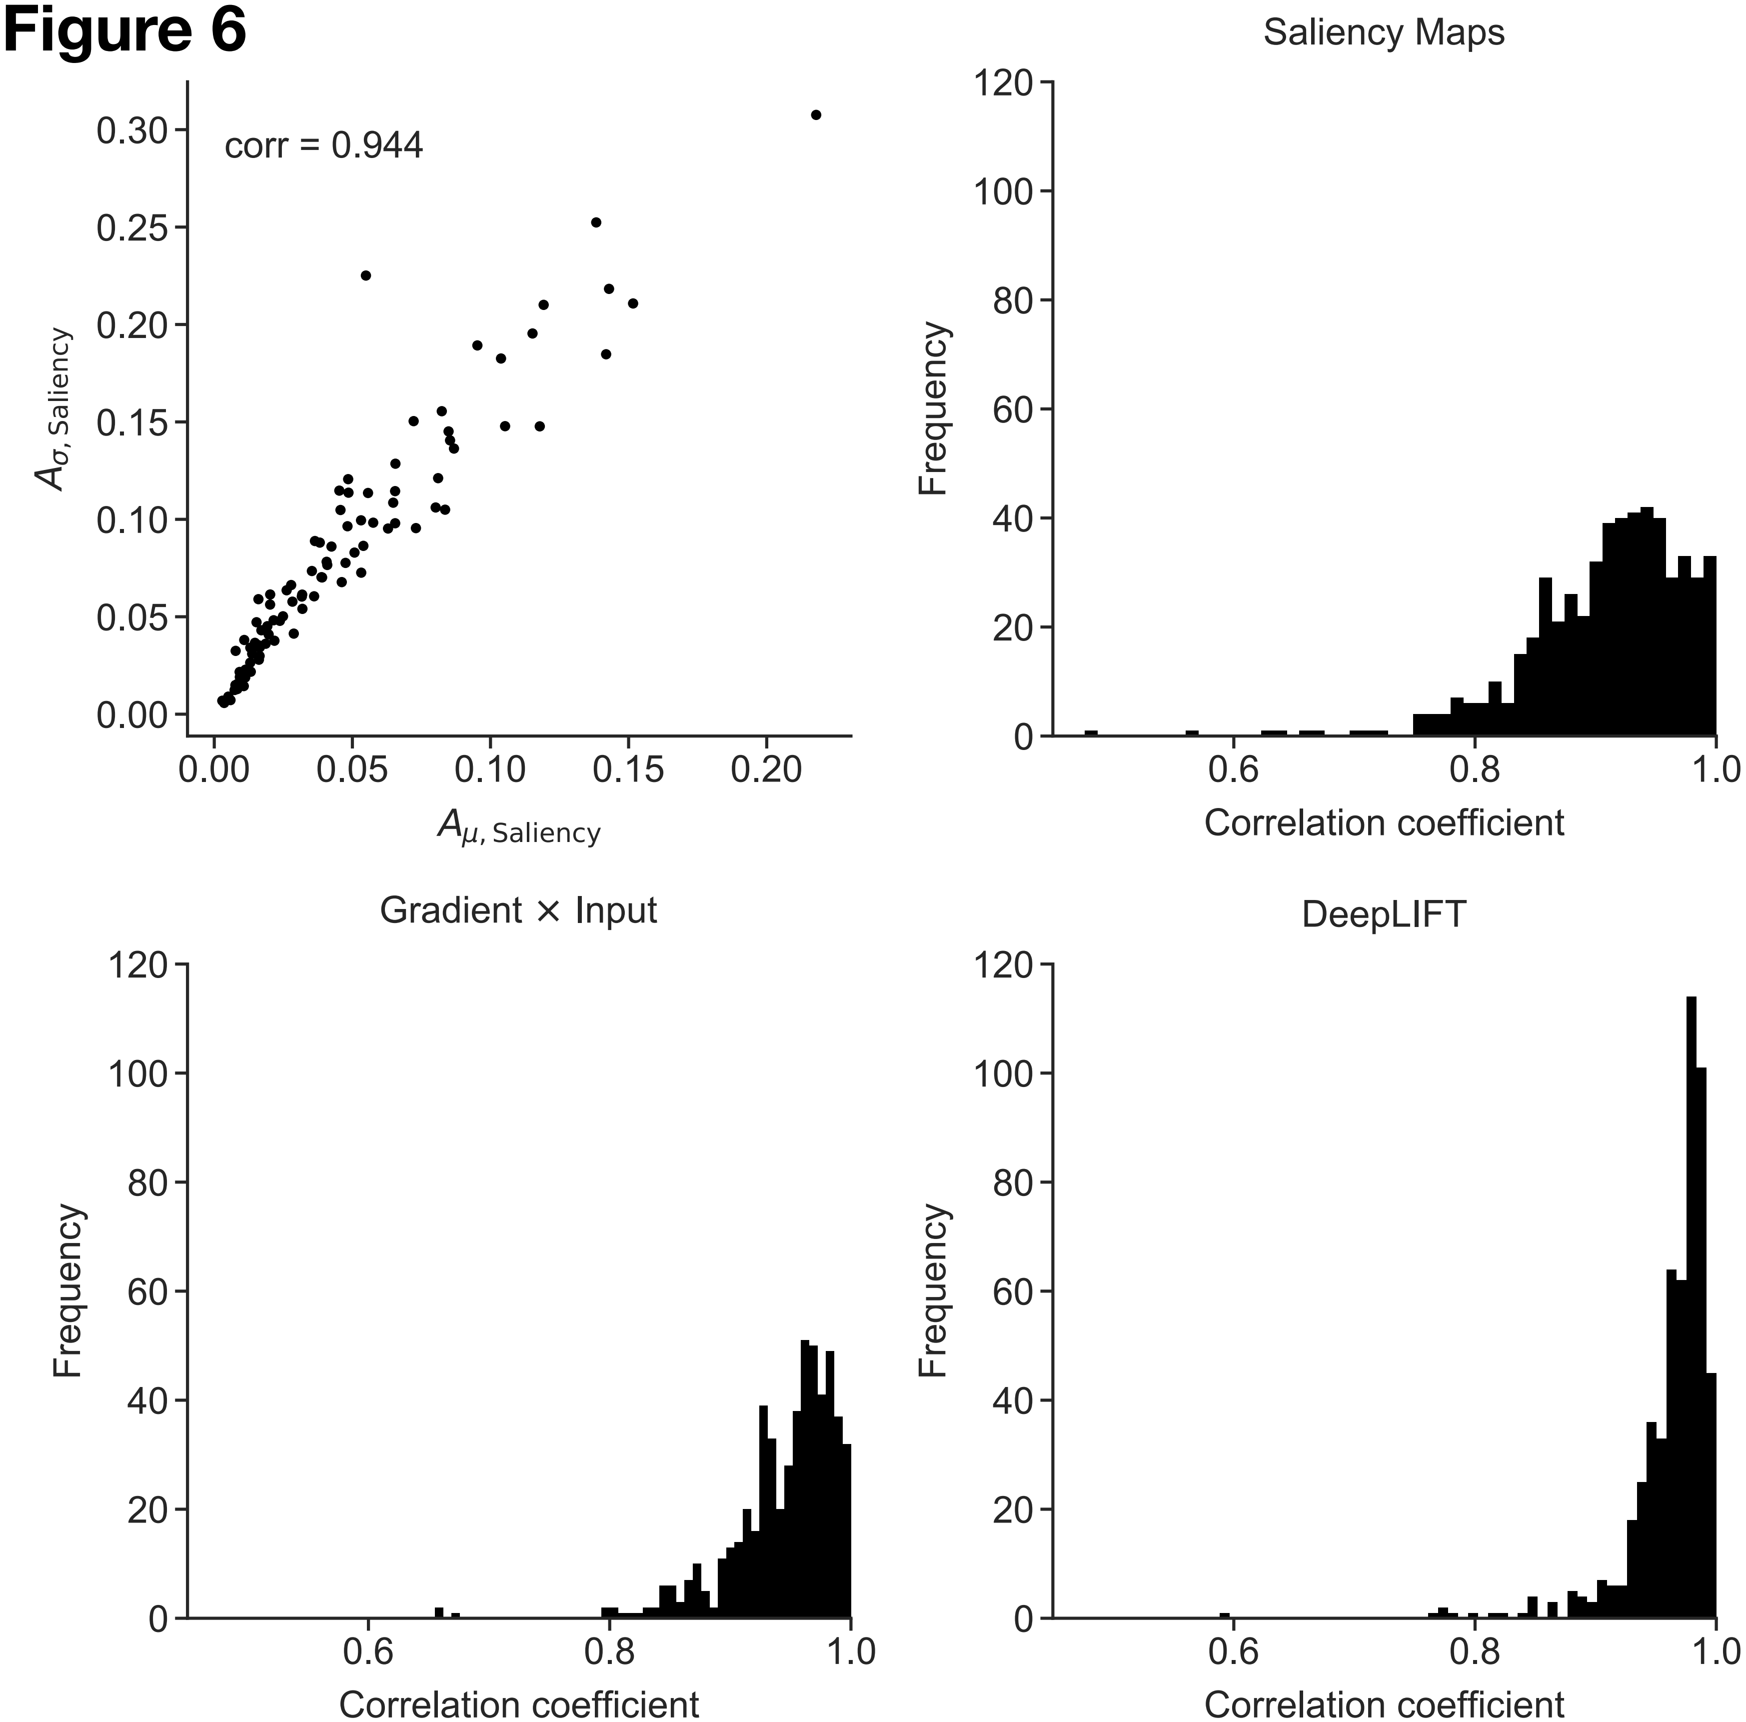

# Accounting of stimulus as a confound

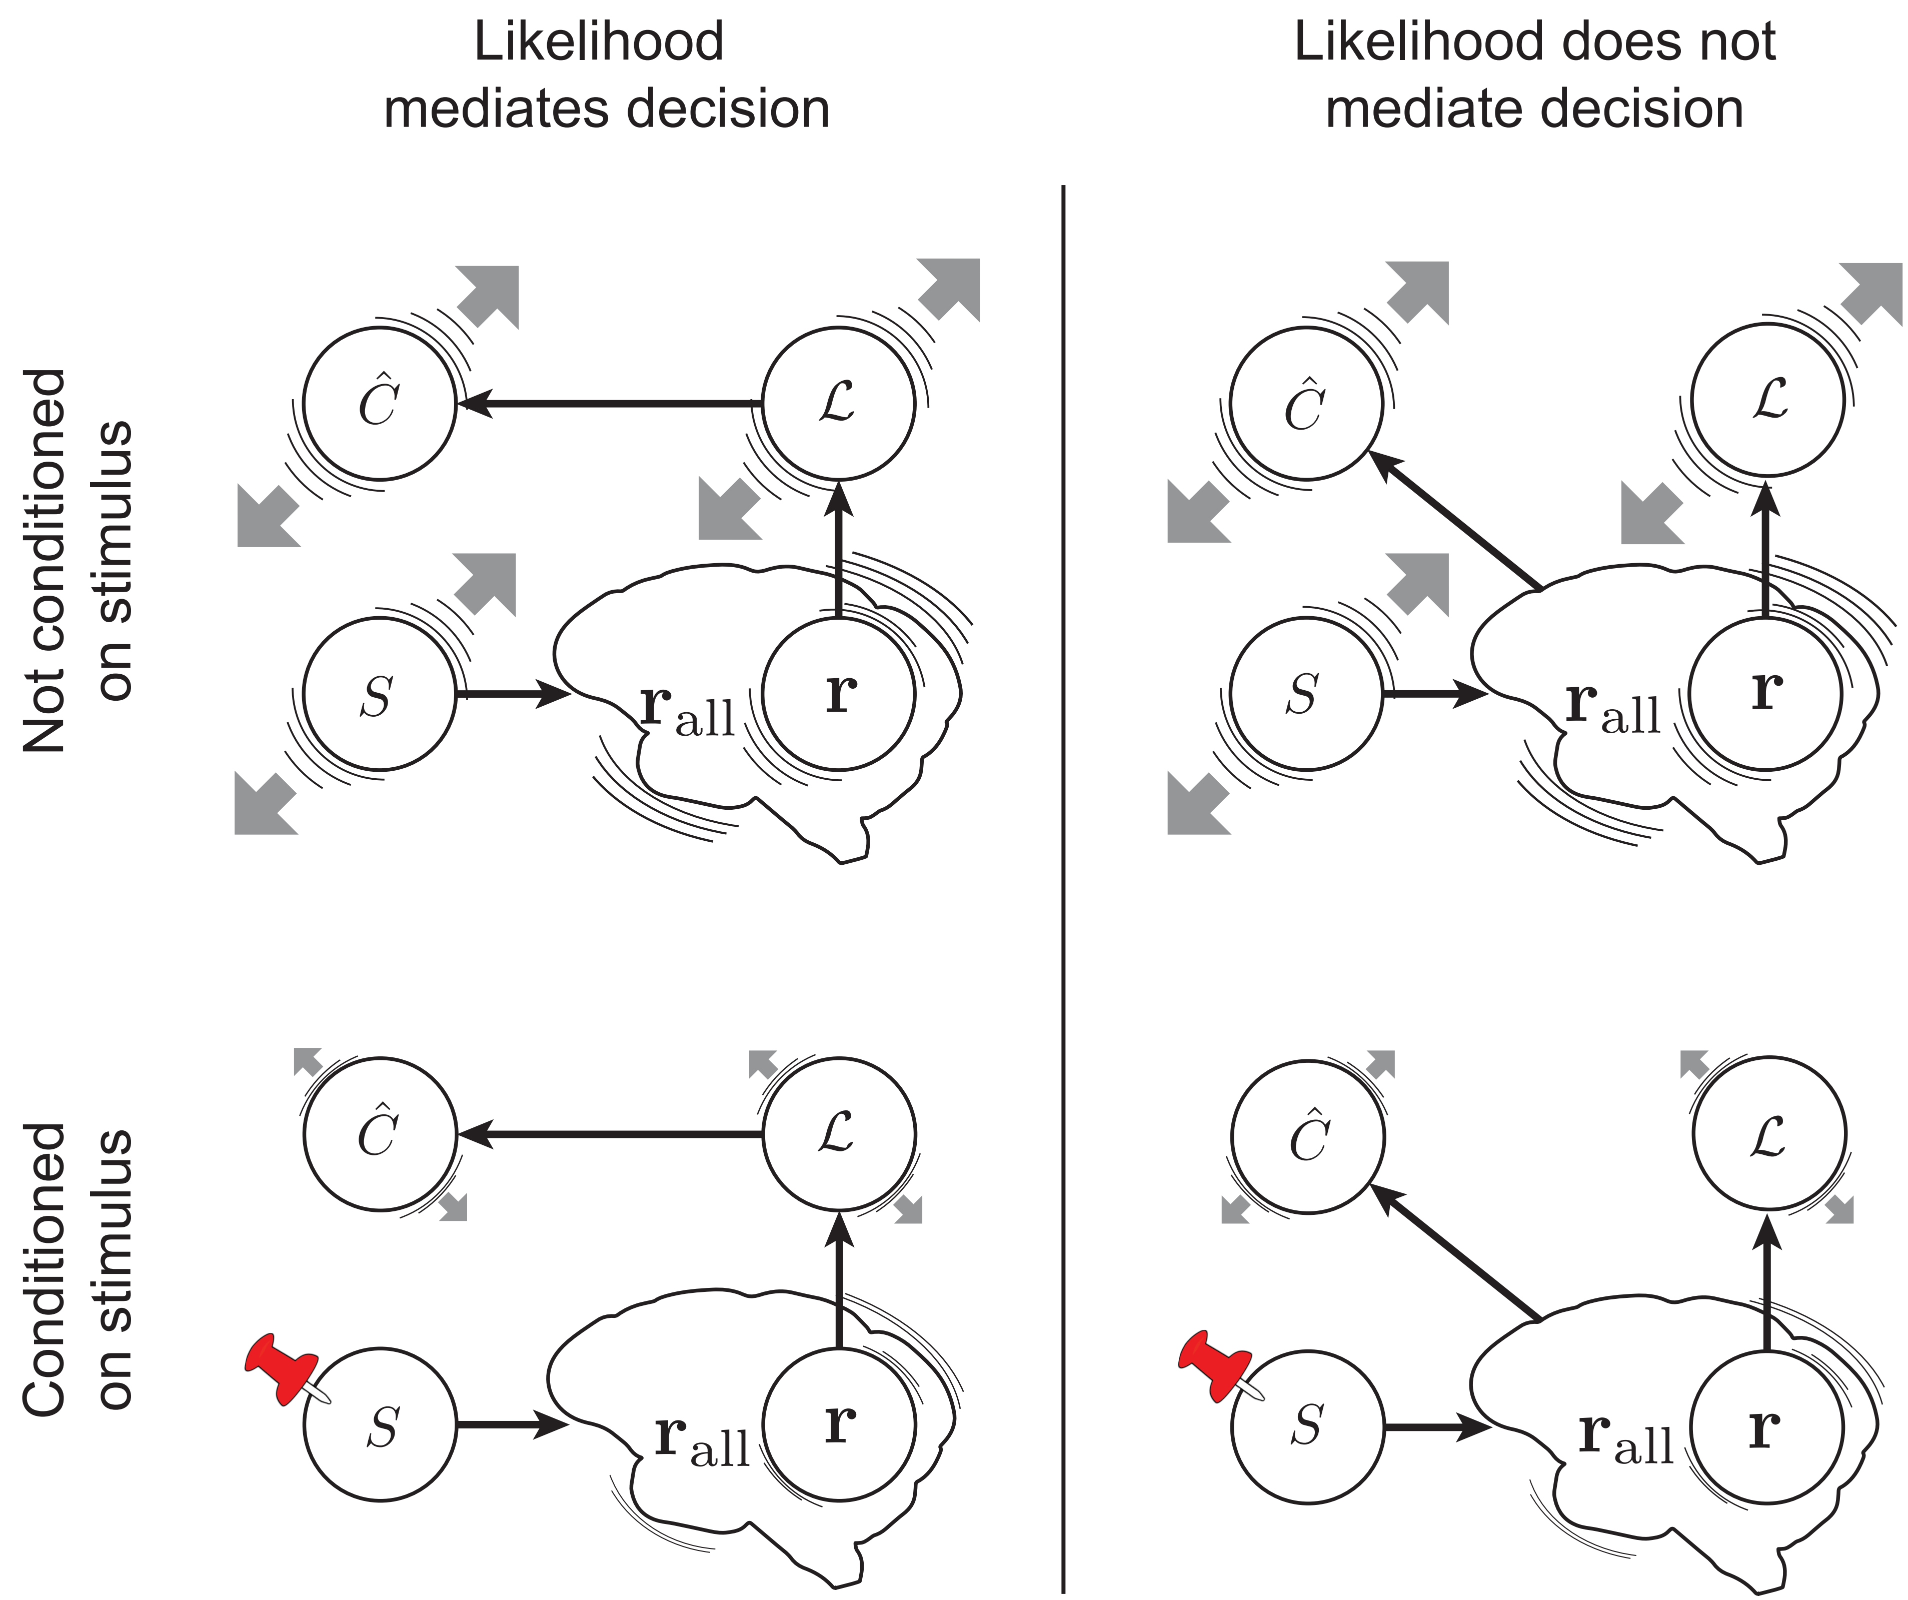


