Laboratoire des Systémes Perceptifs
École Normale Supérieure

Texture Interpolation for Probing Visual Perception

Jonathan Vacher

Collaborators: Aida Davila, Pascal Mamassian, Adam Kohn, Ruben Coen-Cagli

SMAI - Minisymposium: Transport Optimal pour l'Inférence Statistique

Vision Studies: Neurophysiology and Psychophysics

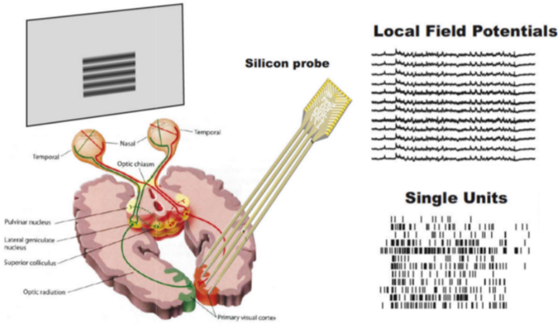

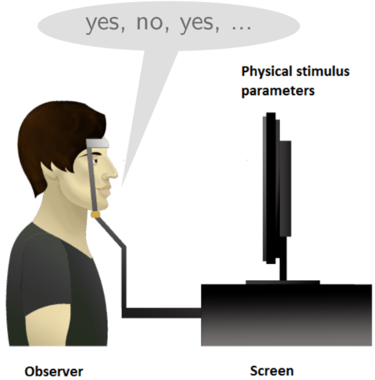

stimulus $\rightarrow$ $\boxed{??}$ $\rightarrow$ responses

Visual Stimuli

Noise

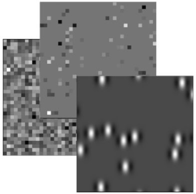

Artificial

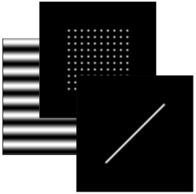

Natural

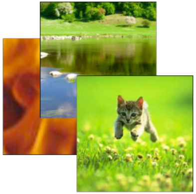

Visual Stimuli

Noise

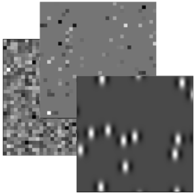

Natural Images

Artificial

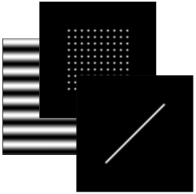

Natural

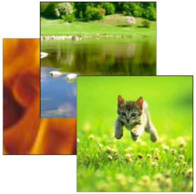

Visual Stimuli

Noise

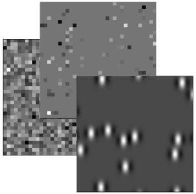

Natural Images

Textures

Artificial

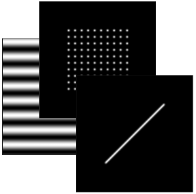

Natural

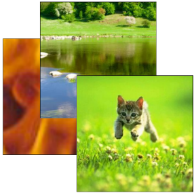

Visual Stimuli

Noise

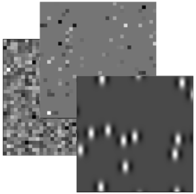

Natural Images

Artificial

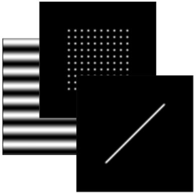

Natural

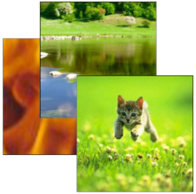

Textures

Segmentation/Contours

How Does the Brain Work ?

Philosophy of mind:

- computational functionalism (Putnam 1967) $\rightarrow$ brain=*information processing system*
- computationalist $\neq$ functionalism (Piccinini 2004; Piccinini 2009; Piccinini and Bahar 2013))

How Does the Brain Work ?

Philosophy of mind:

- computational functionalism (Putnam 1967) $\rightarrow$ brain=*information processing system*
- computationalist $\neq$ functionalism (Piccinini 2004; Piccinini 2009; Piccinini and Bahar 2013))

Hypotheses:

- Perceptions are internal representations (R) of the external world (S): $R = f(S)$
- Measurements are noisy internal representations: $M_R = R + N$
- The Brain infers the causes of these measurements (Bayes rule): $\mathbb{P}_{S|M_R} = \mathbb{P}_{M_R|S}\mathbb{P}_S / \mathbb{P}_M$

How Does the Brain Work ?

Philosophy of mind:

- computational functionalism (Putnam 1967) $\rightarrow$ brain=*information processing system*
- computationalist $\neq$ functionalism (Piccinini 2004; Piccinini 2009; Piccinini and Bahar 2013))

Hypotheses:

- Perceptions are internal representations (R) of the external world (S): $R = f(S)$
- Measurements are noisy internal representations: $M_R = R + N$
- The Brain infers the causes of these measurements (Bayes rule): $\mathbb{P}_{S|M_R} = \mathbb{P}_{M_R|S}\mathbb{P}_S/\mathbb{P}_M$

Consequences for the design of stimuli:

- Use our knowledge about the measurements performed by the Brain (neurophysiology)
- Build generative models $\mathbb{P}_{S|M_R}$

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)
- lack of theoretical and mathematical understanding of their effectiveness

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)
- lack of theoretical and mathematical understanding of their effectiveness

This will limit:

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)
- lack of theoretical and mathematical understanding of their effectiveness

This will limit:

- perceptual study

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)
- lack of theoretical and mathematical understanding of their effectiveness

This will limit:

- perceptual study
- modeling

Naturalistic Texture Synthesis Models Are Crucial to Study Vision

- Visual categorization (Balas 2008)
- Visual crowding (Balas et al. 2009)
- Visual search (Rosenholtz et al. 2012)
- A model of peripheral vision (Freeman and Simoncelli 2011; Rosenholtz 2016)
- Identifying regions with texture-selectivity (Freeman, Ziemba, et al. 2013; Okazawa et al. 2015)
- And more (Ziemba, Freeman, Movshon, et al. 2016; Ziemba, Freeman, Simoncelli, et al. 2018)

However:

- smart engineering design (Portilla et al. 2000; Gatys et al. 2015)
- lack of theoretical and mathematical understanding of their effectiveness

This will limit:

