CNRS - Toulouse INP - UT3 - UT1 - UT2
Institut de Recherche en Informatique de Toulouse

Multi-scale modeling of natural texture images

Sixin Zhang,
Université de Toulouse, INP, IRIT

Geometry and Data summer school, Strasbourg 2023

Scope of talk: image understanding

- How do we represent image data ?
- Is it possible to model natural images mathematically ?
  - Can we generate texture images using machines ?
  - How to understand images from physical models ?
- Given a large amount of data, how can we process them ?

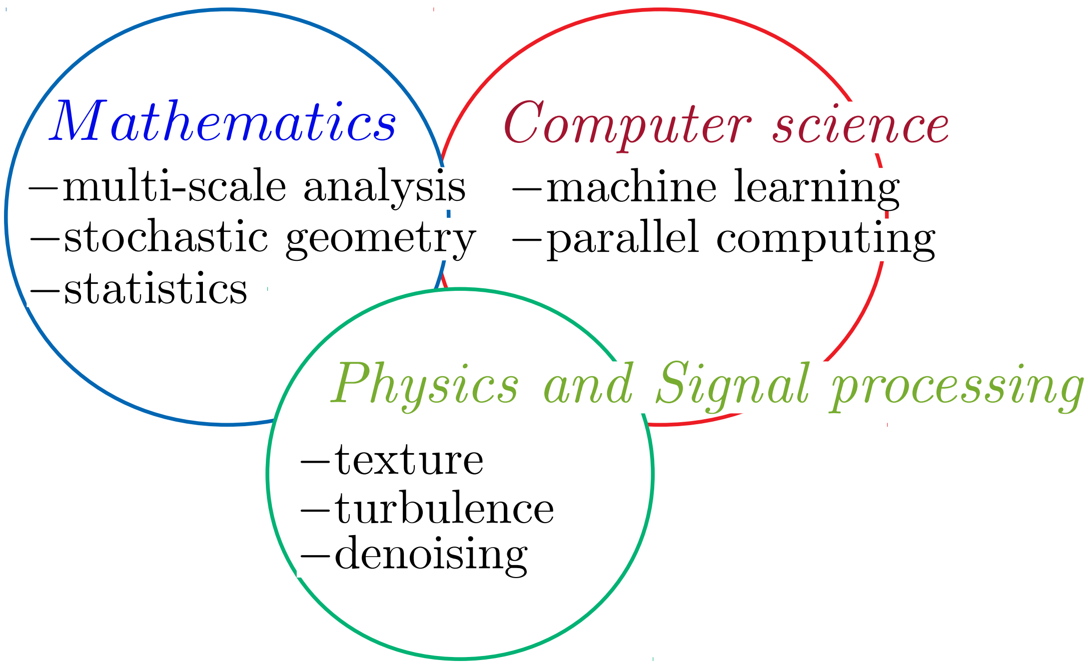

Texture synthesis problem *Julesz, 1962*

- Textures are spatially homogeneous images, consisting of similar patterns forming a coherent ensemble.

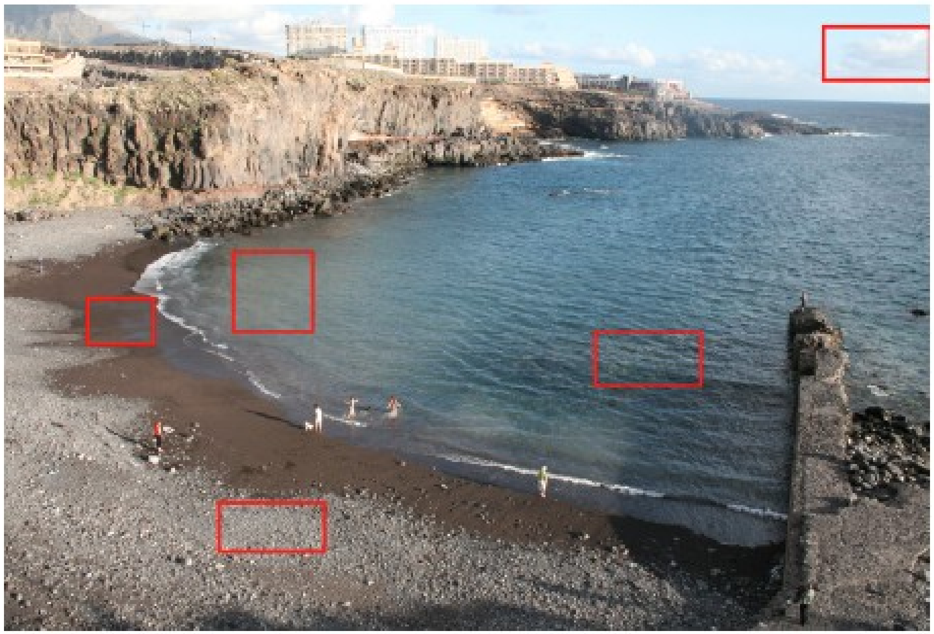

Galerne et al. 2011

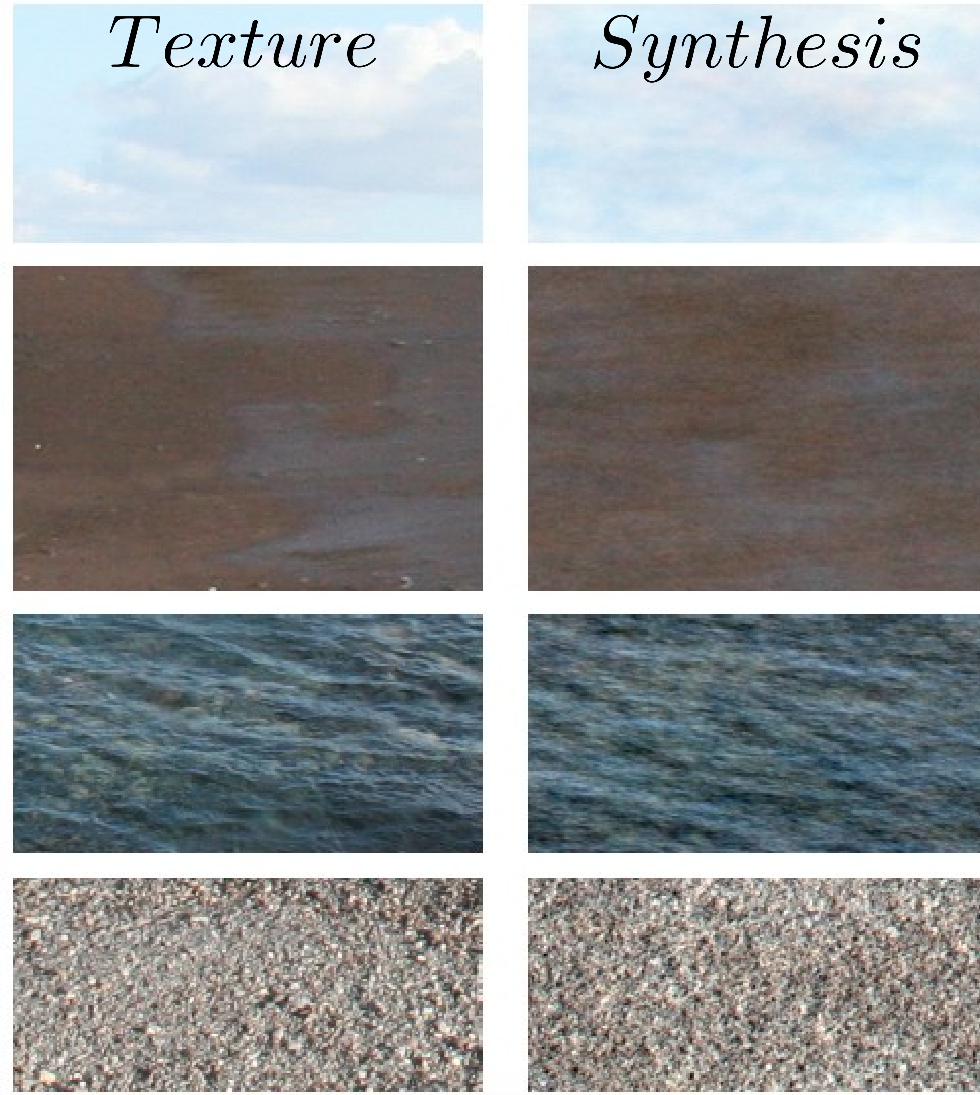

Turbulence modeling

Can we characterize coherent structures in turbulent flows?

- Simulate fluid vorticity by PDE models (Navier-Stokes)
- Use texture models to synthesize vorticity (images)

Turbulence modeling

Can we characterize coherent structures in turbulent flows?