**Extended Figure 4**

# Time course of a trial

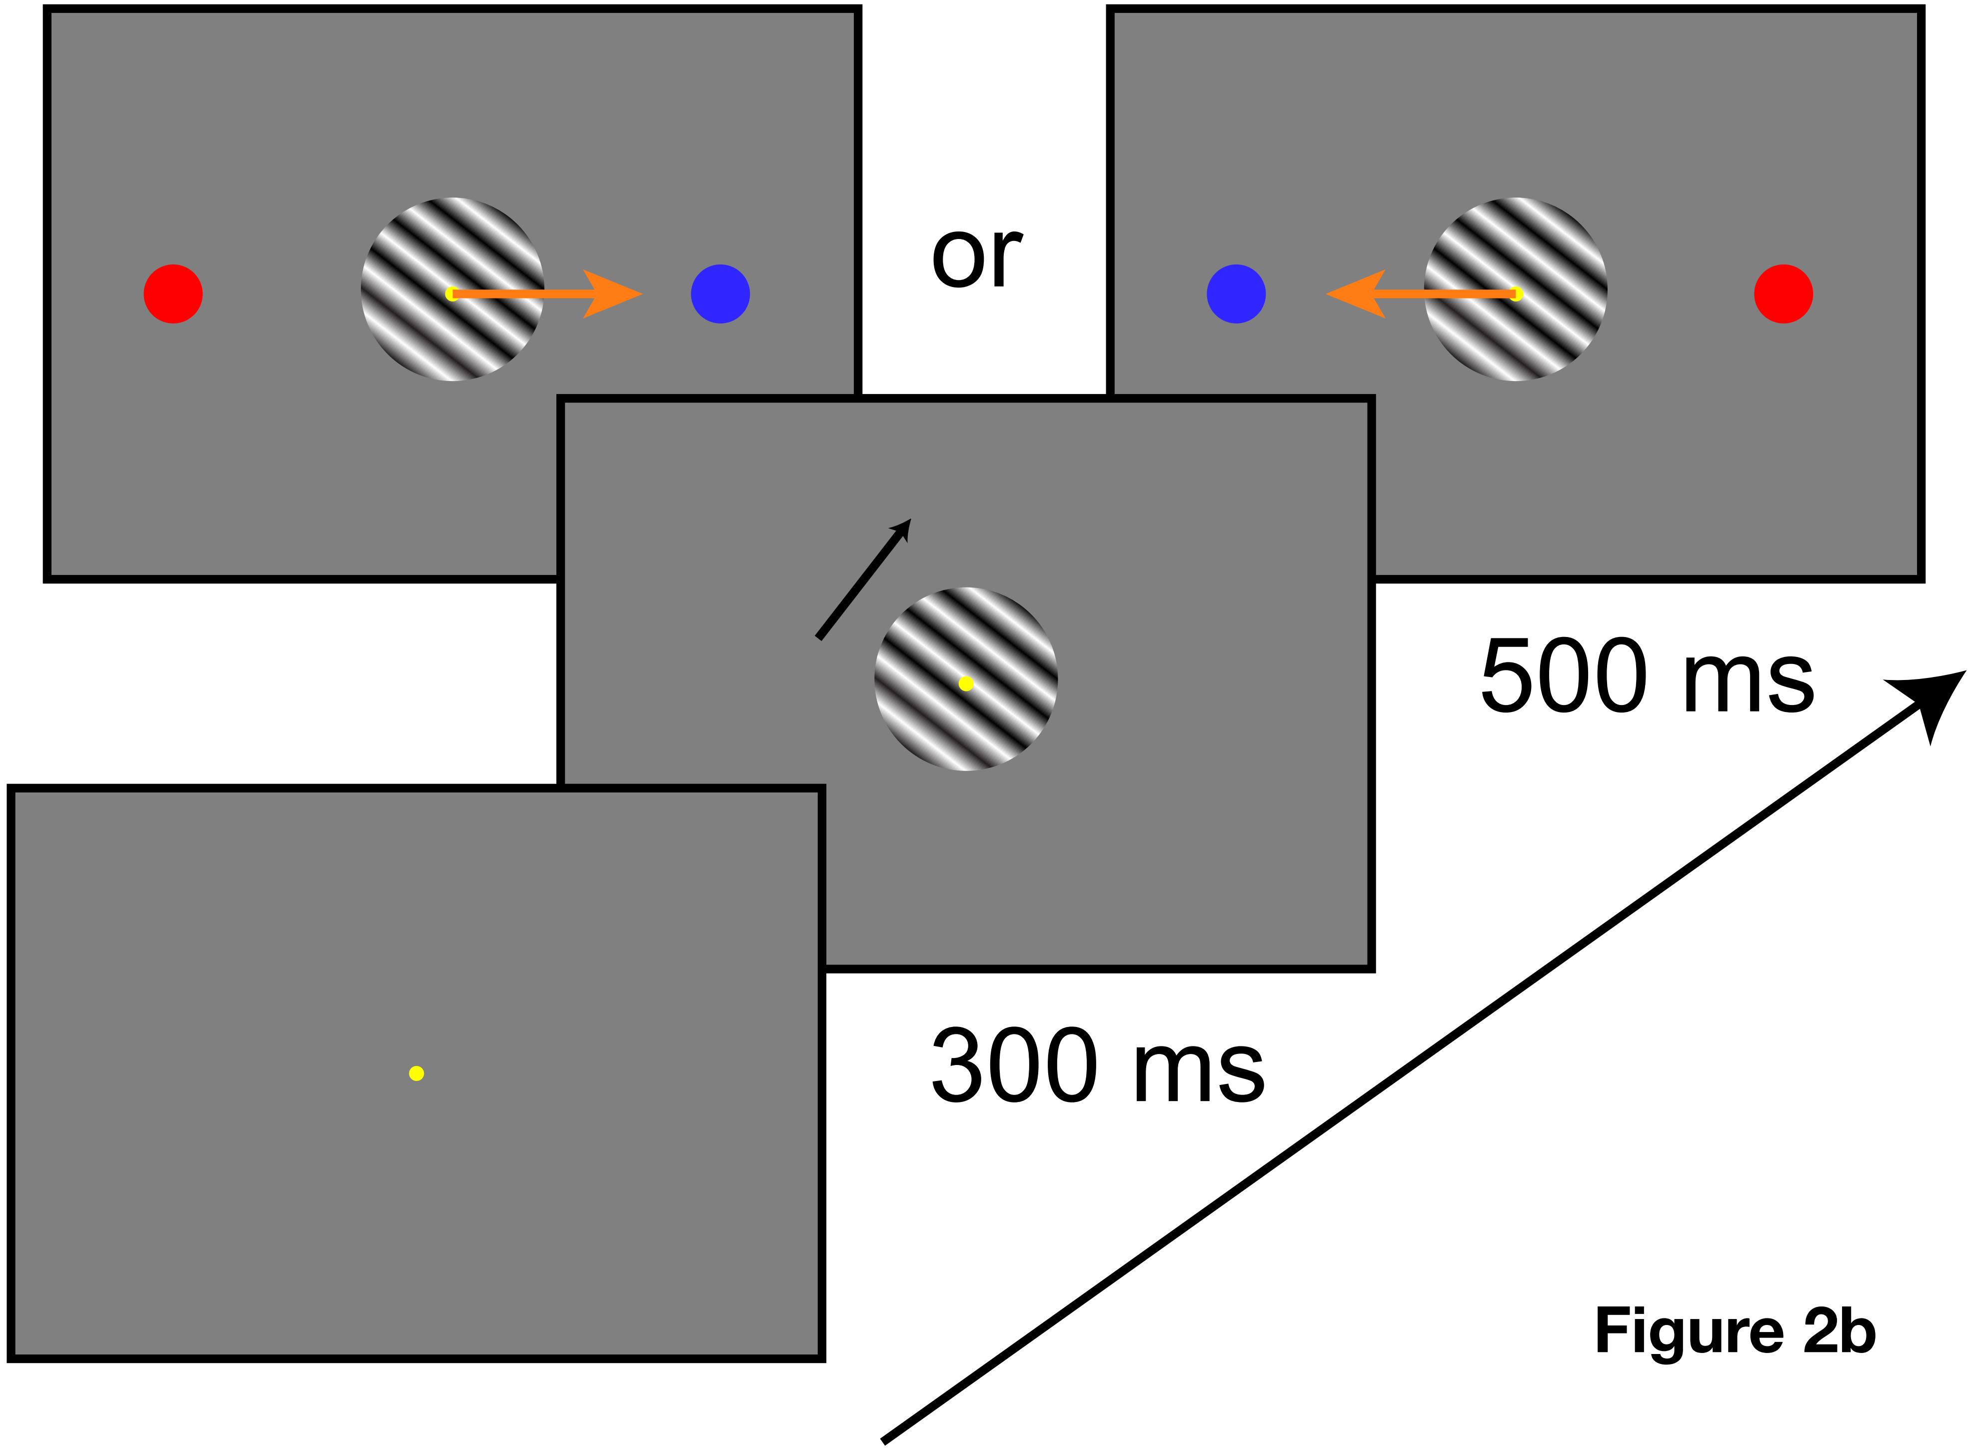


**Figure 2b**

Walker, Cotton et