- perceptual study
- modeling
- understanding of neural activity

I. A statistical explanation of CNN-based texture synthesis.

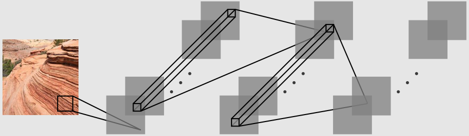

CNN activations have elliptical distributions

Exploring the Space of Textures

I. A statistical explanation of CNN-based texture synthesis.

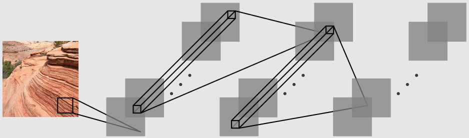

CNN activations have elliptical distributions

II. A mathematically grounded theory of texture interpolation.

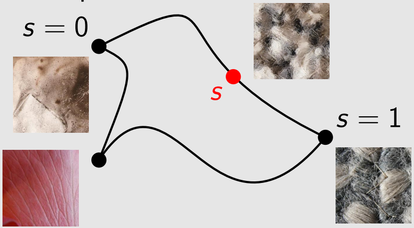

Exploring the Space of Textures

I. A statistical explanation of CNN-based texture synthesis.

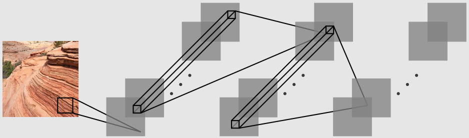

CNN activations have elliptical distributions

II. A mathematically grounded theory of texture interpolation.

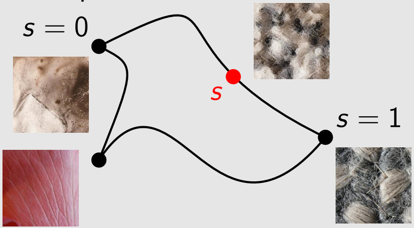

III. How are interpolation paths perceived ?

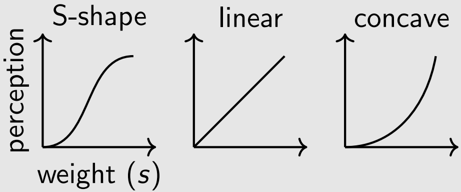

Exploring the Space of Textures

I. A statistical explanation of CNN-based texture synthesis.

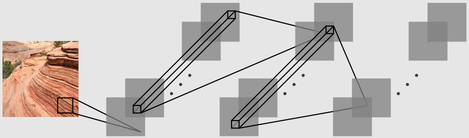

CNN activations have elliptical distributions

II. A mathematically grounded theory of texture interpolation.

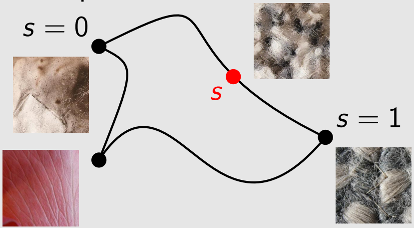

III. How are interpolation paths perceived ?

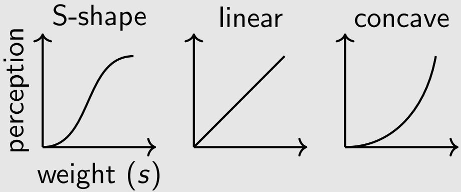

IV. How does neural population responses change along interpolation paths?

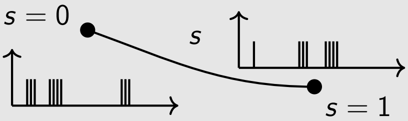

Distributions of Texture Activations

A distribution is elliptical if

$$\mathbb{P}_X(\mathbf{x}; \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X, g) \propto g((\mathbf{x} - \boldsymbol{\mu}_X)^{\mathrm{T}} \boldsymbol{\Sigma}_X^{-1} (\mathbf{x} - \boldsymbol{\mu}_X))$$

Texture *vs* Natural Images

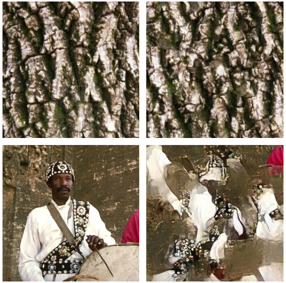

A distribution is elliptical if

$$\mathbb{P}_X(\mathbf{x}; \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X, g) \propto g((\mathbf{x} - \boldsymbol{\mu}_X)^{\mathrm{T}} \boldsymbol{\Sigma}_X^{-1} (\mathbf{x} - \boldsymbol{\mu}_X))$$

Texture *vs* Natural Images

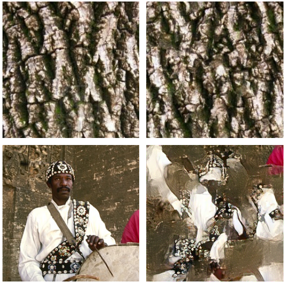

Steerable pyramid (Portilla et al. 2000)

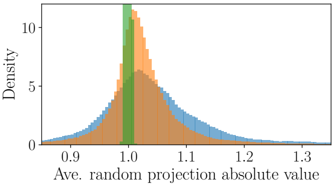

VGG19 (Simonyan et al. 2014)

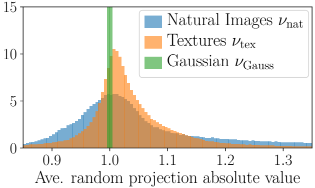

Distributions of Texture Activations

A distribution is elliptical if

$$\mathbb{P}_X(\mathbf{x}; \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X, g) \propto g((\mathbf{x} - \boldsymbol{\mu}_X)^{\mathrm{T}} \boldsymbol{\Sigma}_X^{-1} (\mathbf{x} - \boldsymbol{\mu}_X))$$

Texture *vs* Natural Images

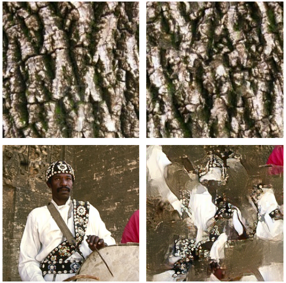

Steerable pyramid (Portilla et al. 2000)

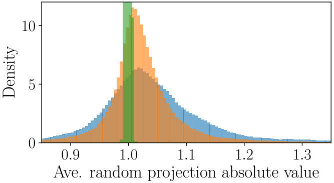

VGG19 (Simonyan et al. 2014)

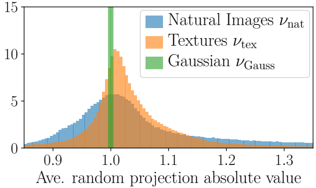

Texture CNN activations are more elliptically distributed than natural image CNN activations.

Theory of Texture Interpolation: Optimal Transport

Optimal transport $\rightarrow$ metrics over probability distributions,

Theory of Texture Interpolation: Optimal Transport

Optimal transport $\rightarrow$ metrics over probability distributions, *e.g.* the 2-Wasserstein distance:

$$W_2(\mathbb{P}_X, \mathbb{P}_Y)^2 = \inf_{\substack{\mathbb{P}_{X,Y} \\ X \sim \mathbb{P}_X, Y \sim \mathbb{P}_Y}} \mathbb{E}_{X,Y}\left(\|X - Y\|^2\right).$$

Theory of Texture Interpolation: Optimal Transport

Optimal transport $\rightarrow$ metrics over probability distributions, *e.g.* the 2-Wasserstein distance:

$$W_2(\mathbb{P}_X, \mathbb{P}_Y)^2 = \inf_{\substack{\mathbb{P}_{X,Y} \\ X \sim \mathbb{P}_X, Y \sim \mathbb{P}_Y}} \mathbb{E}_{X,Y}\left(\|X - Y\|^2\right).$$

For elliptical distributions (with the same radial function) we have

$$W_2(\mathbb{P}_X, \mathbb{P}_Y)^2 = \|\boldsymbol{\mu}_X - \boldsymbol{\mu}_Y\|^2 + \mathcal{B}(\boldsymbol{\Sigma}_X, \boldsymbol{\Sigma}_Y)^2$$

where $\mathcal{B}(\boldsymbol{\Sigma}_X, \boldsymbol{\Sigma}_Y)^2 = \mathrm{Tr}(\boldsymbol{\Sigma}_X + \boldsymbol{\Sigma}_Y - 2(\boldsymbol{\Sigma}_X^{1/2}\boldsymbol{\Sigma}_Y\boldsymbol{\Sigma}_X^{1/2})^{1/2})$ (Bures metric).