- Simulate fluid vorticity by PDE models (Navier-Stokes)
- Use texture models to synthesize vorticity (images)

Vorticity

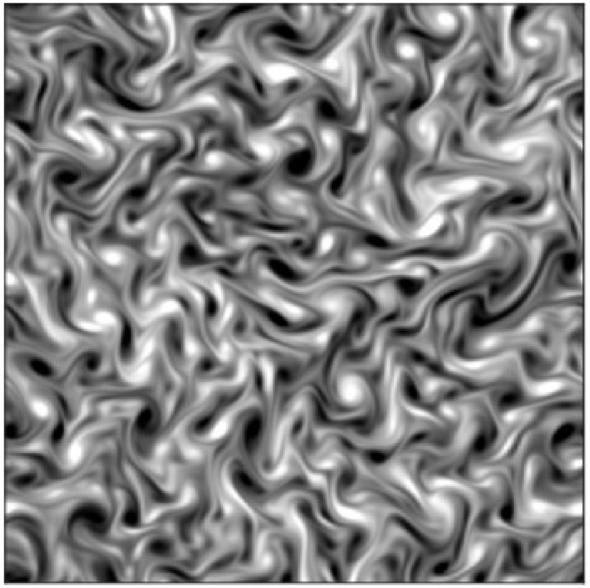

Synthesis

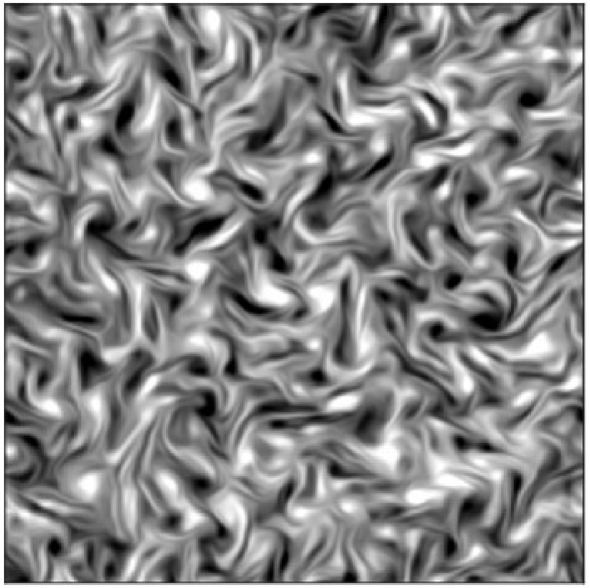

Texture models can capture geometric information in images, why?

Outline

- Multi-scale models for texture synthesis
  - Framework: maximum-entropy models
  - Review of multi-scale approach: wavelet vs. deep learning
  - Main result: *phase harmonic covariance model*
- Texture models and stochastic geometry
  - Point process models and topological data analysis
- Texture models and cosmology
  - From synthesis to denoising problem
- Texture style transfer and relation to AI

Maximum-entropy model

- Maximize the entropy of model $\tilde{p}$ under moment constraints

$$\max_{\tilde{p}} \text{ Entropy}(\tilde{p}) \text{ s.t. } \mathbb{E}_{x\sim p}(\Phi(x)) = \mathbb{E}_{x\sim \tilde{p}}(\Phi(x))$$

Maximum-entropy model

- Maximize the entropy of model $\tilde{p}$ under moment constraints

$$\max_{\tilde{p}} \text{Entropy}(\tilde{p}) \text{ s.t. } \mathbb{E}_{x\sim p}(\Phi(x)) = \mathbb{E}_{x\sim\tilde{p}}(\Phi(x))$$

- Dual solution: Gibbs distribution

$$\tilde{p}(x) = Z^{-1} e^{-\lambda^\intercal \Phi(x)}, \quad \Phi(x) \in \mathbb{R}^{d'}$$

Maximum-entropy model

- Maximize the entropy of model $\tilde{p}$ under moment constraints

$$\max_{\tilde{p}} \text{Entropy}(\tilde{p}) \text{ s.t. } \mathbb{E}_{x\sim p}(\Phi(x)) = \mathbb{E}_{x\sim\tilde{p}}(\Phi(x))$$

- Dual solution: Gibbs distribution

$$\tilde{p}(x) = Z^{-1} e^{-\lambda^\intercal \Phi(x)}, \quad \Phi(x) \in \mathbb{R}^{d'}$$

- Evaluation: compare similarity between samples of p and $\tilde{p}$

Texture: sample from p

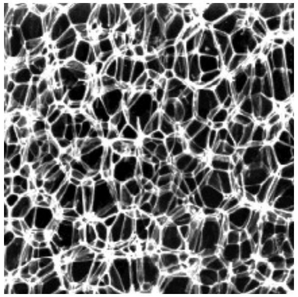

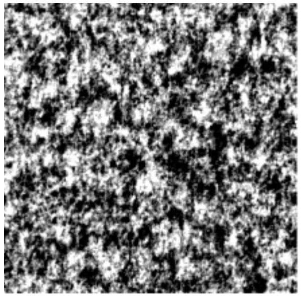

Synthesis: sample from $\tilde{p}$

Maximum-entropy model

- Maximize the entropy of model $\tilde{p}$ under moment constraints

$$\max_{\tilde{p}} \text{Entropy}(\tilde{p}) \text{ s.t. } \mathbb{E}_{x\sim p}(\Phi(x)) = \mathbb{E}_{x\sim\tilde{p}}(\Phi(x))$$

- Dual solution: Gibbs distribution

$$\tilde{p}(x) = Z^{-1} e^{-\lambda^\intercal \Phi(x)}, \quad \Phi(x) \in \mathbb{R}^{d'}$$

- Evaluation: compare similarity between samples of p and $\tilde{p}$

Texture: sample from p

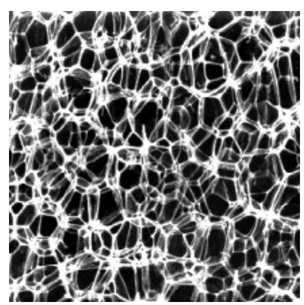

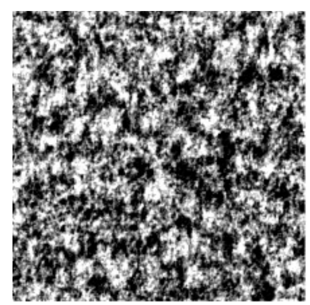

Synthesis: sample from $\tilde{p}$

- Key question: 1. how to specify $\Phi : \mathbb{R}^d \to \mathbb{R}^{d'}$ so that $\tilde{p} \approx p$?
  2. draw samples from $\tilde{p}$ when d is large

Choice of moments

Large amount of Φ do not always produce similar samples

Texture
$d = 6 \times 10^4$

Synthesis
$d' = 6 \times 10^4$

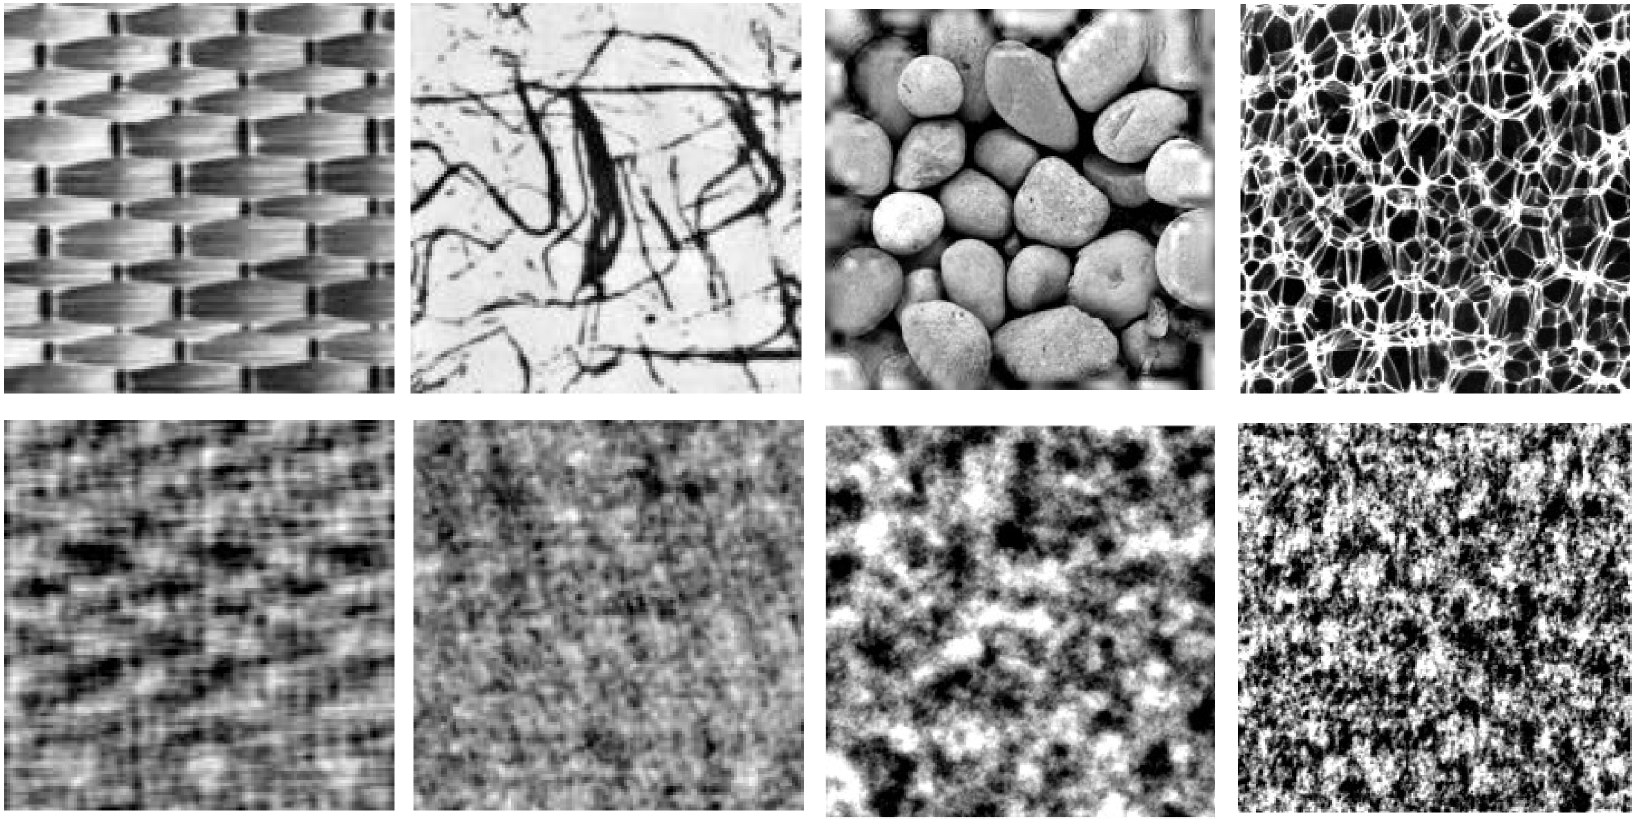

Φ: 2nd-order moments (covariance between pixel values)