al., 2020 *Nature Neuro*

# Prediction of monkey's decisions

**Extended Figure 7a—d**

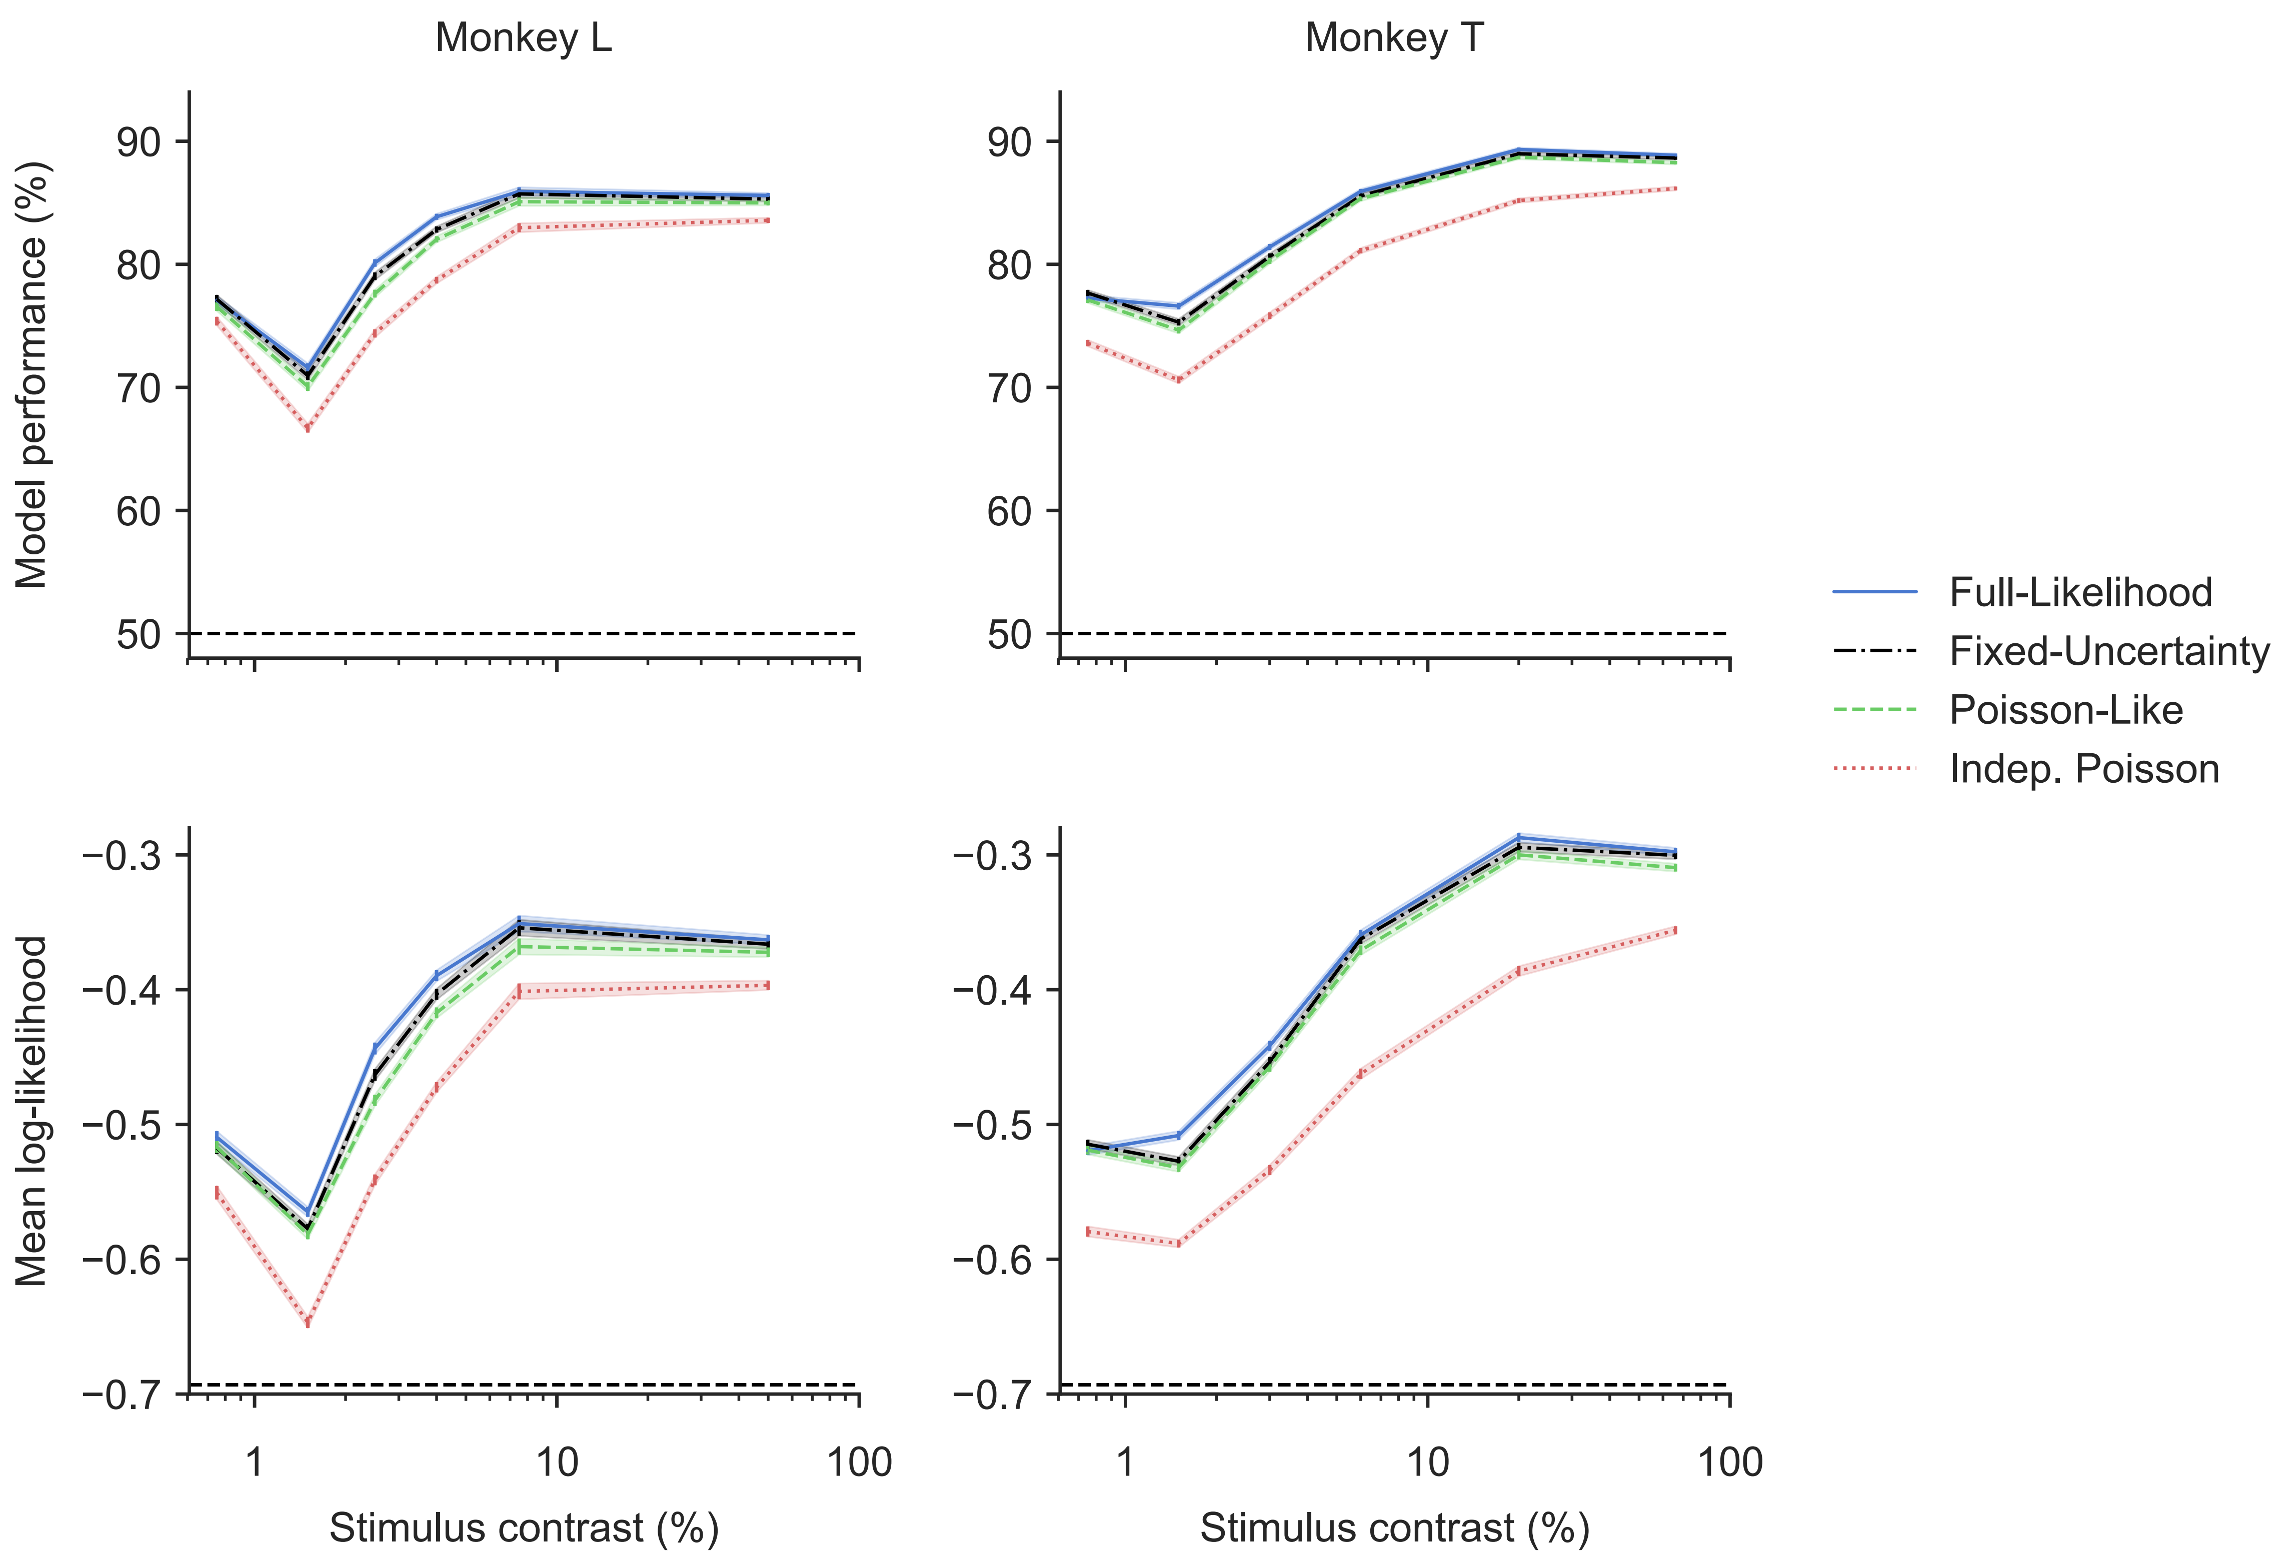