Theory of Texture Interpolation: Optimal Transport

Optimal transport $\rightarrow$ metrics over probability distributions, *e.g.* the 2-Wasserstein distance:

$$W_2(\mathbb{P}_X, \mathbb{P}_Y)^2 = \inf_{\substack{\mathbb{P}_{X,Y} \\ X \sim \mathbb{P}_X, Y \sim \mathbb{P}_Y}} \mathbb{E}_{X,Y}\left(\|X - Y\|^2\right).$$

For elliptical distributions (with the same radial function) we have

$$W_2(\mathbb{P}_X, \mathbb{P}_Y)^2 = \|\boldsymbol{\mu}_X - \boldsymbol{\mu}_Y\|^2 + \mathcal{B}(\boldsymbol{\Sigma}_X, \boldsymbol{\Sigma}_Y)^2$$

where $\mathcal{B}(\boldsymbol{\Sigma}_X, \boldsymbol{\Sigma}_Y)^2 = \mathrm{Tr}(\boldsymbol{\Sigma}_X + \boldsymbol{\Sigma}_Y - 2(\boldsymbol{\Sigma}_X^{1/2}\boldsymbol{\Sigma}_Y\boldsymbol{\Sigma}_X^{1/2})^{1/2})$ (Bures metric).

To get some intuition, when $\boldsymbol{\Sigma}_X$ and $\boldsymbol{\Sigma}_Y$ commute (*i.e.* $\boldsymbol{\Sigma}_X\boldsymbol{\Sigma}_Y = \boldsymbol{\Sigma}_Y\boldsymbol{\Sigma}_X$):

$$\mathcal{B}(\boldsymbol{\Sigma}_X, \boldsymbol{\Sigma}_Y) = \left\|\boldsymbol{\Sigma}_X^{1/2} - \boldsymbol{\Sigma}_Y^{1/2}\right\|_f.$$

Given a texture example u and an input white noise image v, we write

Given a texture example u and an input white noise image v, we write

The Wasserstein Loss

$$L_{\mathrm{W}}(u,v) = \sum_{l=1}^{L} \| M_{N_l}(\mathbf{X}_l^v) - \boldsymbol{\mu}_l^u \|^2 + \left\| C_{N_l}(\mathbf{X}_l^v)^{1/2} - \boldsymbol{\Sigma}_l^{u\,1/2} \right\|_f^2$$

Texture Synthesis and Interpolation

Given a texture example u and an input white noise image v, we write

The Wasserstein Loss

$$L_{\mathrm{W}}(u,v) = \sum_{l=1}^{L} \|M_{N_l}(\mathbf{X}_l^v) - \boldsymbol{\mu}_l^u\|^2 + \left\| C_{N_l}(\mathbf{X}_l^v)^{1/2} - \boldsymbol{\Sigma}_l^{u\,1/2} \right\|_f^2$$

Weighted barycenter $\rightarrow$ interpolation. For $\mathrm{m} \in \{\mathrm{G}, \mathrm{W}\}$:

$$\bar{v}_{\mathrm{m}} = \underset{v}{\mathrm{argmin}} \sum_{k=1}^{K} \lambda_k L_{\mathrm{m}}(u_k, v) \quad \text{where} \quad \sum_{i=1}^{K} \lambda_i = 1$$

