Information geometry of the retinal representation manifold

Xuehao Ding, Dongsoo Lee, Joshua B. Melander, George Sivulka, Surya Ganguli, and Stephen A. Baccus

Stanford University

Some important stimuli to discriminate in nature

Discriminability vs Sensitivity

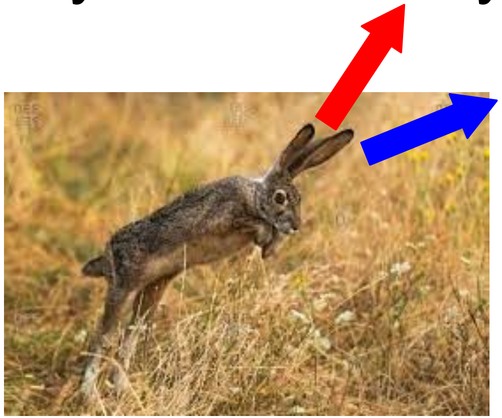

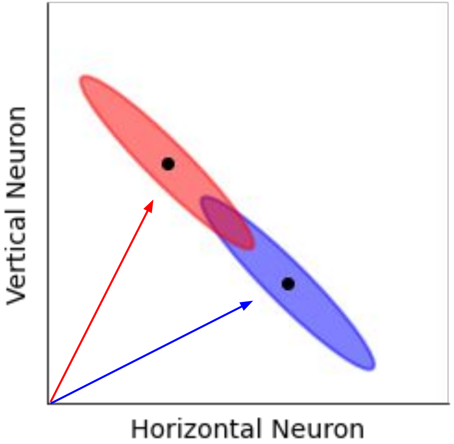

Only knowing sensitivity is not enough, we also need stochasticity to study the discriminability

Sensitivity, Stochasticity, and Discriminability (1d)

1d stimulus
1d response

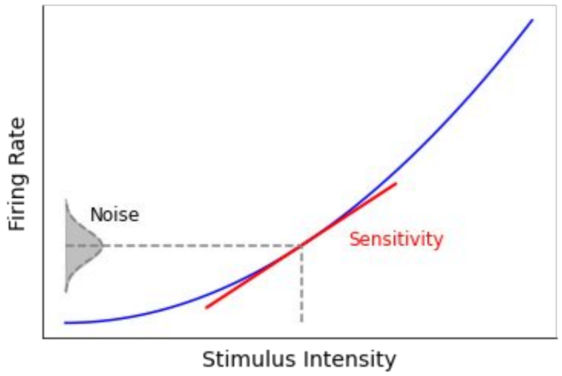

Notations:

- x: stimulus
- y: response
- P(y|x)
- μ(x):=E(y|x) firing rates

Fisher information

$$I(x) = \int P(y|x) \left(\frac{\partial \ln P(y|x)}{\partial x} \right)^2 dy$$

Discriminability

$$= \left(\frac{d\mu(x)}{dx} \right)^2 / \sigma^2$$

(when P(y|x) Gaussian and noise locally independent of x)