Choice of moments

Large amount of Φ do not always produce similar samples

$Texture$ $d = 6 \times 10^4$

$Synthesis$ $d' = 6 \times 10^4$

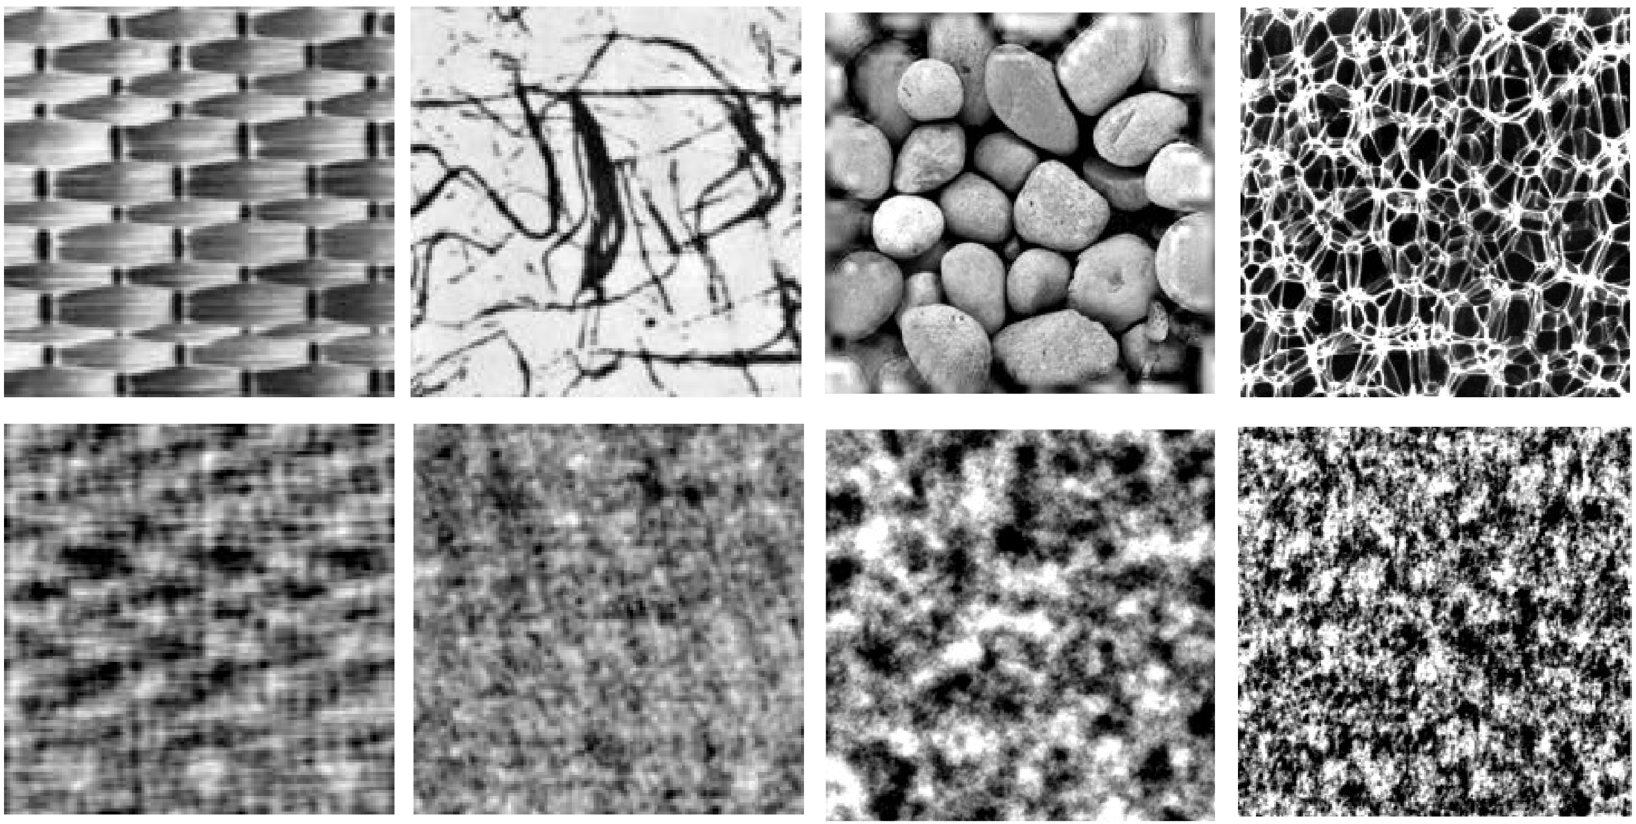

Φ: 2nd-order moments (covariance between pixel values)

$\Rightarrow$ **Problem:** $\tilde{p}$ is a Gaussian distribution, but p is not

Capture non-Gaussian information

- Goal: Specify Φ to capture info. beyond 2nd order moments

2nd order moments $\quad x(u)$ is pixel value at position u

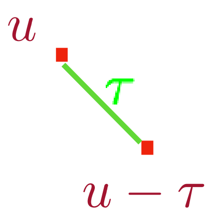

$$\Phi(x) = \{x(u)x(u-\tau)\}_\tau$$

higher order moments (e.g. 3rd order)

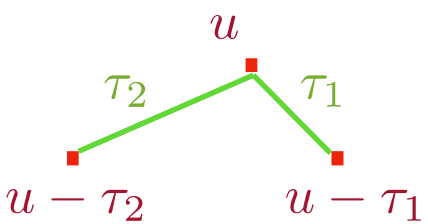

$$\Phi(x) = \{x(u)x(u-\tau_1)x(u-\tau_2)\}_{\tau_1,\tau_2}$$

Problem: estimation variance of $\mathbb{E}_{x\sim p}(\Phi(x))$

Capture non-Gaussian information

- Goal: Specify Φ to capture info. beyond 2nd order moments

2nd order moments $\quad x(u)$ is pixel value at position u

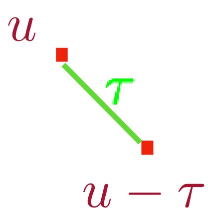

$$\Phi(x) = \{x(u)x(u-\tau)\}_\tau$$

higher order moments (e.g. 3rd order)

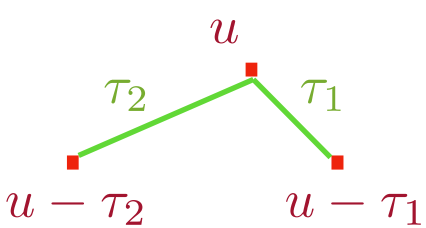

$$\Phi(x) = \{x(u)x(u-\tau_1)x(u-\tau_2)\}_{\tau_1,\tau_2}$$

Problem: estimation variance of $\mathbb{E}_{x\sim p}(\Phi(x))$

- State-of-the-art: Wavelet-based vs. Deep-learning based Φ

Idea : capture 1st and 2nd order moments in a transform domain

$\Rightarrow$ non-Gaussian info. without too large variance

Wavelet and deep learning, 1989

S. Mallat. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation

Y. LeCun et al. Backpropagation applied to handwritten zip code recognition

Wavelet and deep learning, 1989

S. Mallat. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation
Y. LeCun et al. Backpropagation applied to handwritten zip code recognition

- Decompose images into multi-scale using self-similar filters
- Image recognition using a cascade of learnt filters

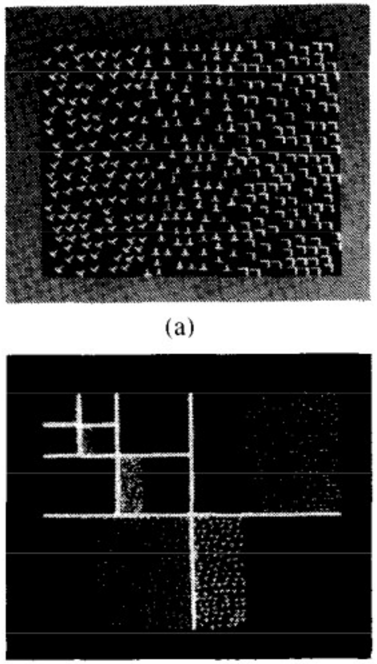

Fig. 17. (a) J. Beck textures: only the left texture is preattentively discriminable by a human observer. (b) These images show the absolute value of the wavelet coefficients of image (a), computed on three resolution levels. The left texture can be discriminated with a first-order statistical analysis of the detail signals amplitude. The two other textures can not be discriminated with such a technic.