Interpolation Examples

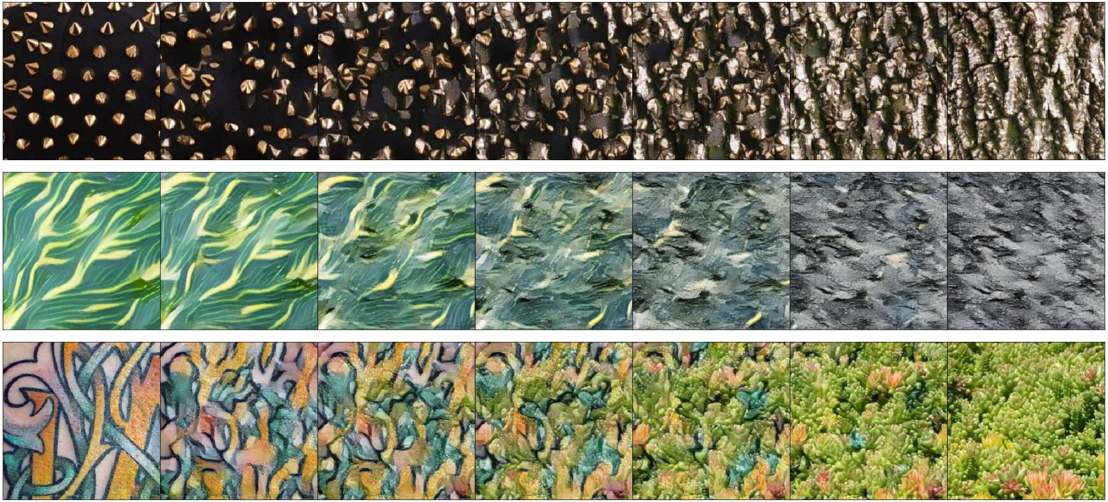

Interpolation Examples

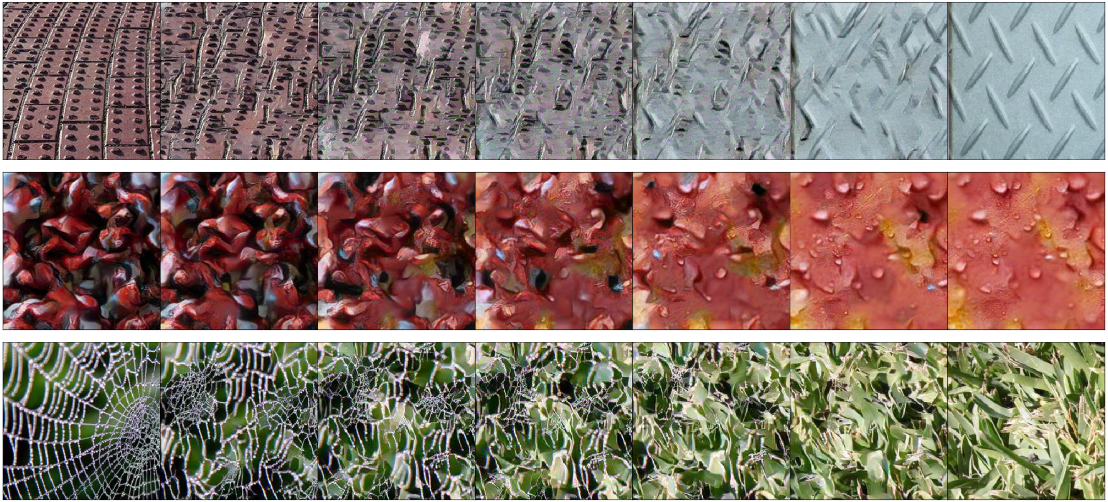

Interpolation Examples

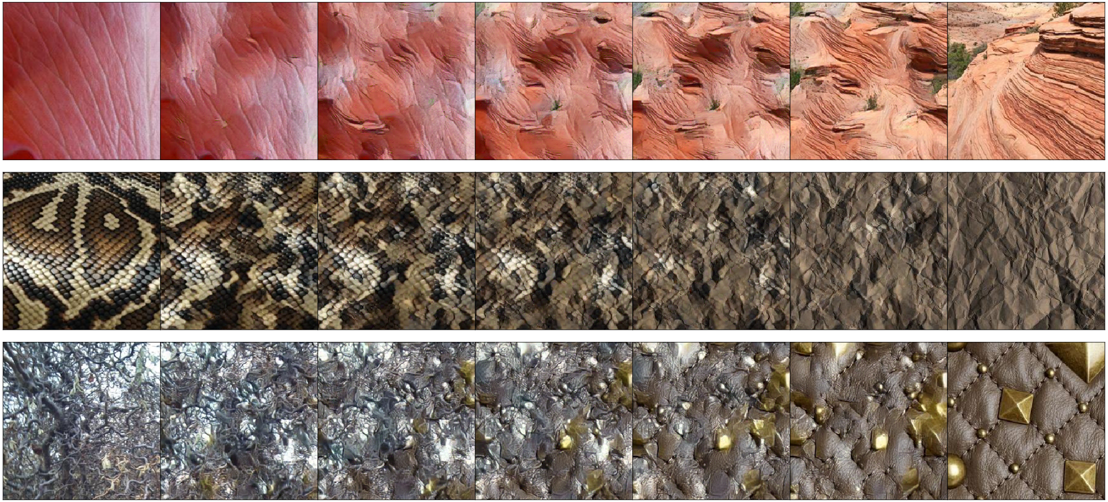

The Wasserstein Loss

$$L_{\mathrm{W}}(u,v) = \sum_{l=1}^{L} \|M_{N_l}(\mathbf{X}_l^v) - \boldsymbol{\mu}_l^u\|^2 + \left\|C_{N_l}(\mathbf{X}_l^v)^{1/2} - \boldsymbol{\Sigma}_l^{u\,1/2}\right\|_f^2$$

$M_{N_l}(X)$ and $C_{N_l}(X)$: empirical mean, covariance of X.

Weighted barycenter $\rightarrow$ interpolation. Given a collection of textures $(u_1, \dots, u_K)$:

$$\text{For m} \in \{\mathrm{G}, \mathrm{W}\}, \quad \bar{v}_{\mathrm{m}} = \underset{v}{\operatorname{argmin}} \sum_{k=1}^{K} \lambda_k L_{\mathrm{m}}(u_k, v) \quad \text{where} \quad \sum_{i=1}^{K} \lambda_i = 1$$

jonathan.vacher@ens.fr

https://jonathanvacher.github.io

Texture Synthesis and Interpolation

The Wasserstein Loss

$$L_{\mathrm{W}}(u, v) = \sum_{l=1}^{L} \| M_{N_l}(\mathbf{X}_l^v) - \boldsymbol{\mu}_l^u \|^2 + \left\| C_{N_l}(\mathbf{X}_l^v)^{1/2} - \boldsymbol{\Sigma}_l^{u\,1/2} \right\|_f^2$$

$M_{N_l}(X)$ and $C_{N_l}(X)$: empirical mean, covariance of X.

Weighted barycenter $\to$ interpolation. Given a collection of textures $(u_1, \dots, u_K)$:

$$\text{For } \mathrm{m} \in \{\mathrm{G}, \mathrm{W}\}, \quad \bar{v}_{\mathrm{m}} = \operatorname*{argmin}_v \sum_{k=1}^{K} \lambda_k L_{\mathrm{m}}(u_k, v) \quad \text{where} \quad \sum_{i=1}^{K} \lambda_i = 1$$

The Gram Loss (Gatys et al. 2015)

$$L_{\mathrm{G}}(u, v) = \sum_{l=1}^{L} \| G_{N_l}(\mathbf{X}_l^v) - \mathbf{G}_l^u \|_f^2$$

$G_{N_l}(X)$: empirical Gram matrix of X.

Texture Synthesis and Interpolation

The Wasserstein Loss

$$L_{\mathrm{W}}(u,v) = \sum_{l=1}^{L} \|M_{N_l}(\mathbf{X}_l^v) - \boldsymbol{\mu}_l^u\|^2 + \left\| C_{N_l}(\mathbf{X}_l^v)^{1/2} - \boldsymbol{\Sigma}_l^{u\,1/2} \right\|_f^2$$

$M_{N_l}(X)$ and $C_{N_l}(X)$: empirical mean, covariance of X.

Weighted barycenter $\rightarrow$ interpolation. Given a collection of textures $(u_1, \dots, u_K)$:

$$\text{For m} \in \{\mathrm{G}, \mathrm{W}\}, \quad \bar{v}_{\mathrm{m}} = \operatorname*{argmin}_v \sum_{k=1}^{K} \lambda_k L_{\mathrm{m}}(u_k, v) \quad \text{where} \quad \sum_{i=1}^{K} \lambda_i = 1$$

The Gram Loss (Gatys et al. 2015)

$$L_{\mathrm{G}}(u,v) = \sum_{l=1}^{L} \|G_{N_l}(\mathbf{X}_l^v) - \mathbf{G}_l^u\|_f^2$$

$G_{N_l}(X)$: empirical Gram matrix of X.

A parametric algorithm (Portilla et al. 2000)

$$\bar{v} = \mathrm{Synthesize}\left(\sum_{k=1}^{K} \lambda_k P_k\right)$$

(See also Vacher et al. 2020)

Different Loss Functions Lead to Different Interpolation Paths

- Not all interpolation paths are relevant for perceptual studies !

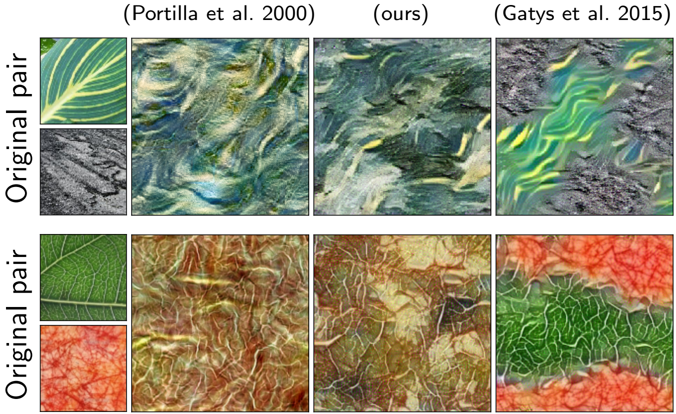

Interpolation between spectrally-matched noise and naturalistic textures:

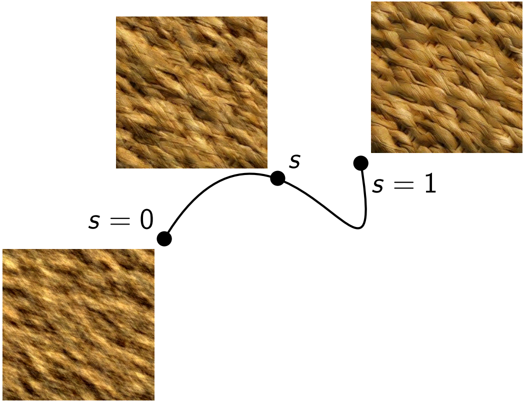

Interpolation between spectrally-matched noise and naturalistic textures:

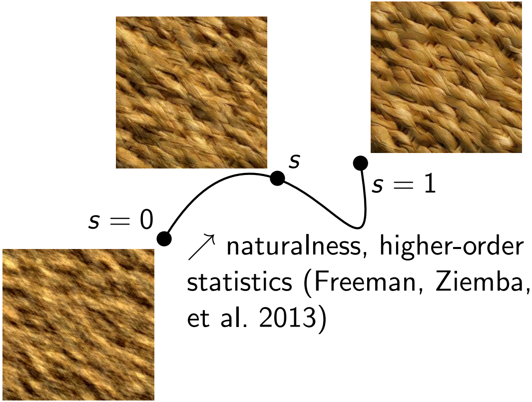

Interpolation between spectrally-matched noise and naturalistic textures:

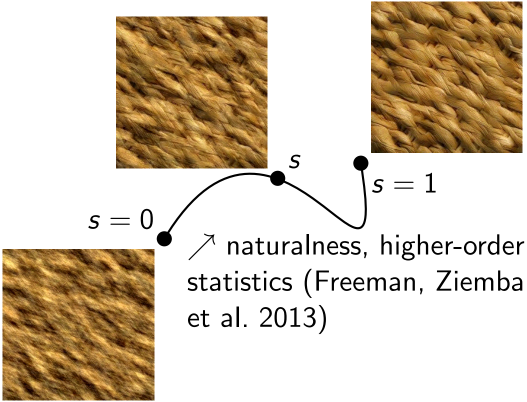

$\nearrow$ naturalness, higher-order statistics (Freeman, Ziemba, et al. 2013)

Interpolation between arbitrary texture pair:

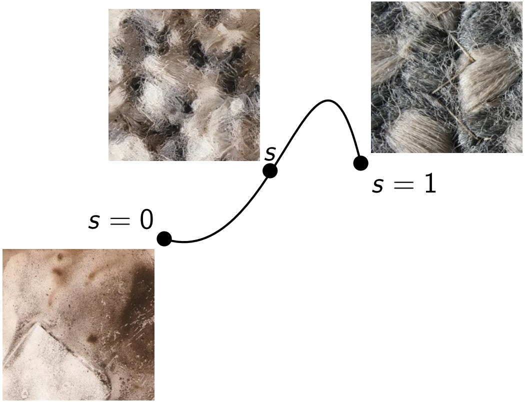

Interpolation between spectrally-matched noise and naturalistic textures:

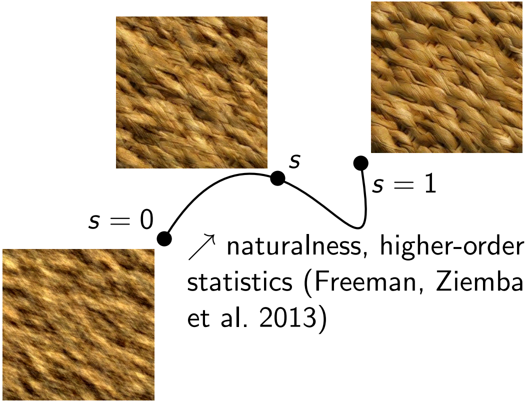

↗ naturalness, higher-order statistics (Freeman, Ziemba, et al. 2013)

Interpolation between arbitrary texture pair:

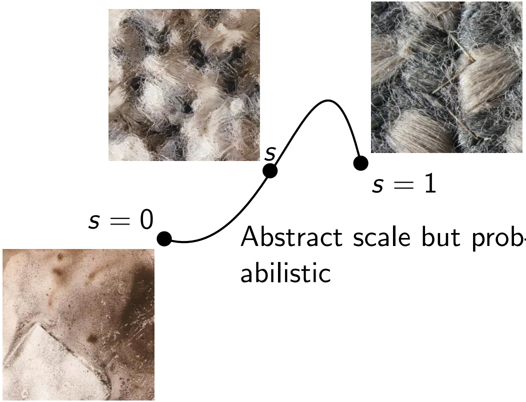

Abstract scale but probabilistic

Visual Perception and Neural Activity

Interpolation between spectrally-matched noise and naturalistic textures:

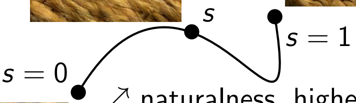

$\nearrow$ naturalness, higher-order statistics (Freeman, Ziemba, et al. 2013)

How are these paths perceived?

Interpolation between arbitrary texture pair:

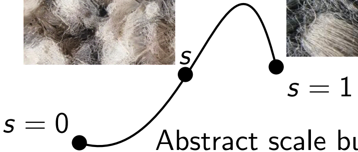

Abstract scale but probabilistic

How does neural decoding performances change along interpolation paths?

A Summary of the Visual Cortex

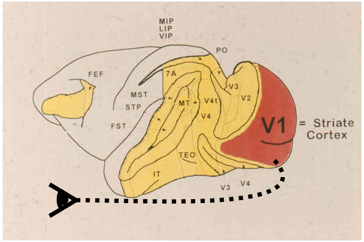

A Summary of the Visual Cortex

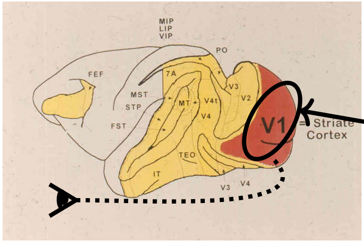

V1: wavelet like features (orientation, spatial/temporal frequencies (Ringach 2002))

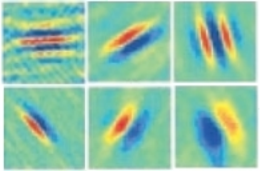

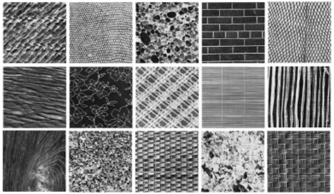

V2/V4: higher-order stats (textures, curvature, crossing (Freeman, Ziemba, et al. 2013))

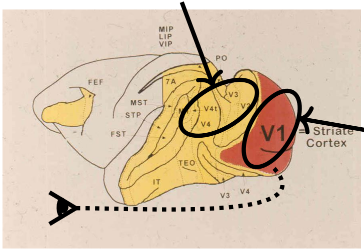

V1: wavelet like features (orientation, spatial/temporal frequencies (Ringach 2002))

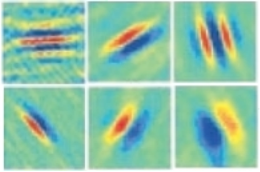

IT: objects (faces, body, animal, natural, man-made (Sato et al. 2013))

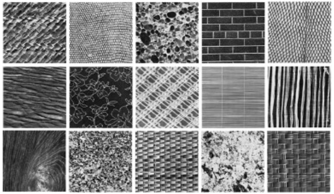

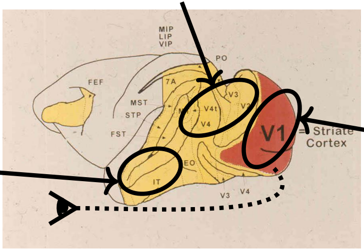

V2/V4: higher-order stats (textures, curvature, crossing (Freeman, Ziemba, et al. 2013))

V1: wavelet like features (orientation, spatial/temporal frequencies (Ringach 2002))