stochasticity

sensitivity

Cramér–Rao bound: $\mathrm{var}(\hat{\theta}) \geq \frac{1}{I(\theta)}$

Stochastic Model

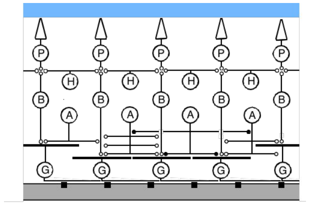

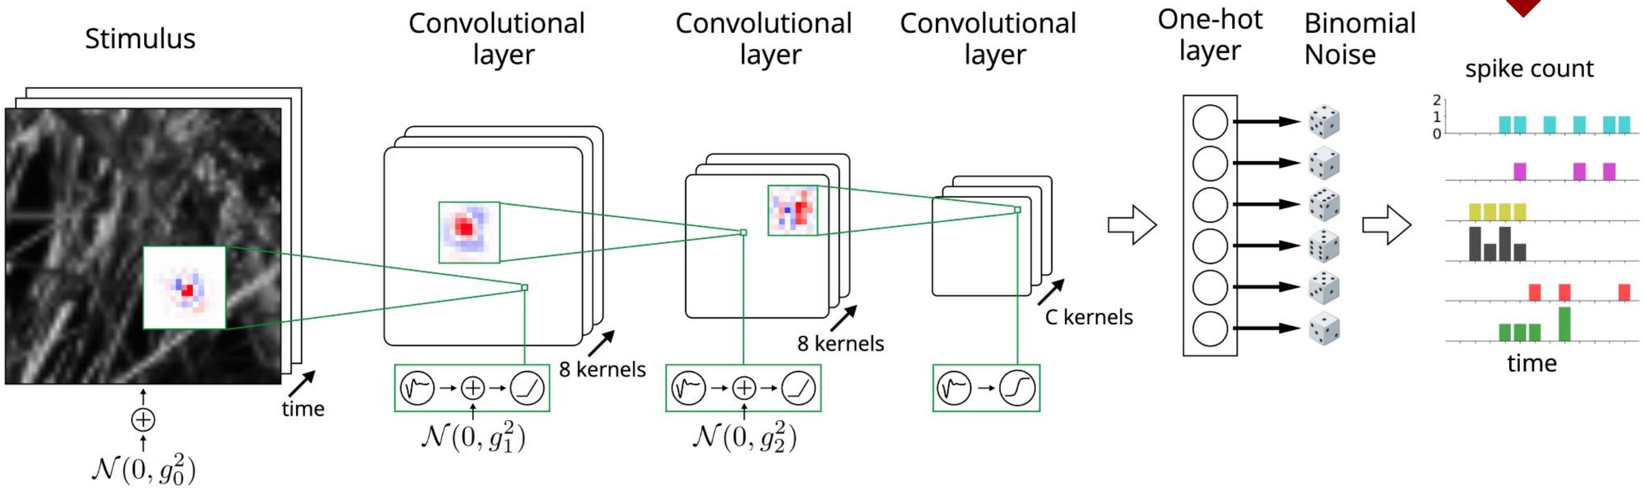

Binomial Noise:

$$p(n; r, k, N) \propto \frac{\Gamma(k \cdot N + 1)}{\Gamma(k \cdot (N - n) + 1) \cdot \Gamma(k \cdot n + 1)} r^{k \cdot n} \cdot (1 - r)^{k \cdot (N - n)}, \quad n = 0, 1, \ldots, N$$

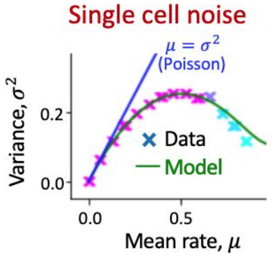

The model captures second-order statistics

Correlations:

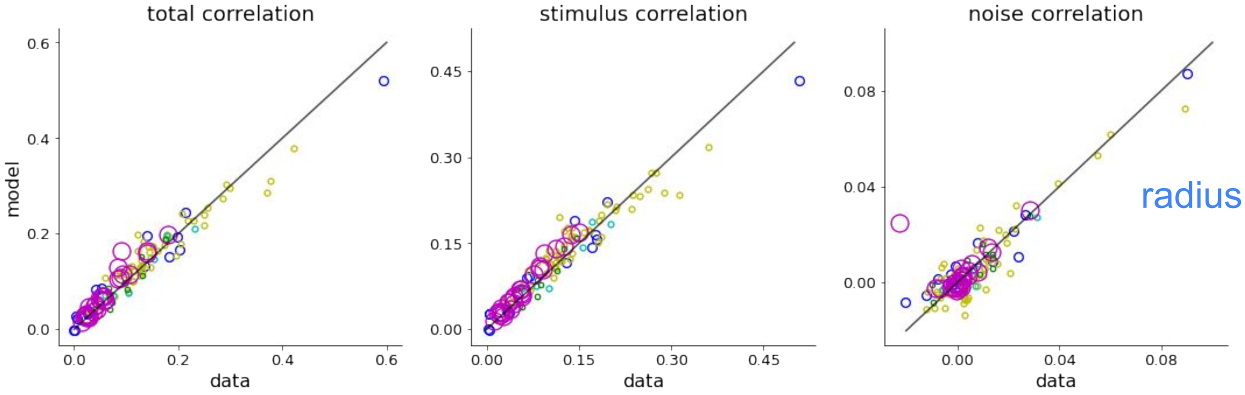

radius ~ sqrt(# trials * length)

Single-cell variability:

High dimensional case

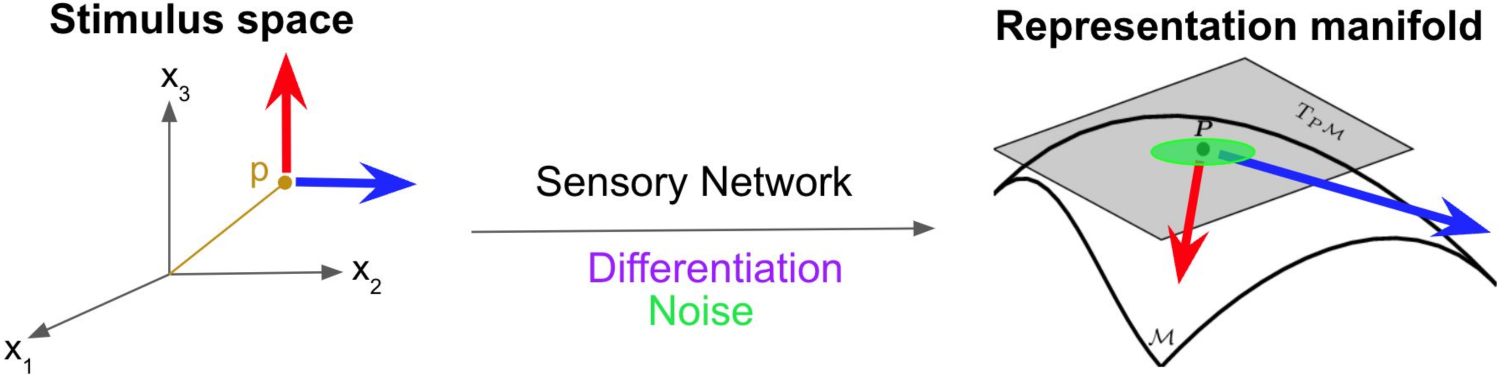

Dim = # pixels
(spatio-temporal 40*50*50)

Full Dim = # neurons (a few thousands)
Manifold Dim = a few hundreds

Fisher information matrix

$$I_{ij}(x) = \int P(y|x) \frac{\partial \ln P(y|x)}{\partial x_i} \frac{\partial \ln P(y|x)}{\partial x_j} dy$$
$$= \left(\frac{d\mu(x)}{dx}^T \Sigma^{-1}(x) \frac{d\mu(x)}{dx} \right)_{ij}$$

$$\frac{d\mu(x)}{dx} = \begin{pmatrix} \frac{d\mu_1(x)}{dx_1} & \frac{d\mu_1(x)}{dx_2} & \dots & \frac{d\mu_1(x)}{dx_n} \\ \frac{d\mu_2(x)}{dx_1} & \frac{d\mu_2(x)}{dx_2} & \dots & \frac{d\mu_2(x)}{dx_n} \\ \dots & \dots & \dots & \dots \\ \frac{d\mu_m(x)}{dx_1} & \frac{d\mu_m(x)}{dx_2} & \dots & \frac{d\mu_m(x)}{dx_n} \end{pmatrix} = \begin{pmatrix} InstRF_1 \\ InstRF_2 \\ \dots \\ InstRF_m \end{pmatrix}$$

Information geometry

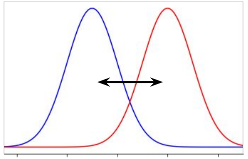

Metric tensor

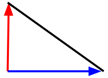

Euclidean: $$ds^2 = dx^2 + dy^2 + dz^2$$

$$= \begin{pmatrix} dx & dy & dz \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} dx \\ dy \\ dz \end{pmatrix}$$

Spherical: $$ds^2 = r^2 d\phi^2 + r^2 \sin^2 \phi d\theta^2$$

$$= \begin{pmatrix} d\phi & d\theta \end{pmatrix} \begin{pmatrix} r^2 & 0 \\ 0 & r^2 \sin^2 \phi \end{pmatrix} \begin{pmatrix} d\phi \\ d\theta \end{pmatrix}$$

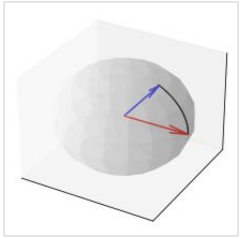

Information geometry:

$$ds^2 := 2 \cdot D_{KL}[P_\theta(y) || P_{\theta+d\theta}(y)]$$

Linearization:

$$\begin{aligned} 2 \cdot D_{KL}[P_\theta(y)||P_{\theta+d\theta}(y)] &= 2 \int P_\theta(y) \ln \frac{P_\theta(y)}{P_{\theta+d\theta}(y)} dy \\ &= -2 \int P_\theta(y) \ln[1 + \frac{dP_\theta(y)}{P_\theta(y)}] dy \\ &\approx \int P_\theta(y) [d \ln P_\theta(y)]^2 dy \\ &= \sum_{ij} \int P_\theta(y) \frac{\partial \ln P_\theta(y)}{\partial \theta_i} \frac{\partial \ln P_\theta(y)}{\partial \theta_j} dy d\theta_i d\theta_j \\ &= d\theta^T \cdot I(\theta) \cdot d\theta, \end{aligned}$$

Fisher information matrix as Riemannian metric

Theory for computing Fisher information matrix

Challenge: Dimension of natural scene stimuli is huge

Theorem: the top N most discriminable directions can all be linearly combined by instantaneous receptive fields, where N is the number of output neurons.