# Full-likelihood model performs better than likelihood functions based on parametric generative models

**Extended Figure 7e—f**

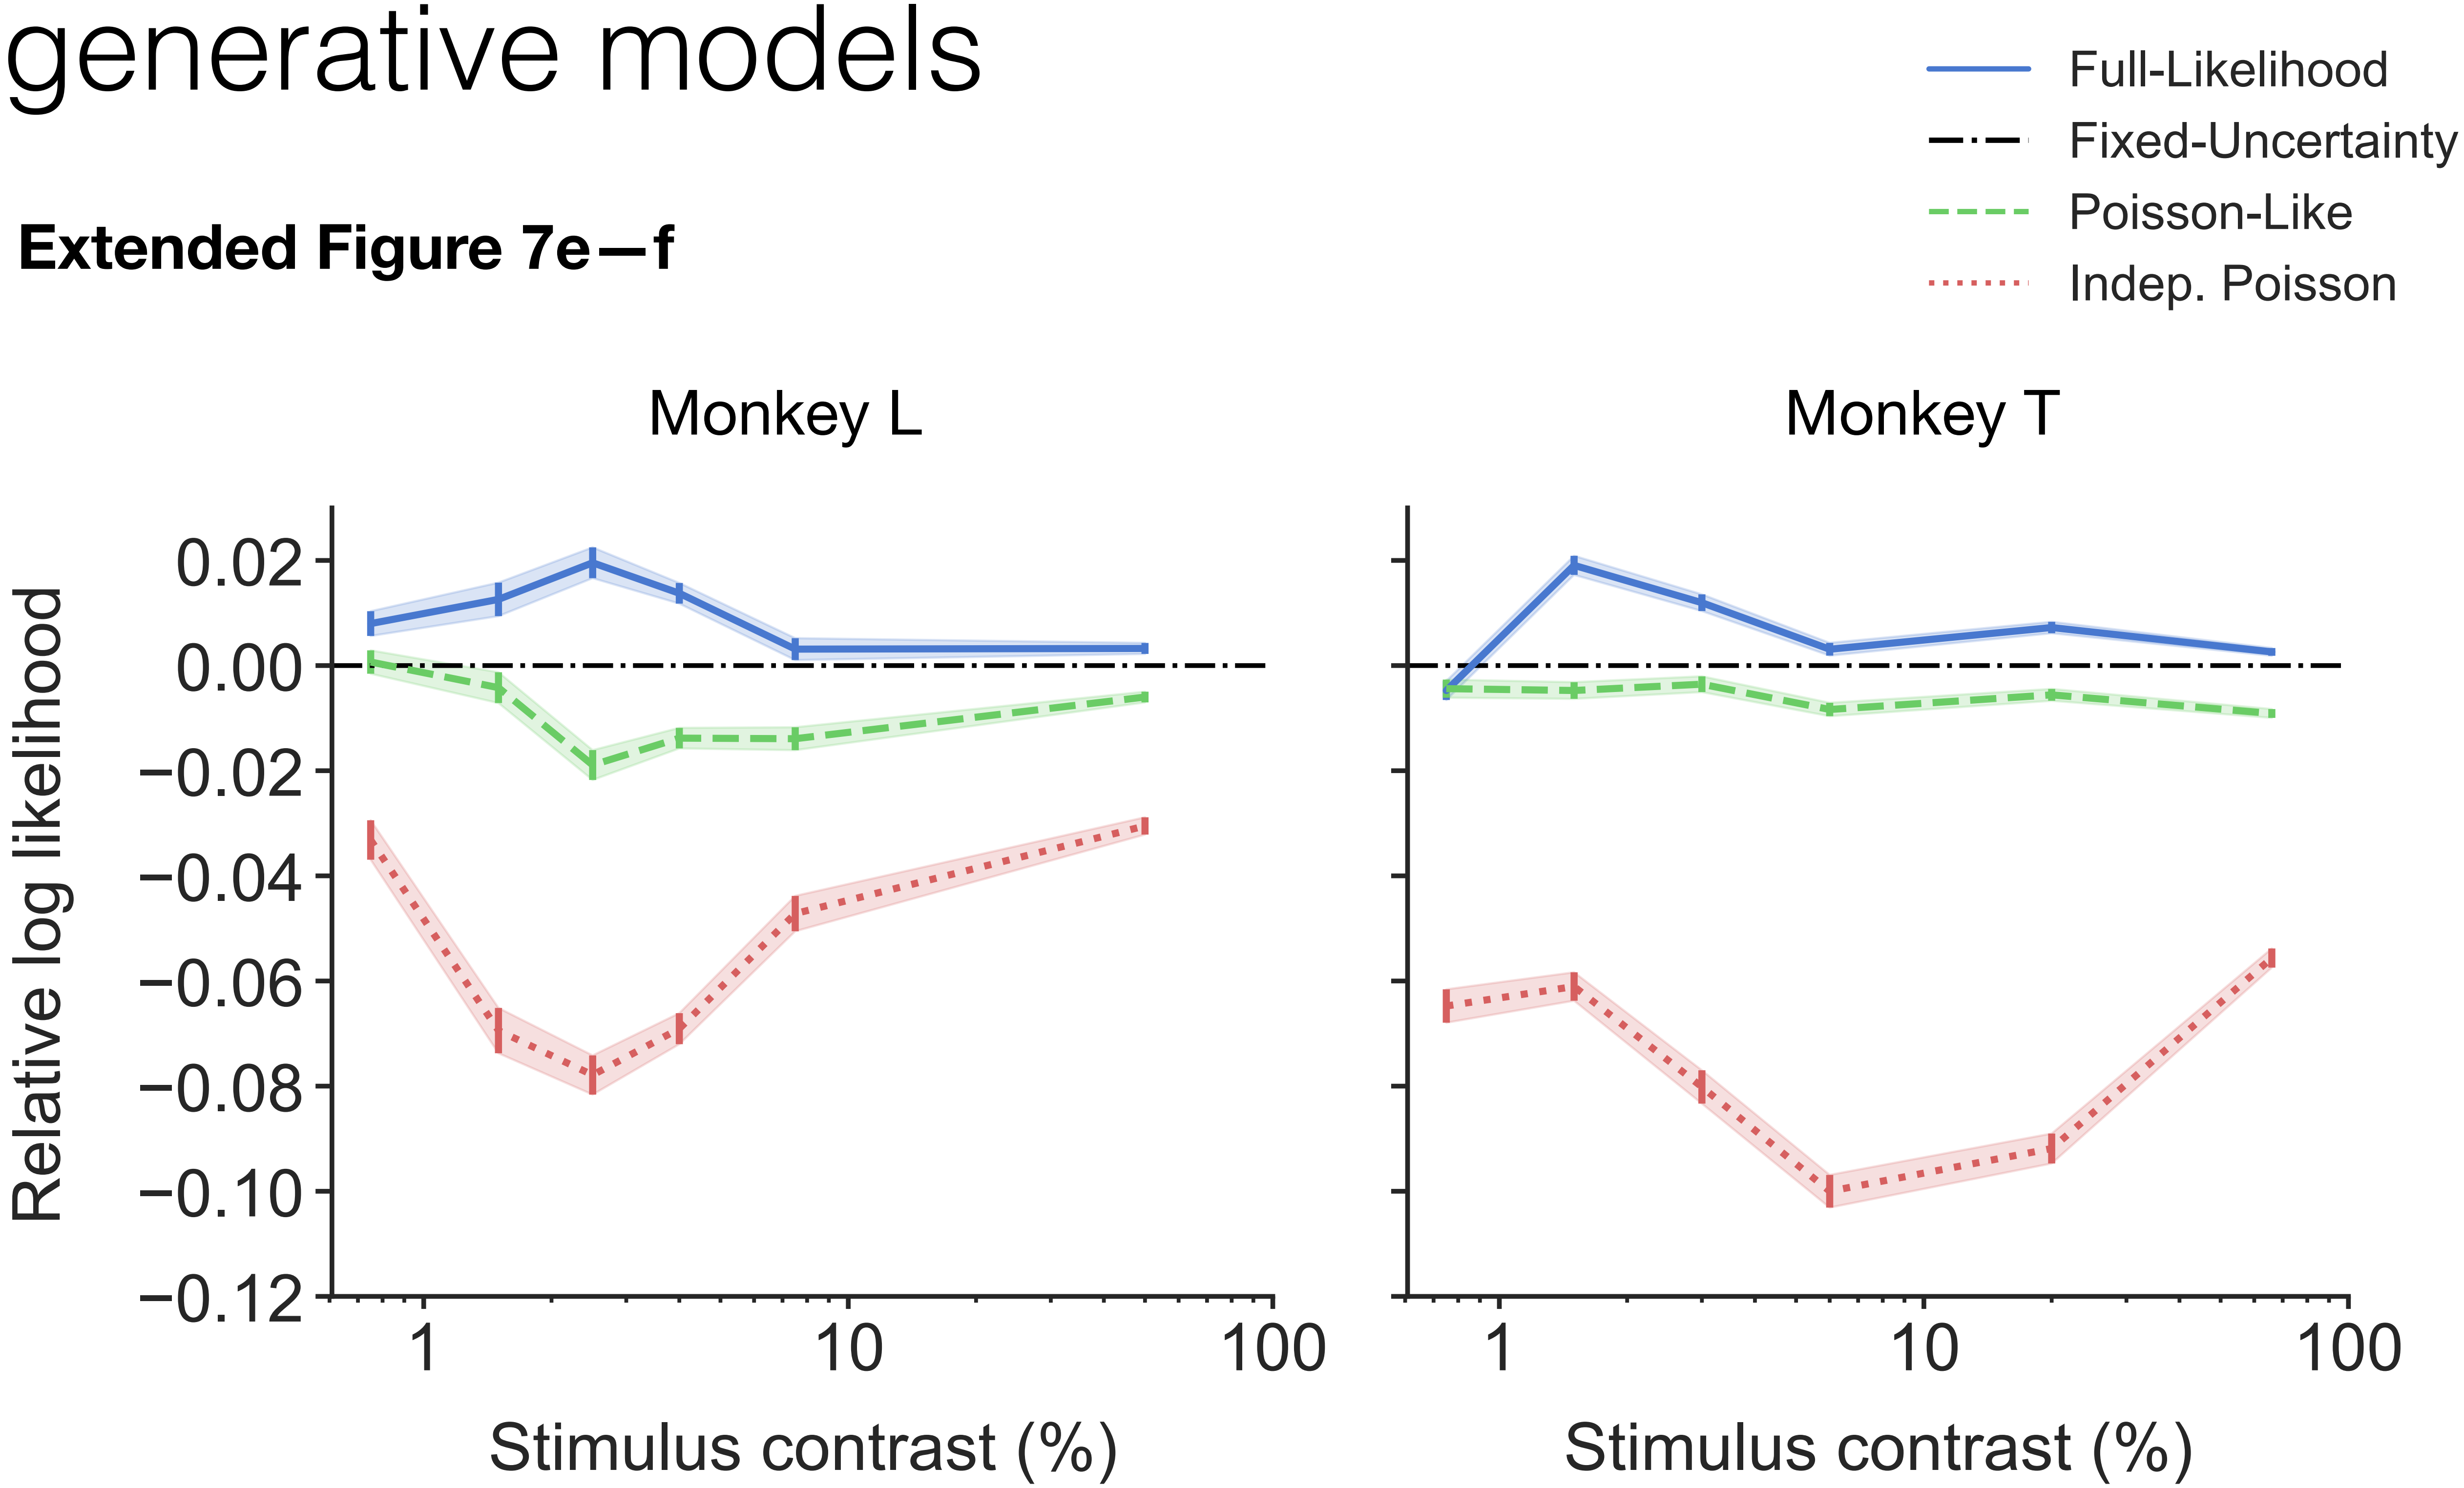

# Likelihood functions under different models

**Extended Figure 2**