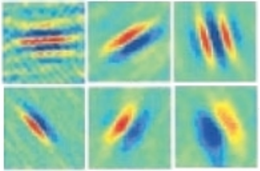

Sensitivity of Mid-level Cortical Neurons to Texture Interpolation

- Neuronal responses to 5 weights $s \in \{0.0, 0.3, 0.5, 0.7, 1.0\}$ on the interpolation path.
- Recordings: awake macaque monkey implanted with “Utah” arrays in V1 and V4

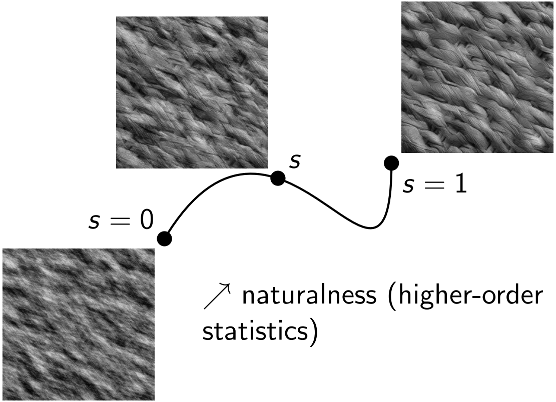

$\nearrow$ naturalness (higher-order statistics)

- Neuronal responses to 5 weights $s \in \{0.0, 0.3, 0.5, 0.7, 1.0\}$ on the interpolation path.
- Recordings: awake macaque monkey implanted with "Utah" arrays in V1 and V4

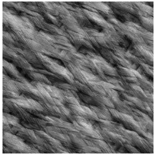

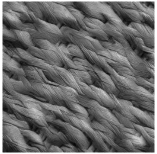

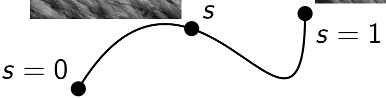

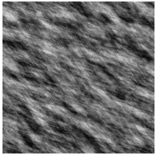

$\nearrow$ naturalness (higher-order statistics)

Decoding of weights 0 and $s = 0.3, 0.5, 0.7$ and 1.0, averaged over 3 textures.

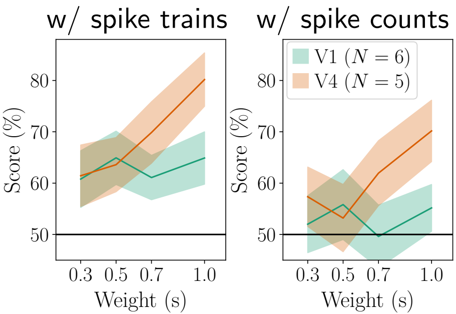

Measuring a Perceptual Scale

- Maximum Likelihood Difference Scaling (MLDS, Maloney et al. 2003)

- Maximum Likelihood Difference Scaling (MLDS, Maloney et al. 2003)

Hypothesis: there exists a monotonic function $f : [0,1] \longrightarrow [0,1]$ that maps the interpolation path to perception.

Measuring a Perceptual Scale

- Maximum Likelihood Difference Scaling (MLDS, Maloney et al. 2003)

Hypothesis: there exists a monotonic function $f : [0,1] \longrightarrow [0,1]$ that maps the interpolation path to perception.

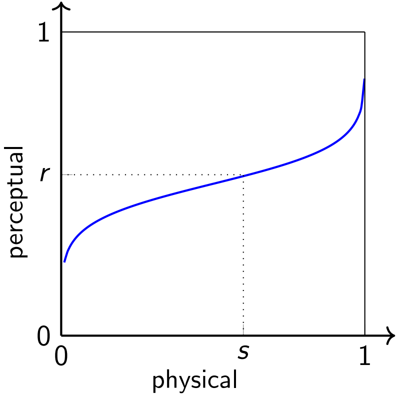

Measuring a Perceptual Scale

- Maximum Likelihood Difference Scaling (MLDS, Maloney et al. 2003)

Hypothesis: there exists a monotonic function $f : [0,1] \longrightarrow [0,1]$ that maps the interpolation path to perception.

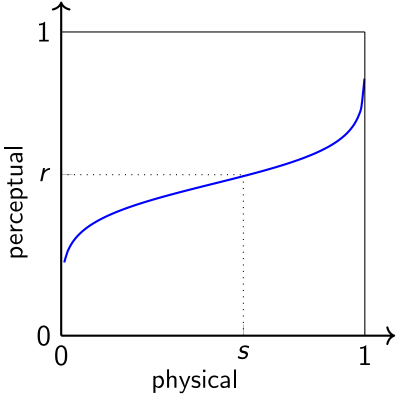

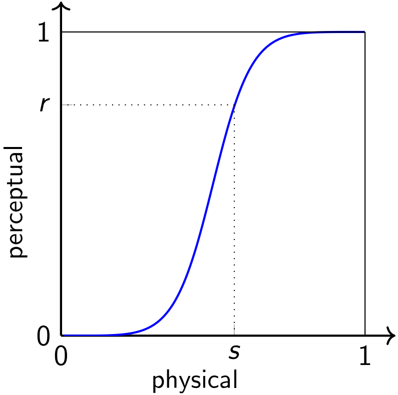

jonathan.vacher@ens.fr

https://jonathanvacher.github.io

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, {\sigma'}^2).$$

Protocol: samples $\{s_1, \ldots, s_N\} \subset [0, 1]$ Perceptual scale inference:

Measuring a Perceptual Scale

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2).$$

Protocol: samples $\{s_1, \dots, s_N\} \subset [0, 1]$

Perceptual scale inference:

- present $s_i < s_j < s_k$ to an observer

Measuring a Perceptual Scale

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2).$$

Protocol: samples $\{s_1, \dots, s_N\} \subset [0, 1]$

Perceptual scale inference:

- present $s_i < s_j < s_k$ to an observer

- ask which of the two pairs (s_i, s_j) and (s_j, s_k) is more similar $\to$ estimating $p_{i,j,k}$

Measuring a Perceptual Scale

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2).$$

Protocol: samples $\{s_1, \ldots, s_N\} \subset [0, 1]$

- present $s_i < s_j < s_k$ to an observer

- ask which of the two pairs (s_i, s_j) and (s_j, s_k) is more similar $\rightarrow$ estimating $p_{i,j,k}$

Perceptual scale inference:

- the observer compares the perceptual differences

$$\Delta_{i,j,k} = |f(s_i) - f(s_j)| - |f(s_j) - f(s_k)| + \varepsilon$$

where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$.

Measuring a Perceptual Scale

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2).$$

Protocol: samples $\{s_1, \ldots, s_N\} \subset [0,1]$

- present $s_i < s_j < s_k$ to an observer

- ask which of the two pairs (s_i, s_j) and (s_j, s_k) is more similar $\rightarrow$ estimating $p_{i,j,k}$

Perceptual scale inference:

- the observer compares the perceptual differences

$$\Delta_{i,j,k} = |f(s_i) - f(s_j)| - |f(s_j) - f(s_k)| + \varepsilon$$

where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$.

- $\mathbb{E}(\Delta_{i,j,k} \geqslant 0) = \Phi\left(\frac{f(s_i) - 2f(s_j) + f(s_k)}{\sigma}\right) \simeq p_{i,j,k}$

Measuring a Perceptual Scale

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2).$$

Protocol: samples $\{s_1, \dots, s_N\} \subset [0, 1]$

- present $s_i < s_j < s_k$ to an observer

- ask which of the two pairs (s_i, s_j) and (s_j, s_k) is more similar $\rightarrow$ estimating $p_{i,j,k}$

Perceptual scale inference:

- the observer compares the perceptual differences

$$\Delta_{i,j,k} = |f(s_i) - f(s_j)| - |f(s_j) - f(s_k)| + \varepsilon$$

where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$.