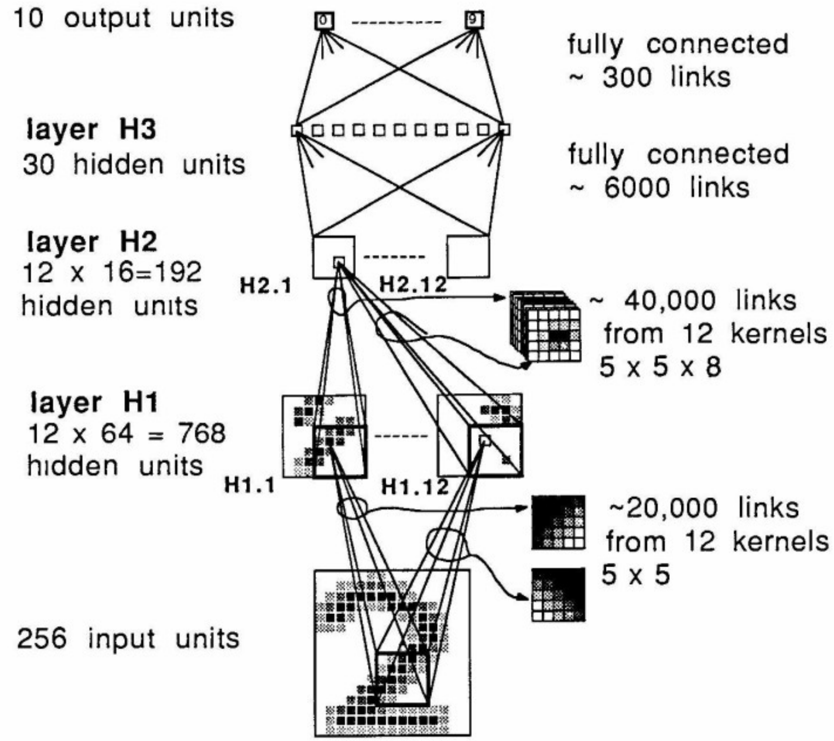

Figure 3 Log mean squared error (MSE) (top) and raw error rate (bottom) versus number of training passes

Two cultures in data science

L. Breiman, 2001

- There are two cultures in the use of statistical modeling to reach conclusions from data.
  - One assumes that the data are generated by a given stochastic data model. ⇒ Simple world
  - The other uses algorithmic models and treats the data mechanism as unknown. ⇒ Complex world

Two cultures in data science

L. Breiman, 2001

- There are two cultures in the use of statistical modeling to reach conclusions from data.
  - One assumes that the data are generated by a given stochastic data model. ⇒ Simple world
  - The other uses algorithmic models and treats the data mechanism as unknown. ⇒ Complex world
- If our goal as a field is to use data to solve problems, then we need to move away from exclusive dependence on data models and adopt a more diverse set of tools.

Statistics, Signal processing ⇒ Machine learning (ML)

Two cultures in data science

L. Breiman, 2001

- There are two cultures in the use of statistical modeling to reach conclusions from data.
  - One assumes that the data are generated by a given stochastic data model. $\Rightarrow$ Simple world
  - The other uses algorithmic models and treats the data mechanism as unknown. $\Rightarrow$ Complex world
- If our goal as a field is to use data to solve problems, then we need to move away from exclusive dependence on data models and adopt a more diverse set of tools.

Statistics, Signal processing $\Rightarrow$ Machine learning (ML)

Challenge: build statistical models of complex data

Texture synthesis problem

Can we model parts of complex data?

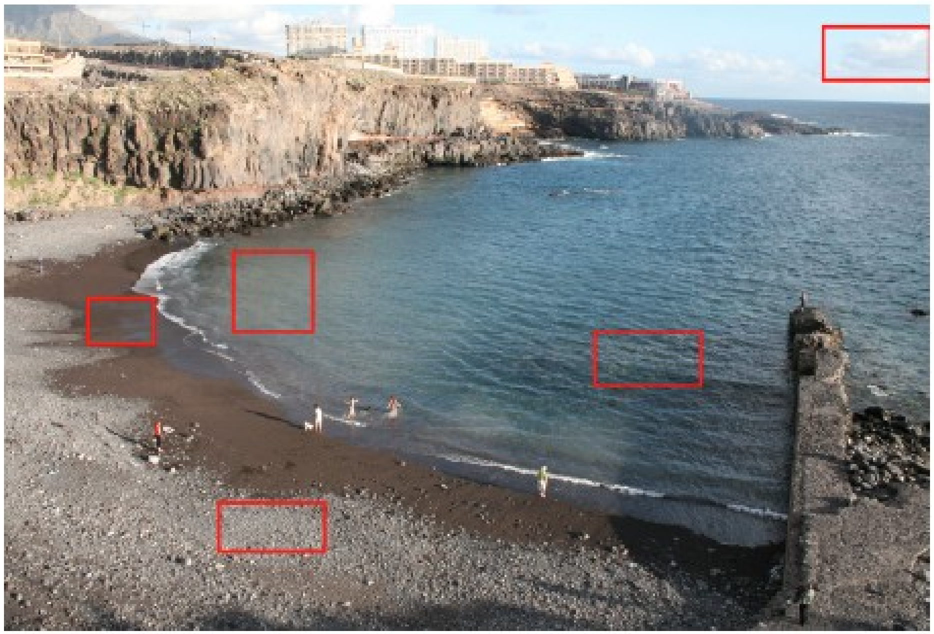

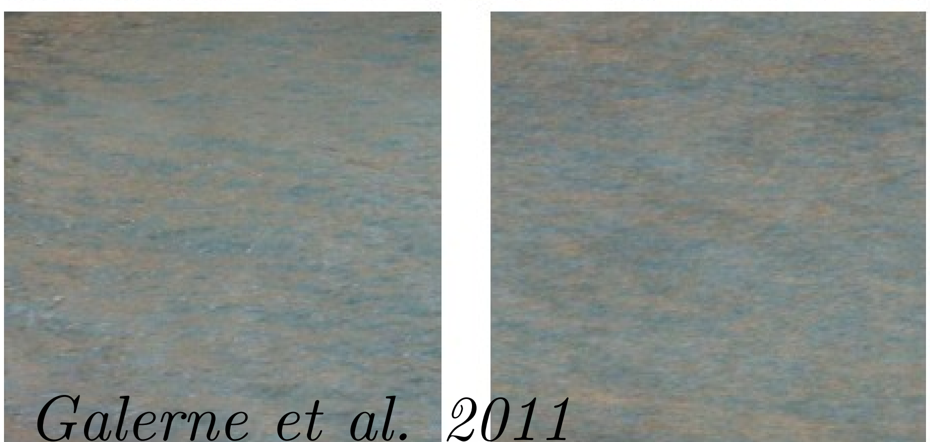

Galerne et al. 2011

Texture *Synthesis*

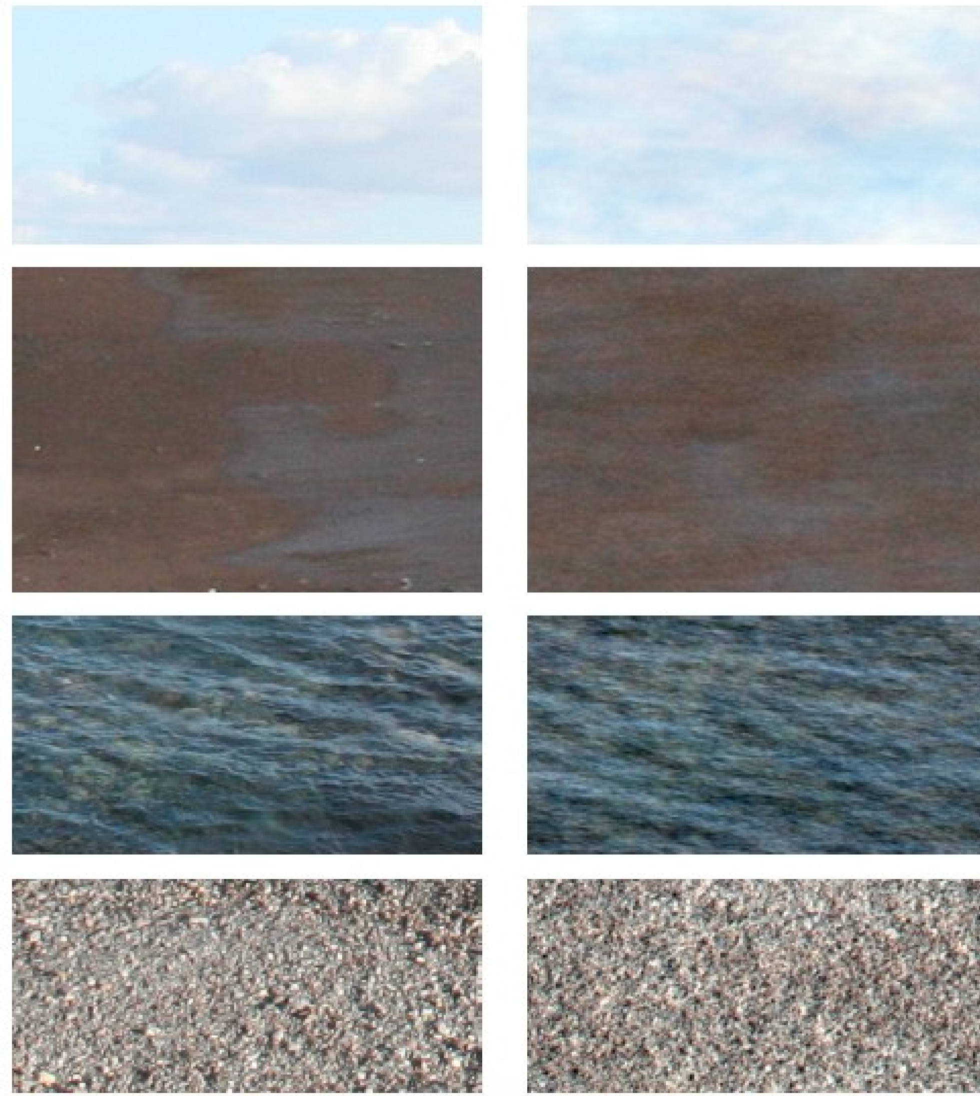

Wavelet-based texture model

PS: Portilla and Simoncelli (2000)