Proof: We use $IR := span(\{\frac{d\mu_i}{dx}\})$ to represent the subspace spanned by instantaneous receptive fields. Suppose the most discriminable direction v_1 is not in IR, then we can decompose $v_1 = v_1^o + \alpha v_1^p$, where v_1^o is the orthogonal component in $IR^\perp$ and αv_1^p is the parallel component in IR, $||v_1|| = ||v_1^p|| = 1$, $|\alpha| < 1$. Then

$$\begin{aligned} g^{stim}(v_1, v_1) &= \alpha^2 g^{stim}(v_1^p, v_1^p) \\ &< g^{stim}(v_1^p, v_1^p), \end{aligned} \tag{6}$$

which contradicts with the fact that v_1 is the most discriminable direction. Therefore v_1 is in IR.

Replacing IR with $IR \cap \{v_1\}^\perp$, with the same analysis we can prove that the second most discriminable direction v_2 is also in IR. Subsequently, the whole theorem can be proven.

MDI varies greatly across stimuli

MDI: Eigenvectors of the Fisher information matrix

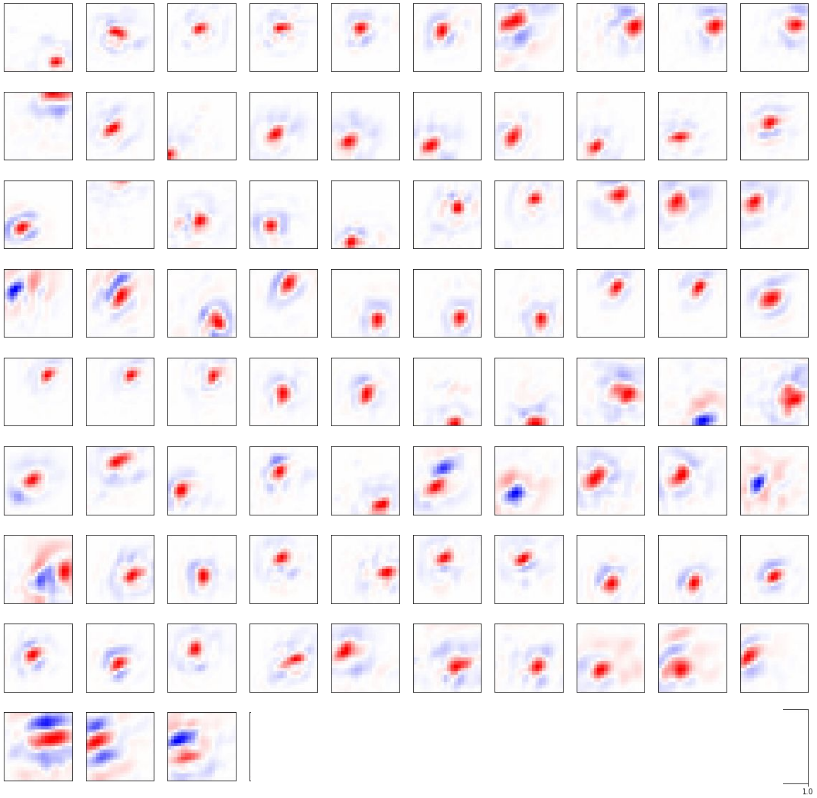

Retinal noise correlation limits information coding

Debate:
Noise correlation limits information coding (da Silveira & Rieke 2021, Kanitscheider et al. 2015)
Noise correlation may not be information-limiting (Moreno-Bote et al. 2014, Cafaro & Rieke 2010)