- $\mathbb{E}(\Delta_{i,j,k} \geqslant 0) = \Phi\left(\frac{f(s_i) - 2f(s_j) + f(s_k)}{\sigma}\right) \simeq p_{i,j,k}$
- solving a constrained linear system (Knoblauch et al. 2008)

$$MF = \Phi^{-1}(P) \quad s.t. \quad F' \geqslant 0$$

A Meaningful Perceptual Scale

- a meaningful perceptual scale compared to a random observer

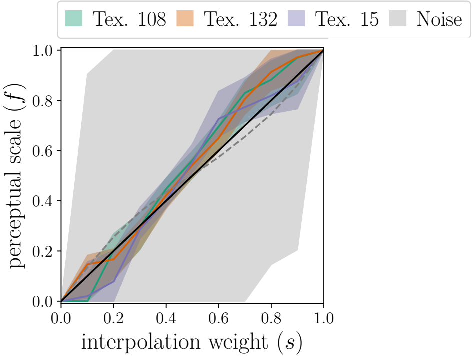

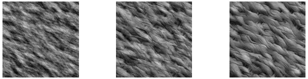

A Meaningful Perceptual Scale

- a meaningful perceptual scale compared to a random observer

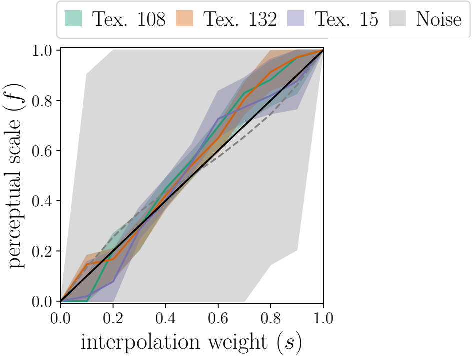

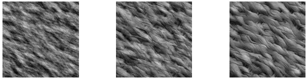

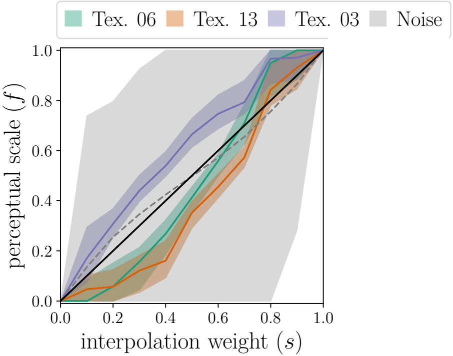

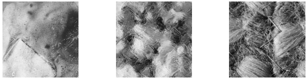

A Meaningful Perceptual Scale

- a meaningful perceptual scale compared to a random observer

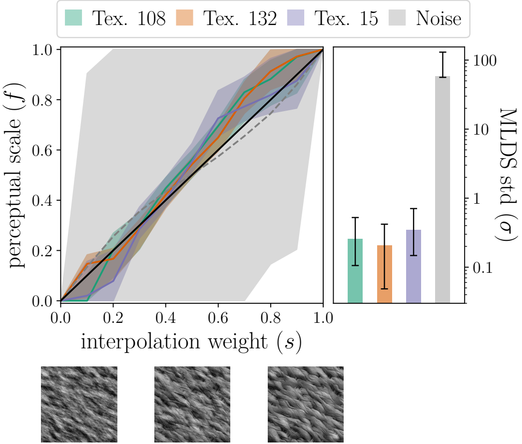

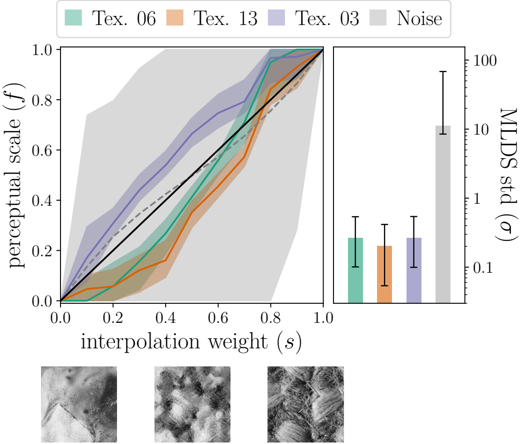

Equation of perception:

$$M_R = R + N = f(S) + N \quad \text{with} \quad N \sim \mathcal{N}(0, \sigma'^2). \tag{1}$$

Fisher Information:

$$\mathcal{I}_S(s) \overset{\text{def.}}{=} \mathbb{E}_{M_R|S}\left(\left(\frac{\partial \log(\mathbb{P}_{M_R|S})}{\partial s}(M_R, s)\right)^2\right) = \mathbb{E}_{M_R|S}\left(\frac{\partial^2 \log(\mathbb{P}_{M_R|S})}{\partial s^2}(M_R, s)\right)$$

$\longrightarrow$ Variance of the log-likelihood derivative or local curvature $\sim$ estimator precision !

Proposition ($\sim$ MLDS assumption)

Assume Equation (1), then $\mathcal{I}_R$ is constant if and only if for all $s \in \mathcal{S} = [0, 1]$

$$f(s) \propto \int_0^s \sqrt{\mathcal{I}_S(t)} \mathrm{d}t.$$

Relation to our Texture Interpolation Model

Interpolation between texture pair:

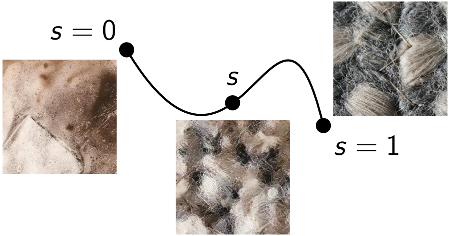

Relation to our Texture Interpolation Model

Interpolation between texture pair:

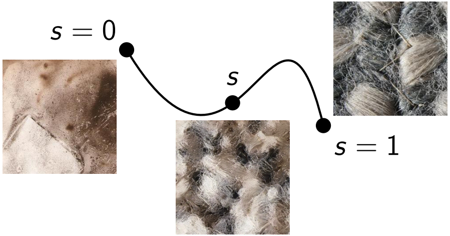

Remain Elliptic along the interpolation path !

$$\mu_s = s\mu_1 + (1-s)\mu_0; \;\; \Sigma_s = \Sigma_1^{-1/2} U_s^2 \Sigma_1^{-1/2}$$

where $U_s = s\Sigma_1 + (1-s)\left(\Sigma_1^{1/2}\Sigma_0\Sigma_1^{1/2}\right)^{1/2}$

Interpolation between texture pair:

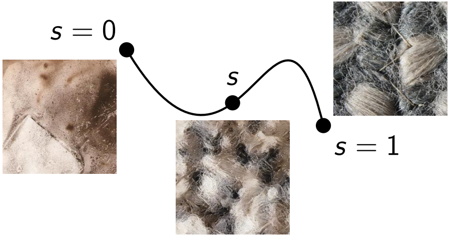

Remain Elliptic along the interpolation path !

$$\mu_s = s\mu_1 + (1-s)\mu_0; \quad \Sigma_s = \Sigma_1^{-1/2} U_s^2 \Sigma_1^{-1/2}$$

where $U_s = s\Sigma_1 + (1-s)\left(\Sigma_1^{1/2}\Sigma_0\Sigma_1^{1/2}\right)^{1/2}$