- **Idea**: Take 1st and 2nd order moments in a wavelet domain
- Wavelet transform: seperate x into multiple scales

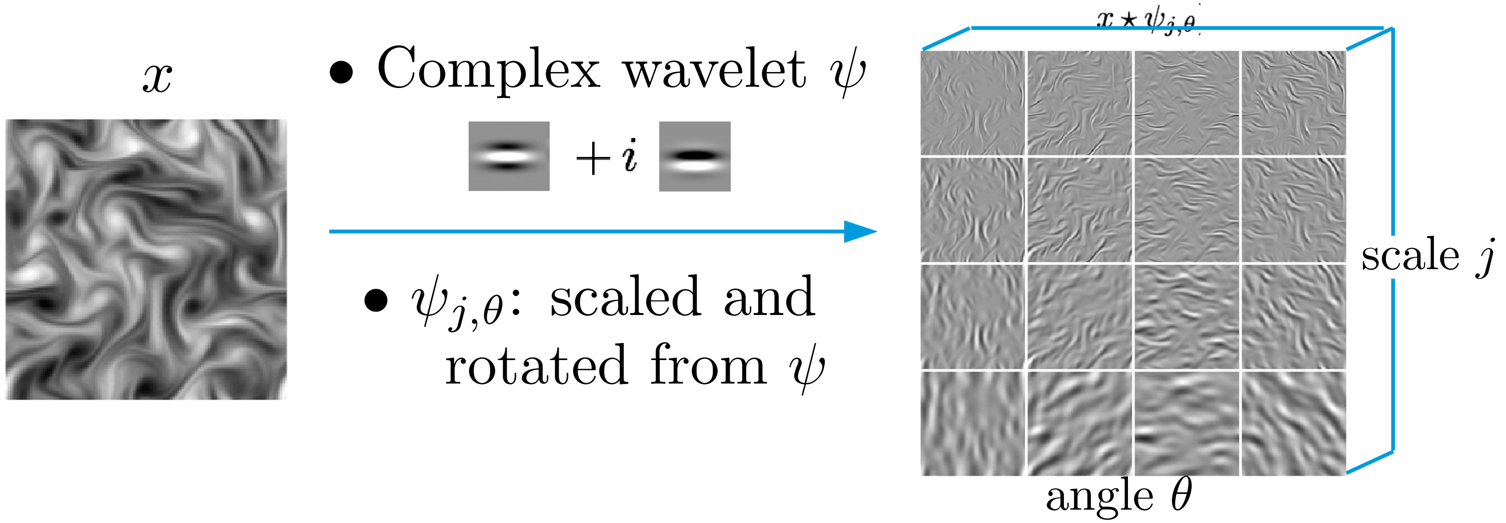

Wavelet-based texture model

PS: Portilla and Simoncelli (2000)

- **Idea**: Take 1st and 2nd order moments in a wavelet domain
- Wavelet transform: seperate x into multiple scales

x

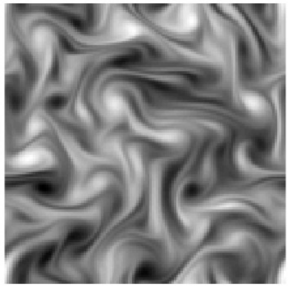

- Complex wavelet ψ

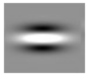 $+\,i$ 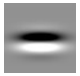

- $\psi_{j,\theta}$: scaled and rotated from ψ

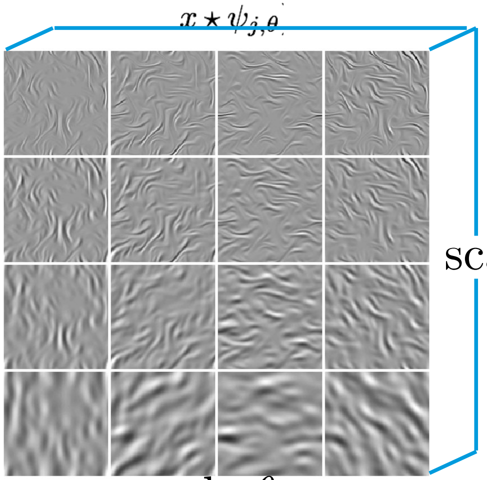

Φ: capture multi-scale angular interactions
respect invariance of p (translation, rotational, scaling)

Wavelet-based texture model

PS: Portilla and Simoncelli (2000)

- **Idea**: Take 1st and 2nd order moments in a wavelet domain
- Wavelet transform: seperate x into multiple scales

x

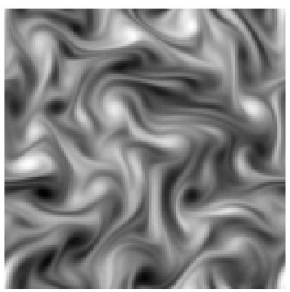

- Complex wavelet ψ

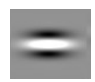 $+i$ 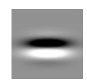

- $\psi_{j,\theta}$: scaled and rotated from ψ

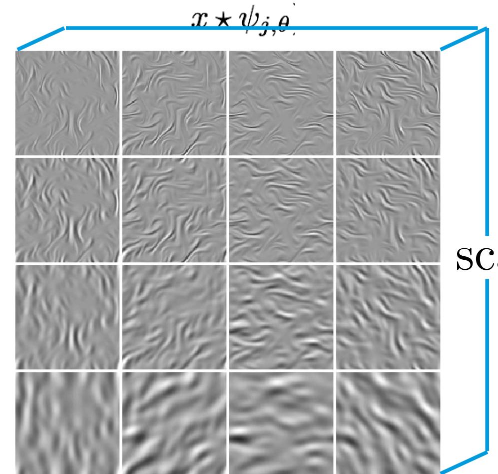

Φ: capture multi-scale angular interactions
respect invariance of p (translation, rotational, scaling)

- Connection with convolutional neural network (CNN)

Image x → Conv2d(x, Real part of $\{\psi_{j,\theta}\}_{(j,\theta)}$) → Non-linearity

Image x → Conv2d(x, Imag part of $\{\psi_{j,\theta}\}_{(j,\theta)}$) → Non-linearity

Deep learning based texture model

VGG: Gatys et al. (2015)

- **Idea**: Take 1st and 2nd order moments in CNN layers

CNN: Image x → Conv2d(x, Filters$_1$) → Non-linearity$_1$ → Conv2d(x, Filters$_2$) → Non-linearity$_2$ → Other layers ...

Deep learning based texture model

VGG: Gatys et al. (2015)

- **Idea**: Take 1st and 2nd order moments in CNN layers

CNN: Image x → Conv2d(x, Filters$_1$) → Non-linearity$_1$

→ Conv2d(x, Filters$_2$) → Non-linearity$_2$ → Other layers ...

- Synthesis comparison

VGG *PS*

Vorticity

Texture

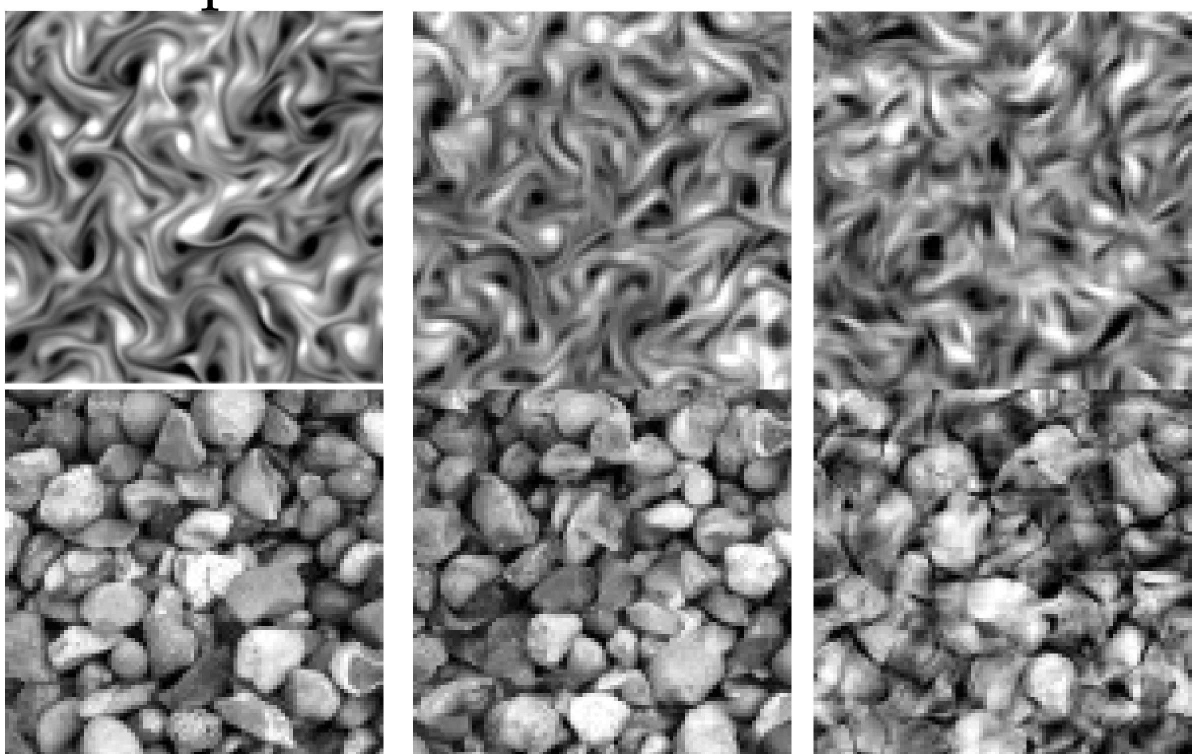

$d = 6 \times 10^4$

$d'_{\text{vgg}} = 18 \times 10^4$

$d'_{\text{ps}} = 0.3 \times 10^4$

PS images: less coherent

Understand deep learning models

RF: Ustyuzhaninov et al., 2017

Deep learning (VGG) performs better by using a large d'

Question: What it takes to generate natural textures?