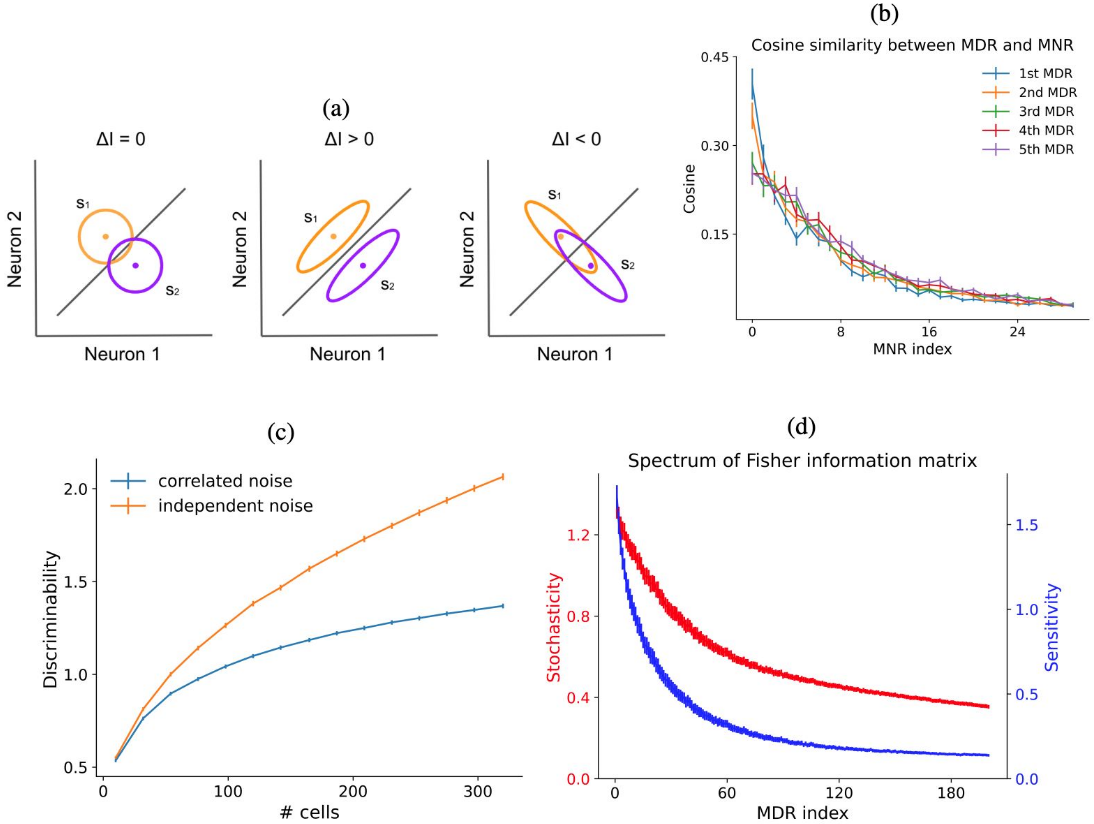

Firing rate dependency

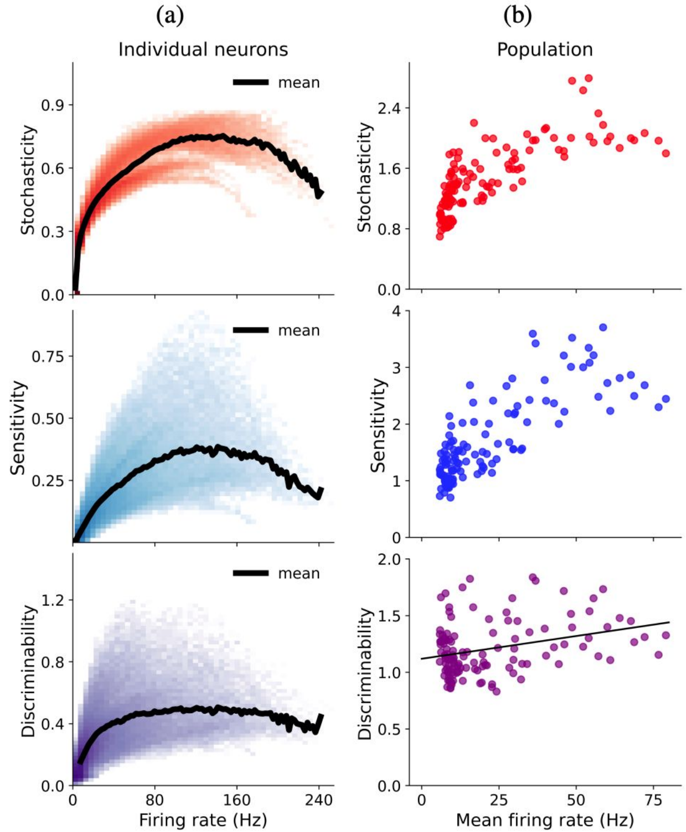

- Complementary coding
- Flat discriminability implies information maximization

Thank you!