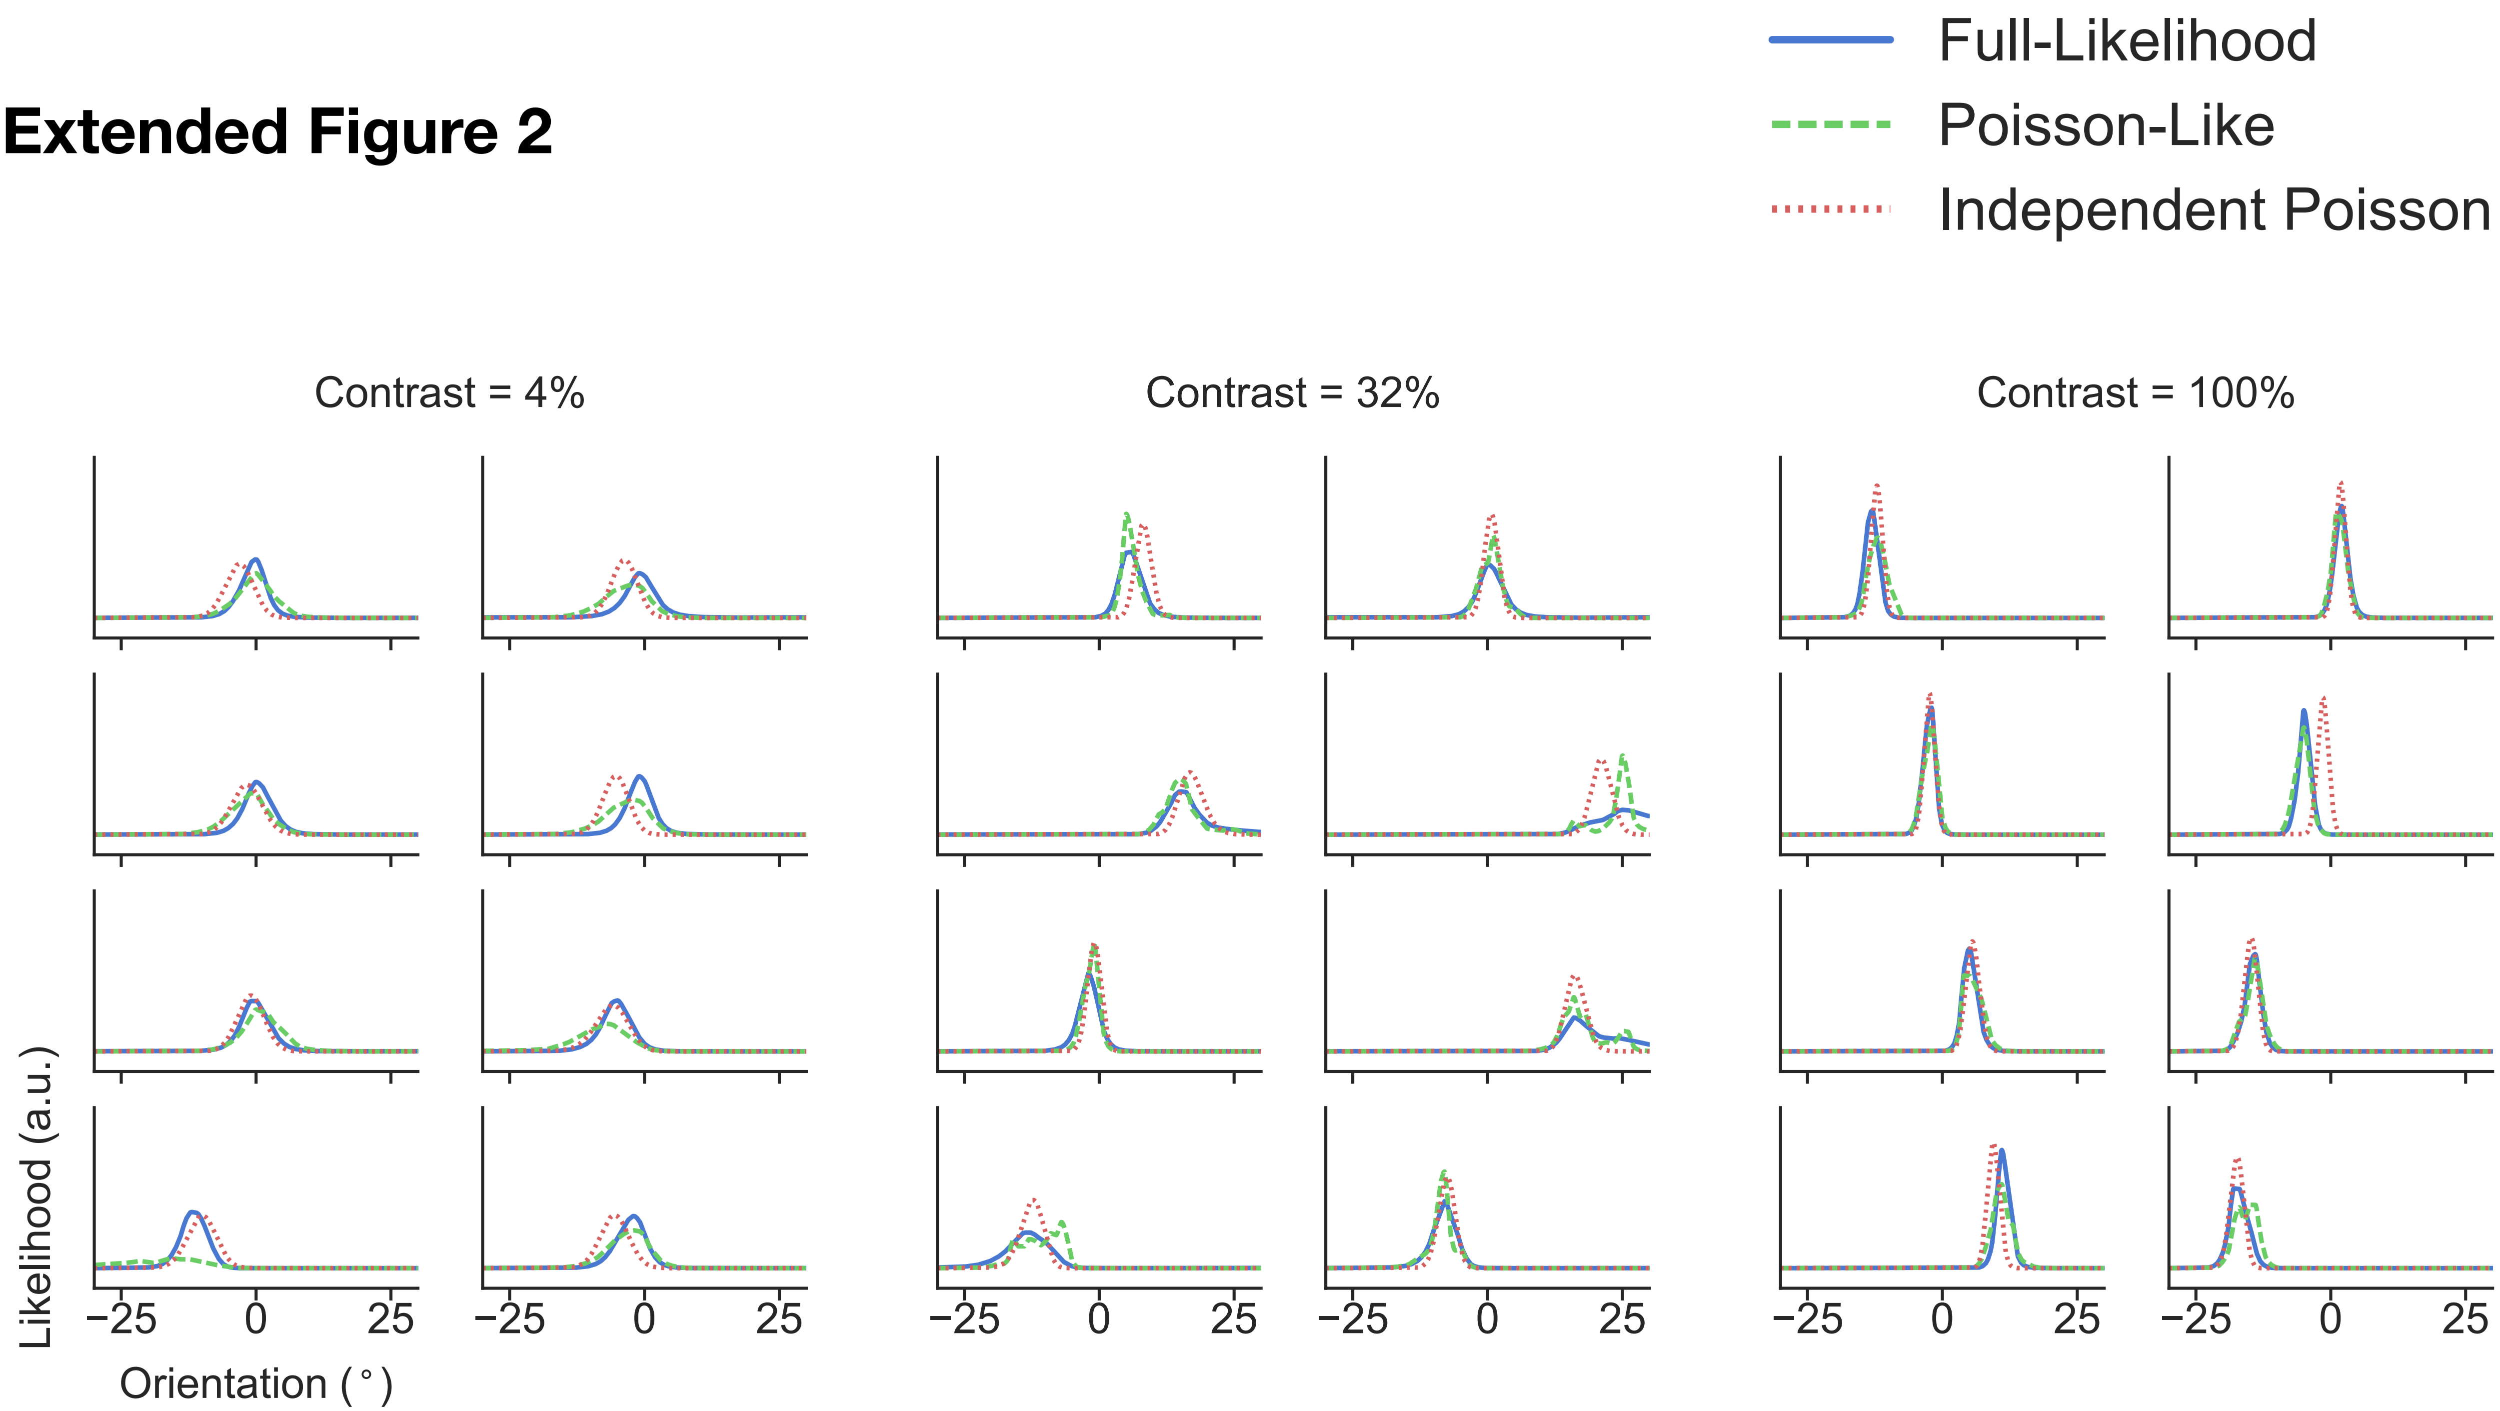

# Likelihood functions under different models

**Extended Figure 5b**