Understand deep learning models

RF: Ustyuzhaninov et al., 2017

Deep learning (VGG) performs better by using a large d'

Question: What it takes to generate natural textures?

- To simplify VGG, RF model is proposed using 1-layer CNN

$$\Phi(x) \approx \{\mathrm{Cov}(\rho(x \star \psi_\lambda(u)), \rho(x \star \psi_{\lambda'}(u-\tau))\}_{\lambda,\lambda',\tau}$$

$\rho(a) = \max(a, 0)$: rectifier non-linearity $\{\psi_\lambda\}_\lambda$: multi-scale random filters

Understand deep learning models

RF: Ustyuzhaninov et al., 2017

Deep learning (VGG) performs better by using a large d'

Question: What it takes to generate natural textures?

- To simplify VGG, RF model is proposed using 1-layer CNN

$$\Phi(x) \approx \{\mathrm{Cov}(\rho(x \star \psi_\lambda(u)), \rho(x \star \psi_{\lambda'}(u-\tau))\}_{\lambda,\lambda',\tau}$$

$\rho(a) = \max(a, 0)$: rectifier non-linearity $\{\psi_\lambda\}_\lambda$: multi-scale random filters

Texture $(d = 65k)$

RF$(d' = 525k)$

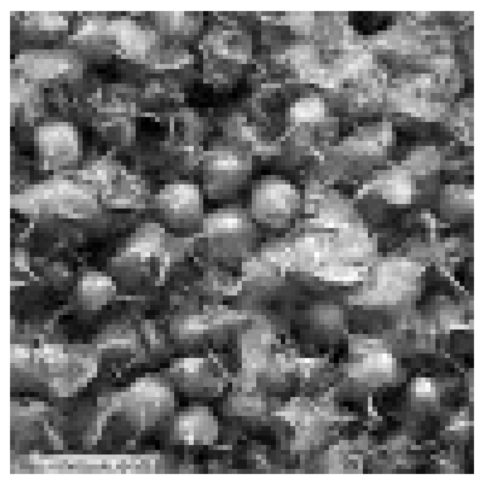

VGG$(d' \approx 177k)$

RF model synthesis is similar to VGG by using large d'

Wavelet vs. deep learning model

Question: What is in common between PS and RF model?

Can we bridge the gap between PS and VGG/RF model?

Wavelet vs. deep learning model

Question: What is in common between PS and RF model?
Can we bridge the gap between PS and VGG/RF model?

- Important differences between RF and PS models
  - ReLU vs. modulus non-linearity
  - Choice of convolutional filters

    PS model uses complex wavelet $\psi_\lambda(u), \lambda = (j, \theta)$

    RF model uses real and random filters ($\psi_\lambda(u) \sim$ i.i.d. Gaussian)

Mallat, Zhang, Gaspar 2020
Zhang, Mallat 2021
Brochard, Zhang, Mallat 2022

Rectifier wavelet covariance

Mallat, Zhang, Gaspar 2020

- Propose a generalized rectifier with phase $\alpha \in [0, 2\pi]$

$$\rho_\alpha(z) = \rho(\mathrm{Real}(e^{i\alpha} z)), \text{z complex number}$$

rectifier: $\rho(a) = \max(a, 0)$

Rectifier wavelet covariance

Mallat, Zhang, Gaspar 2020

- Propose a generalized rectifier with phase $\alpha \in [0, 2\pi]$

$$\rho_\alpha(z) = \rho(\mathrm{Real}(e^{i\alpha} z)), \text{z complex number}$$

rectifier: $\rho(a) = \max(a, 0)$ *Brochard, Zhang, Mallat 2022*

- Wavelet rectifier covariance with spatial shift τ

$$\Phi(x) \approx \{\mathrm{Cov}(\rho_\alpha(x \star \psi_\lambda(u)), \rho_{\alpha'}(x_\tau \star \psi_{\lambda'}(u)))\}_{\alpha,\alpha',\lambda,\lambda',\tau}$$

ρ_α: generalized rectifier $\{\psi_\lambda\}$: complex wavelet filters

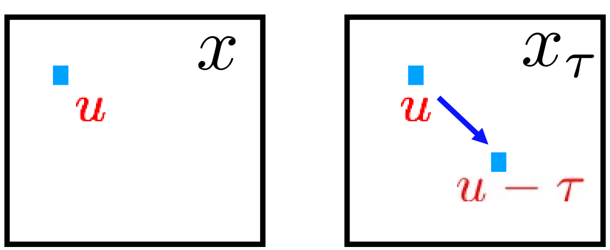

Capture **correlations** between

$\rho_\alpha(x \star \psi_\lambda(u))$ and $\rho_{\alpha'}(x \star \psi_{\lambda'}(u - \tau))$

Relation with phase harmonics

- Fourier transform of ρ_α along α: **phase harmonics**

$$\rho_\alpha(z) = \rho(\mathrm{Real}(e^{i\alpha} z)) \quad \overset{\text{dual}}{\Longleftrightarrow} \quad [z]^k = |z| e^{ik\,\mathrm{phase}(z)}$$

Relation with phase harmonics

- Fourier transform of ρ_α along α: **phase harmonics**

$$\rho_\alpha(z) = \rho(\mathrm{Real}(e^{i\alpha} z)) \quad \overset{\text{dual}}{\longleftrightarrow} \quad [z]^k = |z| e^{ik\,\mathrm{phase}(z)}$$

- Generalize the classical Portilla & Simoncelli (2000) model

$$\mathrm{Cov}([x \star \psi_\lambda(u)]^k, [x_\tau \star \psi_{\lambda'}(u)]^{k'})$$

- **PS moments**: $k = k' = 0, 1$ and $k = 1, k' = 2$

Capture information beyond 2nd order statistics ($k = k' = 1$)

Relation with phase harmonics

- Fourier transform of ρ_α along α: **phase harmonics**

$$\rho_\alpha(z) = \rho(\mathrm{Real}(e^{i\alpha} z)) \quad \overset{\text{dual}}{\Longleftrightarrow} \quad [z]^k = |z| e^{ik\,\mathrm{phase}(z)}$$

- Generalize the classical Portilla & Simoncelli (2000) model

$$\mathrm{Cov}([x \star \psi_\lambda(u)]^k, [x_\tau \star \psi_{\lambda'}(u)]^{k'})$$

- **PS moments**: $k = k' = 0, 1$ and $k = 1, k' = 2$

Capture information beyond 2nd order statistics ($k = k' = 1$)

- **Phase harmonic covariance model** (rectifier form)

$$\mathrm{Cov}(\rho_\alpha(x \star \psi_\lambda(u)), \rho_{\alpha'}(x_\tau \star \psi_{\lambda'}(u)))$$

$\Rightarrow$ Unify PS and RF model (using multi-scale filters)