The Fisher Information is given by

$$\mathcal{I}(s) = \mu'\Sigma_s^{-1}\mu' + \frac{1}{2}\,\mathrm{tr}\left(\Sigma_s^{-1}\Sigma_s'\Sigma_s^{-1}\Sigma_s'\right)$$

where $\mu' = {}^{\mathrm{d}\mu_s}/_{\mathrm{d}s} = \mu_1 - \mu_0$, $\Sigma_s' = {}^{\mathrm{d}\Sigma_s}/_{\mathrm{d}s} = 2s(\Sigma_0 + \Sigma_1 - Q_0 - Q_1) + Q_1 + Q_2 - 2\Sigma_1$ with

$$Q_0 = \Sigma_0^{1/2}\left(\Sigma_0^{1/2}\Sigma_1\Sigma_0^{1/2}\right)\Sigma_0^{-1/2}$$

and

$$Q_1 = \Sigma_0^{-1/2}\left(\Sigma_0^{1/2}\Sigma_1\Sigma_0^{1/2}\right)\Sigma_0^{1/2}.$$

Relation to our Texture Interpolation Model

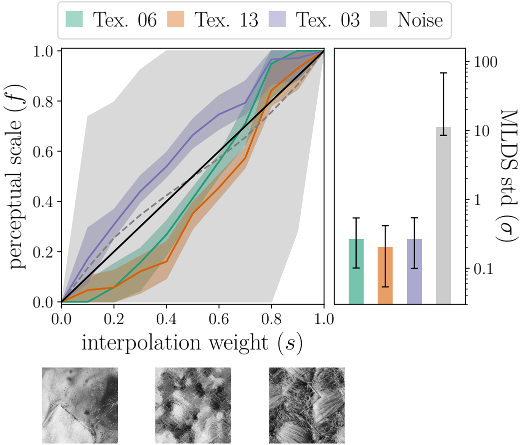

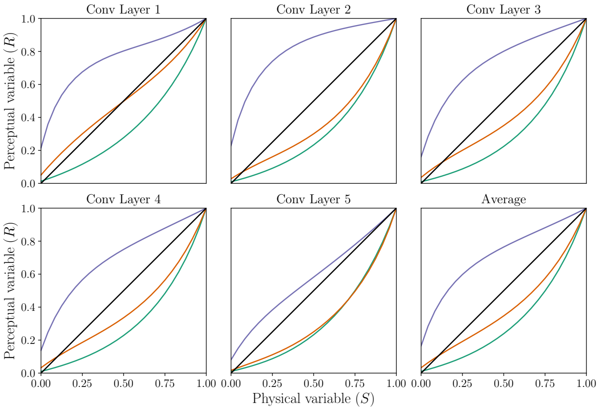

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1 - s)a_0$ with $s \in [0, 1]$.

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1 - s)a_0$ with $s \in [0, 1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1 - s)a_0$ with $s \in [0, 1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Then,

$$f(s) = \frac{\log\left(\frac{a(s)}{a(0)}\right)}{\log\left(\frac{a(1)}{a(0)}\right)} = \frac{\log\left(\frac{a_0+(a_1-a_0)s}{a_0}\right)}{\log\left(\frac{a_1}{a_0}\right)}$$

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1-s)a_0$ with $s \in [0,1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Then,

$$f(s) = \frac{\log\left(\frac{a(s)}{a(0)}\right)}{\log\left(\frac{a(1)}{a(0)}\right)} = \frac{\log\left(\frac{a_0+(a_1-a_0)s}{a_0}\right)}{\log\left(\frac{a_1}{a_0}\right)}$$

$\sim$ Weber-Fechner law: $\mathrm{d}f = \frac{\mathrm{d}s}{s}$.

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1-s)a_0$ with $s \in [0,1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Then,

$$f(s) = \frac{\log\left(\frac{a(s)}{a(0)}\right)}{\log\left(\frac{a(1)}{a(0)}\right)} = \frac{\log\left(\frac{a_0 + (a_1 - a_0)s}{a_0}\right)}{\log\left(\frac{a_1}{a_0}\right)}$$

$\sim$ Weber-Fechner law: $\mathrm{d}f = \frac{\mathrm{d}s}{s}$.

Combining stimuli intensities

$$X(s) = a(s)(U^T, 0)^T + b(s)(0, V^T)^T$$

where $U \sim \mathcal{N}(\mu_U, \Sigma_U)$, $V \sim \mathcal{N}(\mu_V, \Sigma_V)$.

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1-s)a_0$ with $s \in [0, 1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Then,

$$f(s) = \frac{\log\left(\frac{a(s)}{a(0)}\right)}{\log\left(\frac{a(1)}{a(0)}\right)} = \frac{\log\left(\frac{a_0+(a_1-a_0)s}{a_0}\right)}{\log\left(\frac{a_1}{a_0}\right)}$$

$\sim$ Weber-Fechner law: $\mathrm{d}f = \frac{\mathrm{d}s}{s}$.

Combining stimuli intensities

$$X(s) = a(s)(U^T, 0)^T + b(s)(0, V^T)^T$$

where $U \sim \mathcal{N}(\mu_U, \Sigma_U)$, $V \sim \mathcal{N}(\mu_V, \Sigma_V)$. Then,

$$\mathcal{I}(s) = (\mu_U^T \Sigma_U^{-1} \mu_U + d)\frac{a'(s)^2}{a(t)^2} + (\mu_V^T \Sigma_V^{-1} \mu_V + d)\frac{b'(s)^2}{b(t)^2}$$

Future Work: Lower Dimension Parametric Stimuli

Stimulus intensity

$$X(s) = a(s)X \quad \text{where} \quad X \sim \mathcal{N}(\mu, \Sigma)$$

and $a(s) = sa_1 + (1-s)a_0$ with $s \in [0,1]$.

$\sim$ Wasserstein interpolation between $\mathcal{N}(a_0\mu, a_0^2\Sigma)$ and $\mathcal{N}(a_1\mu, a_1^2\Sigma)$ $(a_0, a_1 > 0)$.

Then,

$$f(s) = \frac{\log\left(\frac{a(s)}{a(0)}\right)}{\log\left(\frac{a(1)}{a(0)}\right)} = \frac{\log\left(\frac{a_0+(a_1-a_0)s}{a_0}\right)}{\log\left(\frac{a_1}{a_0}\right)}$$

$\sim$ Weber-Fechner law: $\mathrm{d}f = \frac{\mathrm{d}s}{s}$.

Combining stimuli intensities

$$X(s) = a(s)(U^T, 0)^T + b(s)(0, V^T)^T$$

where $U \sim \mathcal{N}(\mu_U, \Sigma_U)$, $V \sim \mathcal{N}(\mu_V, \Sigma_V)$. Then,

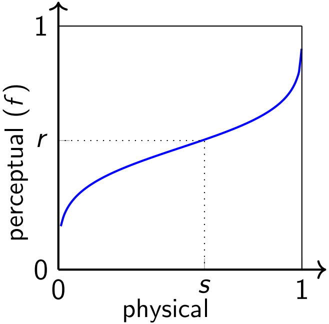

Fin.

Merci pour votre attention !

Empirical Convergence of the Radial Density.