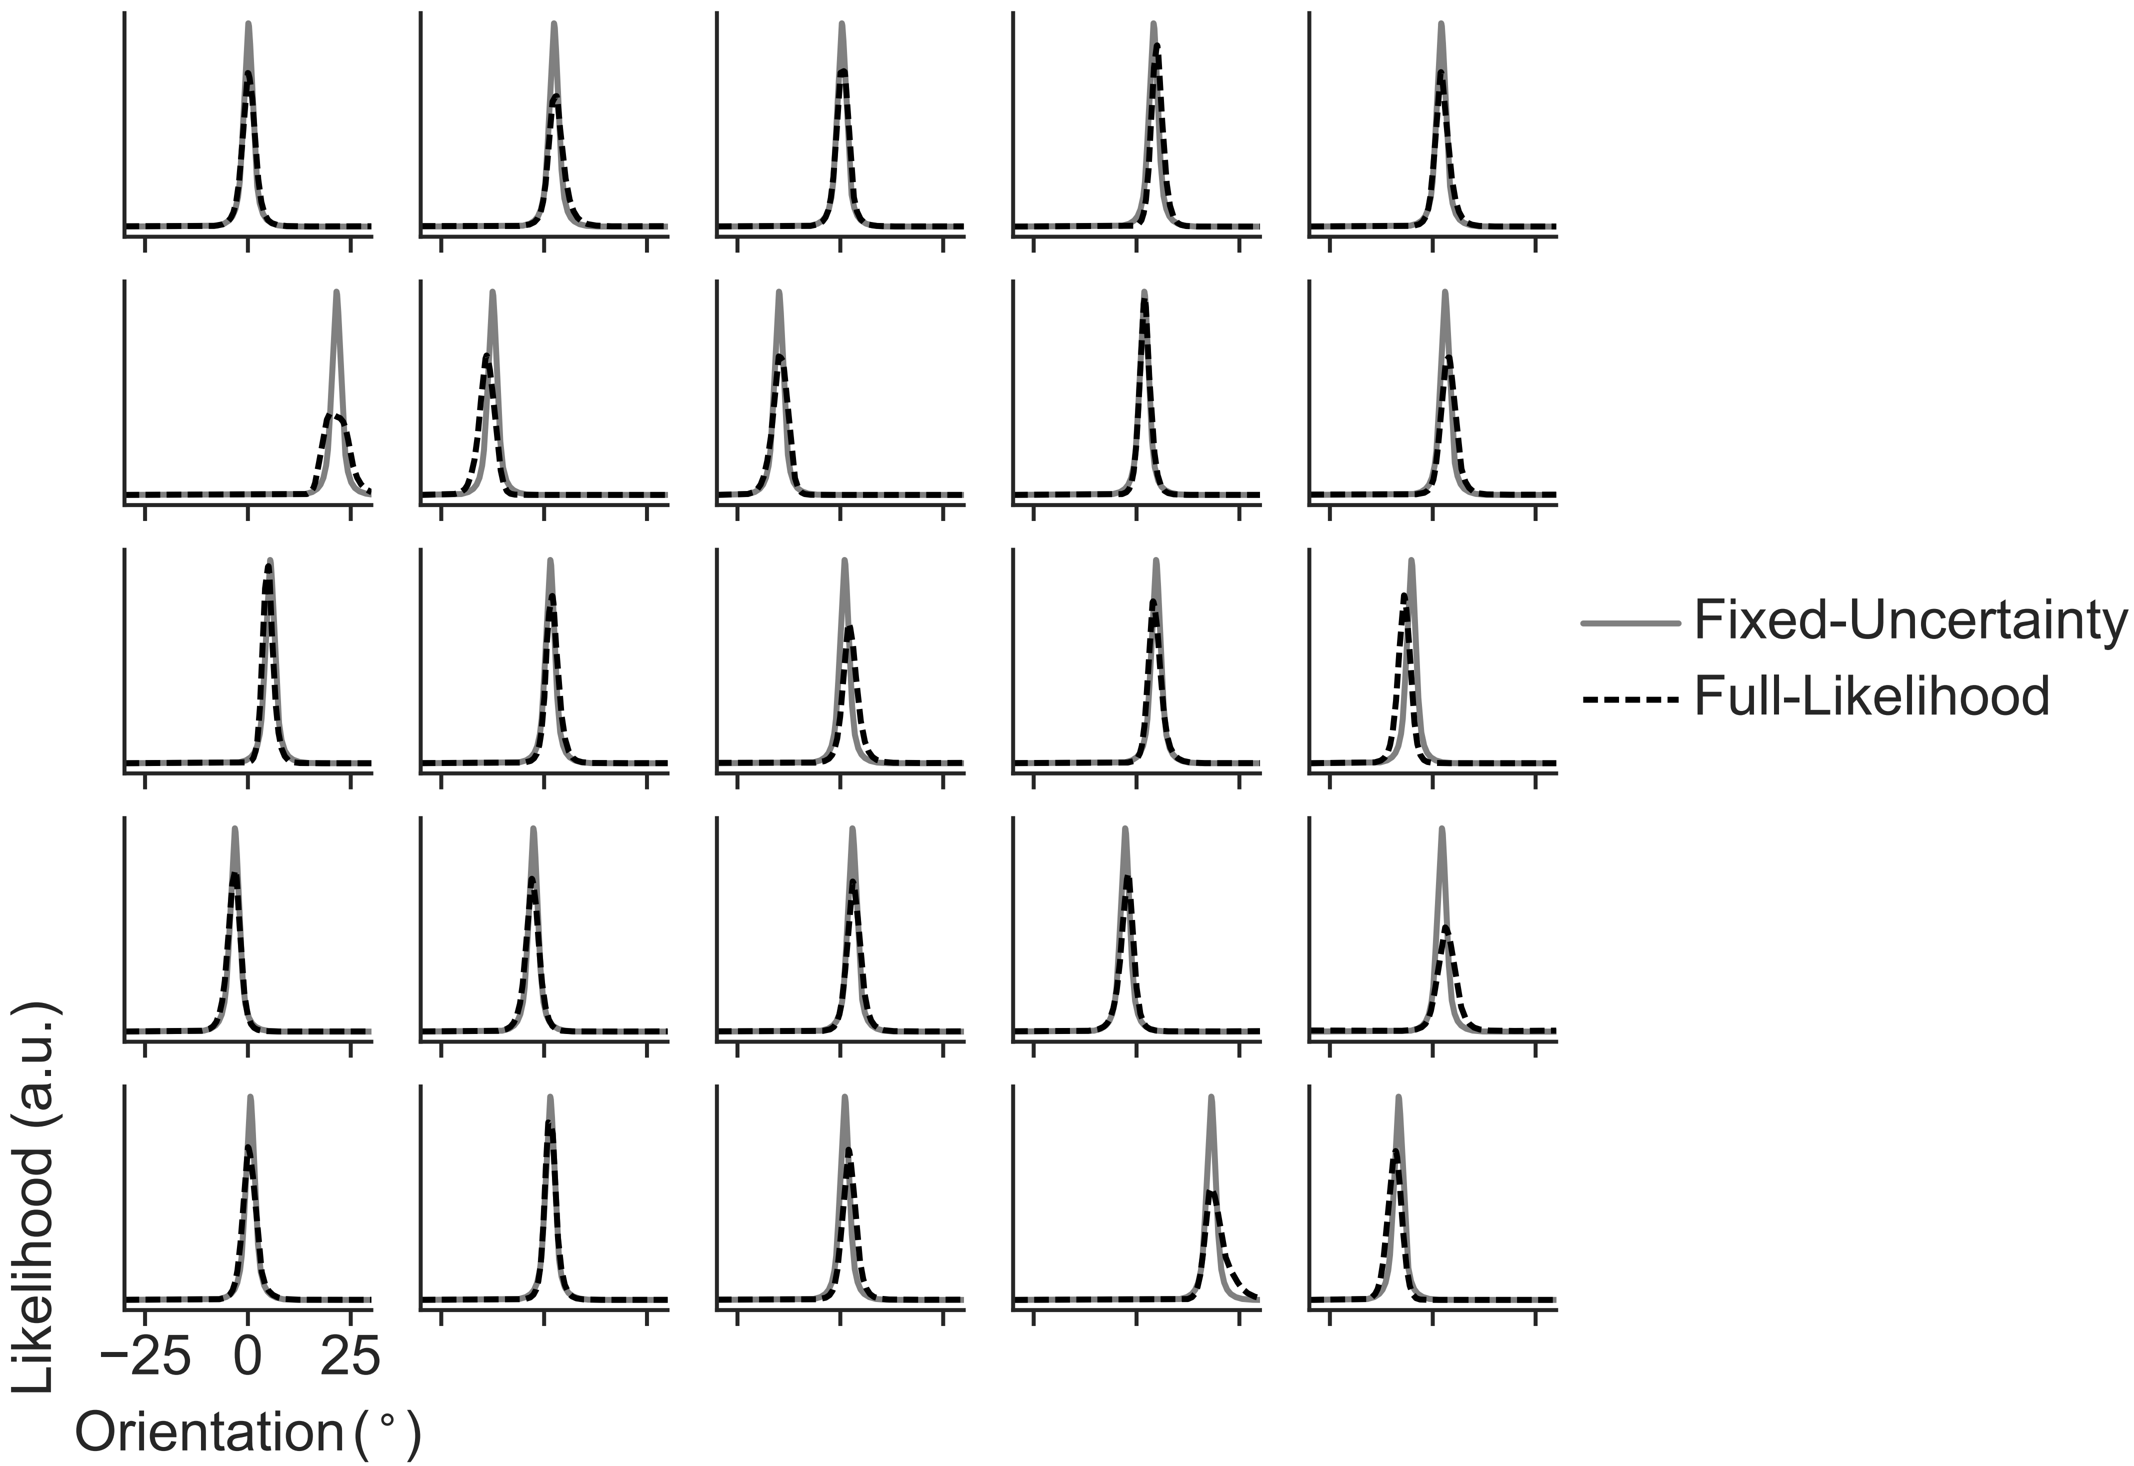

# Monkeys perform the task well

**Figure 2c**

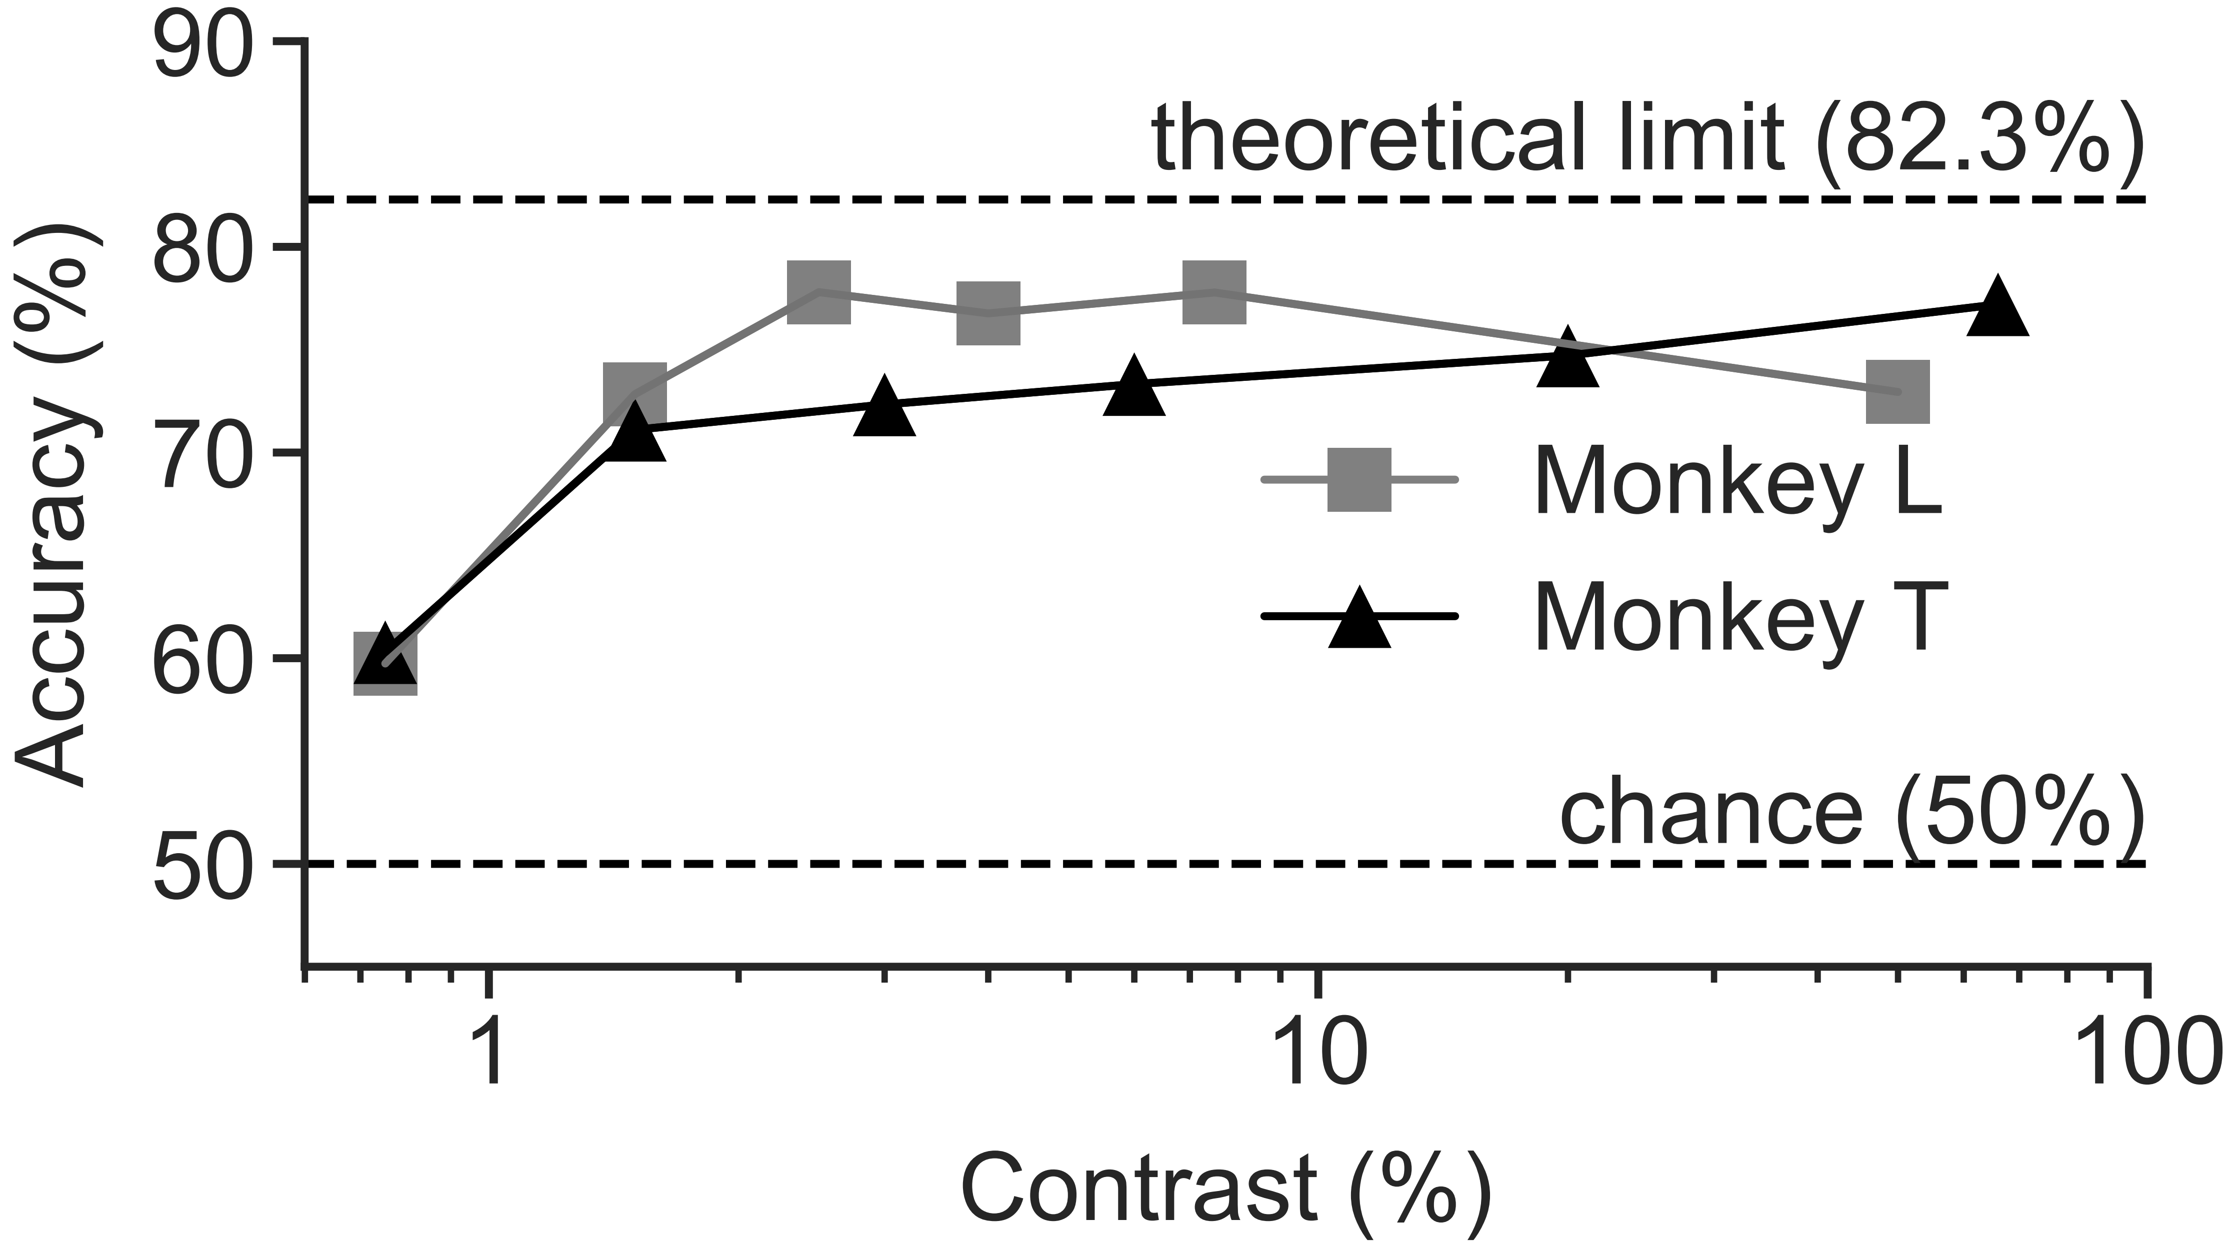