Memorization issues: lack diversity

- Number of moments d': quality/diversity trade-off

Texture(65k) $\Gamma_S(3.5k)$ $\Gamma_M(35k)$ $\Gamma_L(142k)$

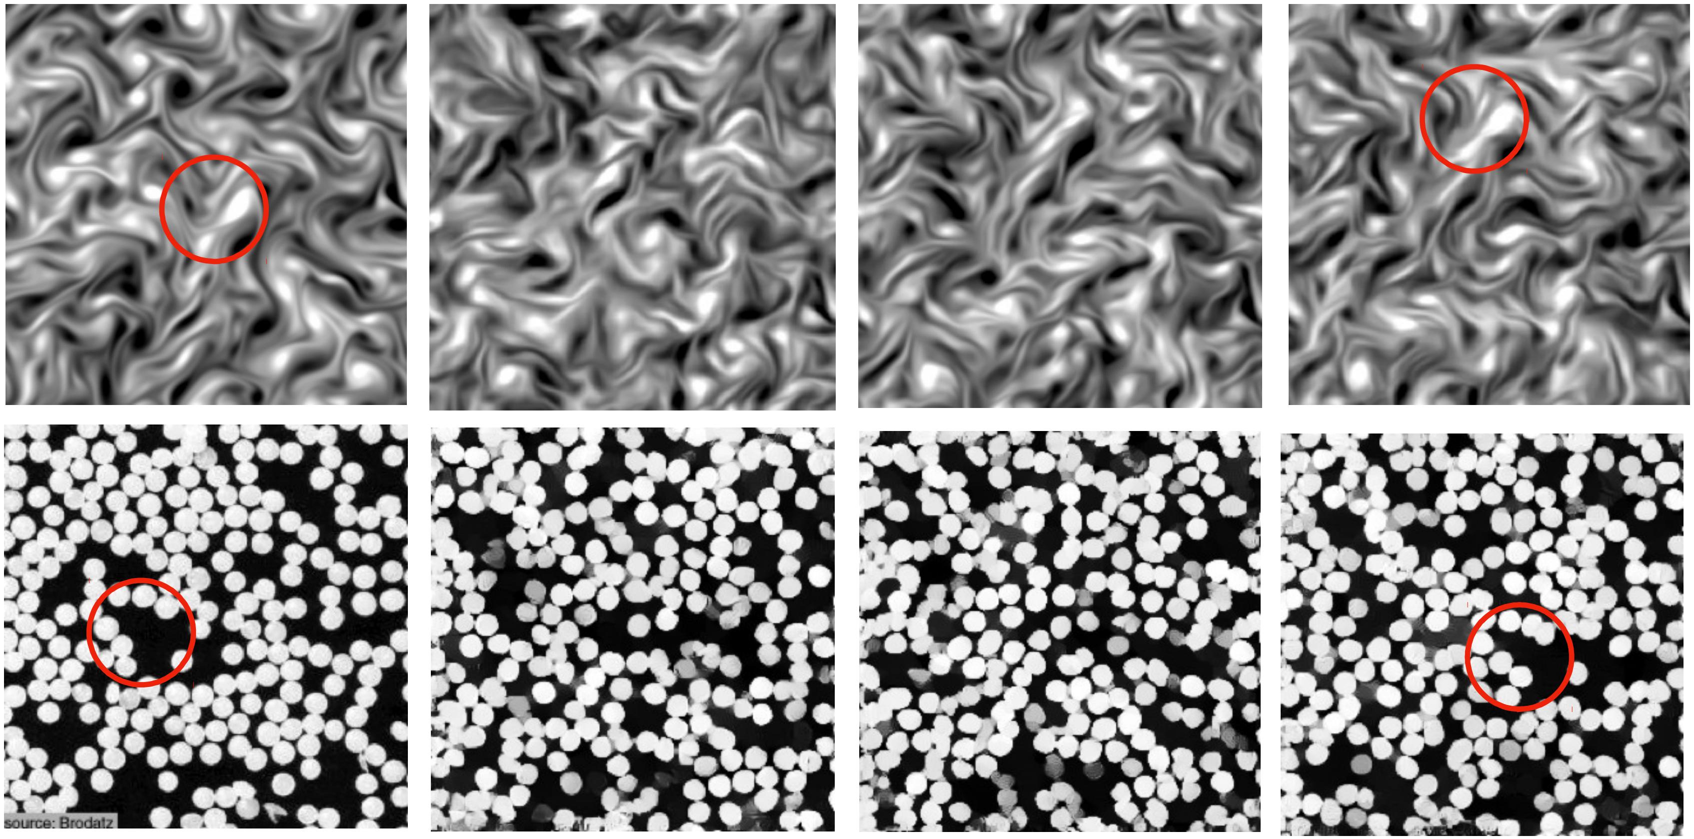

$$\mathrm{Cov}(\rho_\alpha(x \star \psi_\lambda(u)), \rho_{\alpha'}(x_\tau \star \psi_{\lambda'}(u))) \quad (\lambda, \lambda', \alpha, \alpha', \tau) \in \Gamma$$

Increasing range: $\Gamma_S \subset \Gamma_M \subset \Gamma_L$

Choose the covariance sets

How to choose $\Gamma_S, \Gamma_M, \Gamma_L$ of $(\lambda, \lambda', \alpha, \alpha', \tau)$?

- For the wavelets $\{\psi_\lambda\}$, $\lambda = (j, \theta)$, $\lambda' = (j', \theta')$
  - Partial or Full scale interactions: $|j - j'| \leq \Delta$, $\Delta \in \{1, J\}$
  - Full angular interactions: $\forall(\theta, \theta')$

Choose the covariance sets

How to choose $\Gamma_S, \Gamma_M, \Gamma_L$ of $(\lambda, \lambda', \alpha, \alpha', \tau)$?

- For the wavelets $\{\psi_\lambda\}$, $\lambda = (j, \theta)$, $\lambda' = (j', \theta')$
  - Partial or Full scale interactions: $|j - j'| \leq \Delta$, $\Delta \in \{1, J\}$
  - Full angular interactions: $\forall(\theta, \theta')$
- For the phase α and α'
  - Phase interactions: $\alpha \in A_4, \alpha' \in A_1$ or $\alpha' \in A_4$
  - Four phases: $\mathrm{A}_4 = \{0, \pi/4, \pi/2, 3\pi/4\}$
  - One phase: $\mathrm{A}_1 = \{0\}$ $\Gamma_M(35k) : \Delta = J, \alpha' \in A_1$

Choose the covariance sets

How to choose $\Gamma_S, \Gamma_M, \Gamma_L$ of $(\lambda, \lambda', \alpha, \alpha', \tau)$?

- For the wavelets $\{\psi_\lambda\}$, $\lambda = (j, \theta)$, $\lambda' = (j', \theta')$
  - Partial or Full scale interactions: $|j - j'| \leq \Delta$, $\Delta \in \{1, J\}$
  - Full angular interactions: $\forall(\theta, \theta')$
- For the phase α and α'
  - Phase interactions: $\alpha \in A_4, \alpha' \in A_1$ or $\alpha' \in A_4$
  - Four phases: $\mathrm{A}_4 = \{0, \pi/4, \pi/2, 3\pi/4\}$
  - One phase: $\mathrm{A}_1 = \{0\}$ $\quad$ $\Gamma_M(35k) : \Delta = J, \alpha' \in A_1$
- For the spatial shift τ : subsampled grid on $\mathbb{R}^2$

L=2

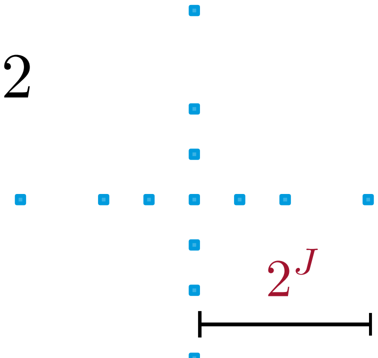

L=4

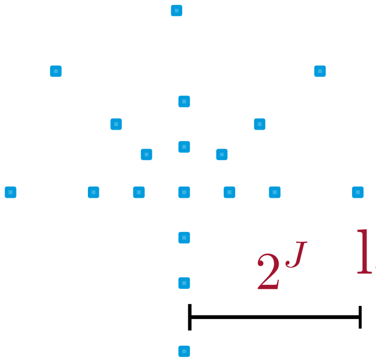

total τ: $O(JL)$

large-scale interactions

Proposed model

Texture (65k) PS(3k) RF(525k) Ours(35k) VGG(177k)