# Monkeys perform the task well

**Figure 2d**

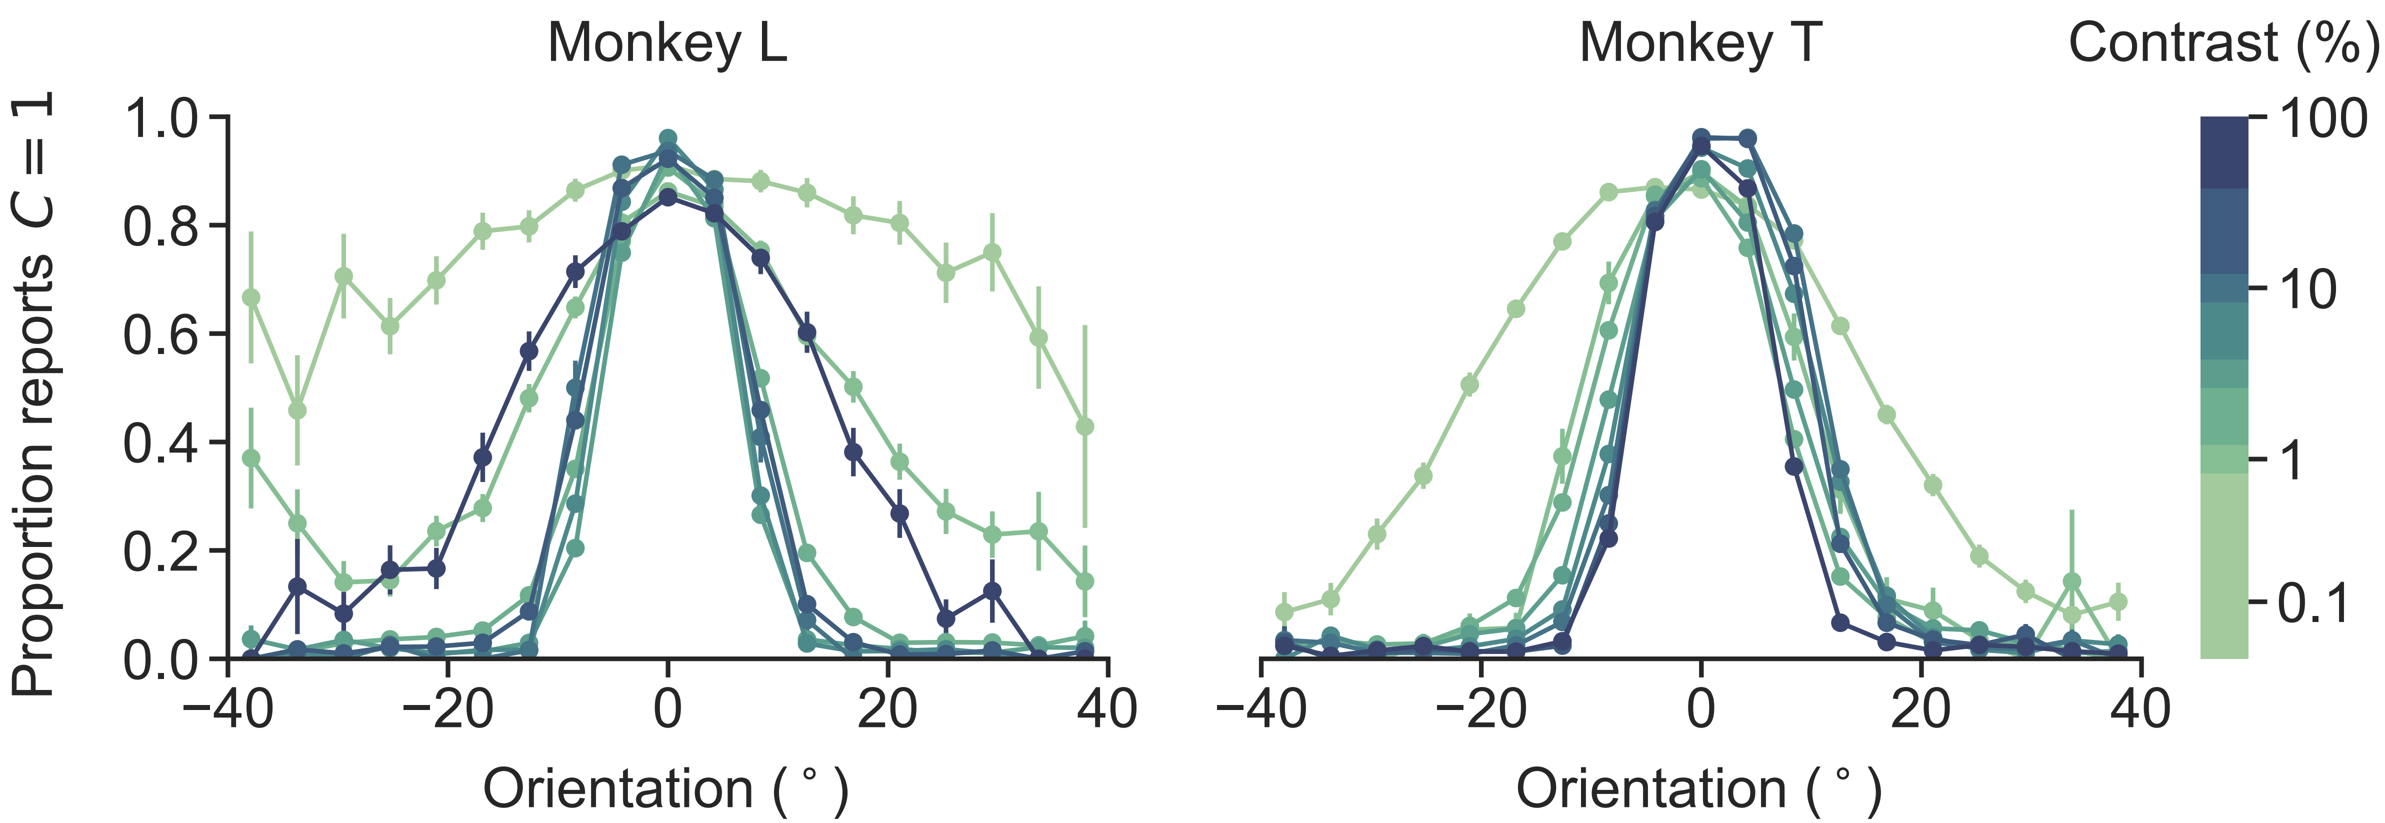

# Monkeys perform the task well

**Figure 2e**