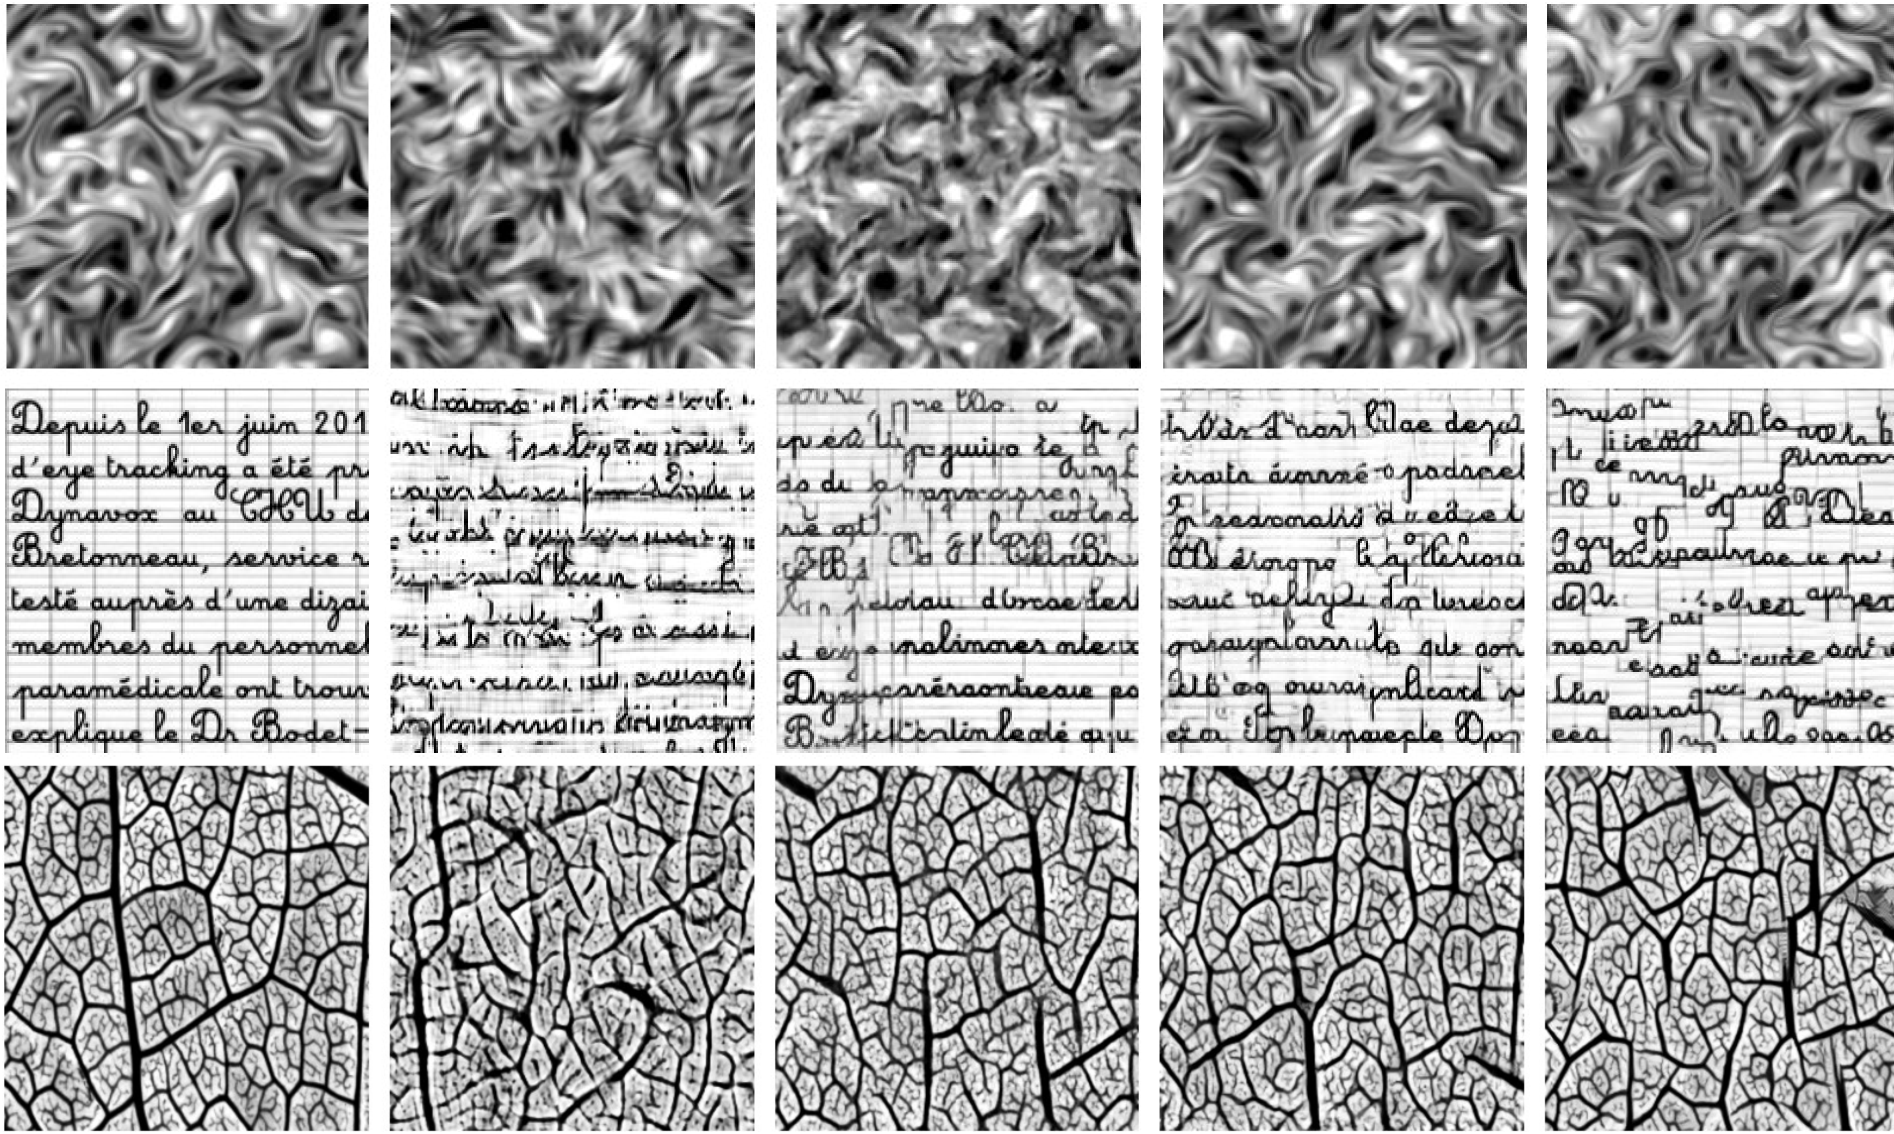

Our model bridges the gap between PS and VGG

Extension to color textures

- Model color texture $x = \{x^c\}_{c=1,2,3}$ with RGB channels

Extension to color textures

- Model color texture $x = \{x^c\}_{c=1,2,3}$ with RGB channels
- Capture color coherence using cross-channel statistics

$$\mathrm{Cov}(\rho_\alpha(x_c \star \psi_\lambda(u)), \rho_{\alpha'}(x_{c',\tau} \star \psi_{\lambda'}(u))) \quad (c, c') \in \{1,2,3\}^2$$

Extension to color textures

- Model color texture $x = \{x^c\}_{c=1,2,3}$ with RGB channels
- Capture color coherence using cross-channel statistics

$$\mathrm{Cov}(\rho_\alpha(x_c \star \psi_\lambda(u)), \rho_{\alpha'}(x_{c',\tau} \star \psi_{\lambda'}(u))) \quad (c, c') \in \{1,2,3\}^2$$

Texture($195k$) PS(17k) RF(525k) Ours(320k) VGG(177k)

Extension to color textures

- Model color texture $x = \{x^c\}_{c=1,2,3}$ with RGB channels
- Capture color coherence using cross-channel statistics

$$\mathrm{Cov}(\rho_\alpha(x_c \star \psi_\lambda(u)), \rho_{\alpha'}(x_{c',\tau} \star \psi_{\lambda'}(u))) \quad (c, c') \in \{1, 2, 3\}^2$$

Texture($195k$) PS(17k) RF(525k) Ours(320k) VGG(177k)

Can we further reduce the number of moments?

Reduced color model

- Capture spatial statistics without color interactions
  - If $c \neq c'$, choose only $\tau = 0$

Original Ours(320k) Reduced(113k)

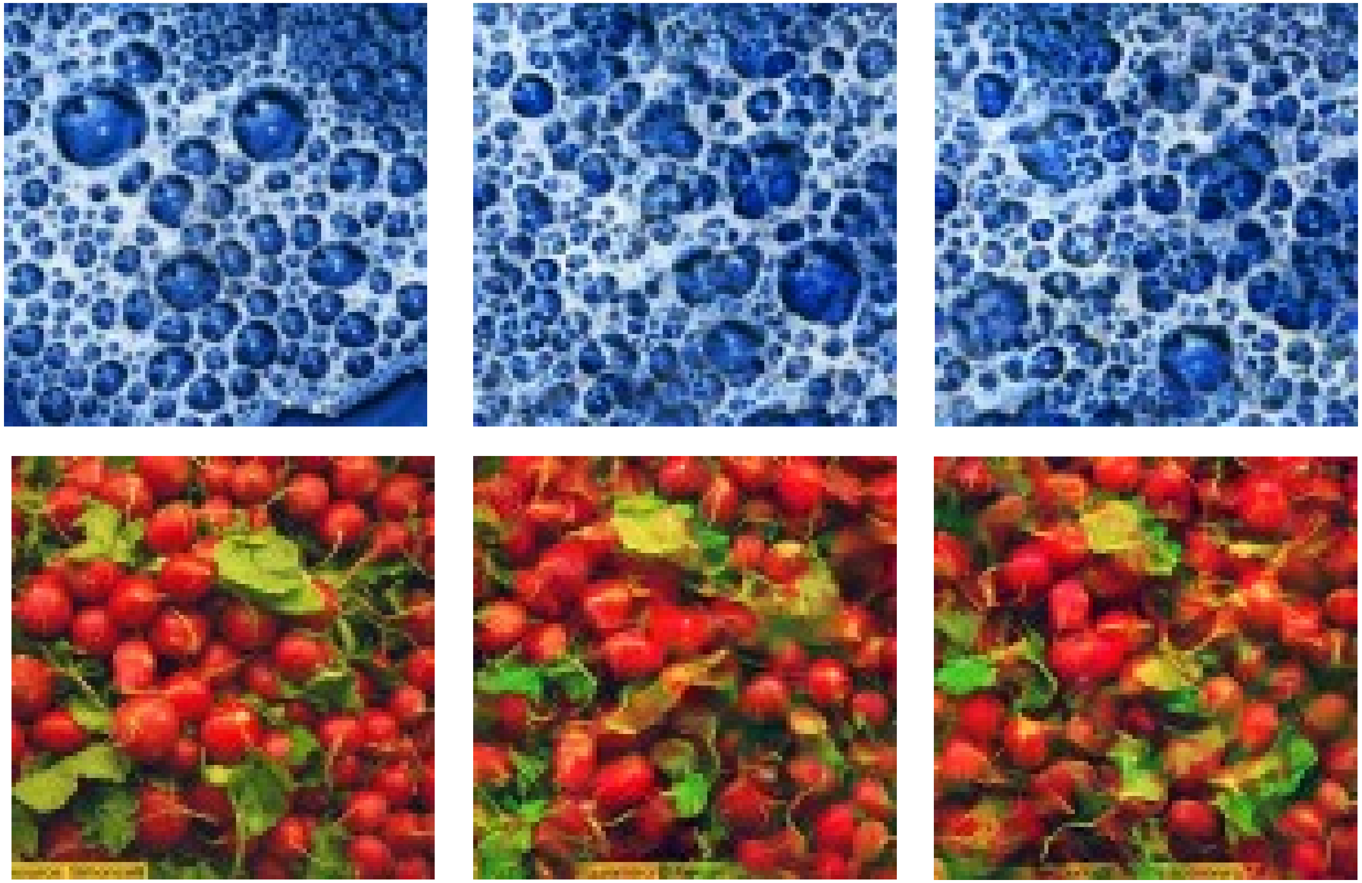

Reduced color model

- Capture spatial statistics without color interactions
  - If $c \neq c'$, choose only $\tau = 0$

Original Ours(320k) Reduced(113k)

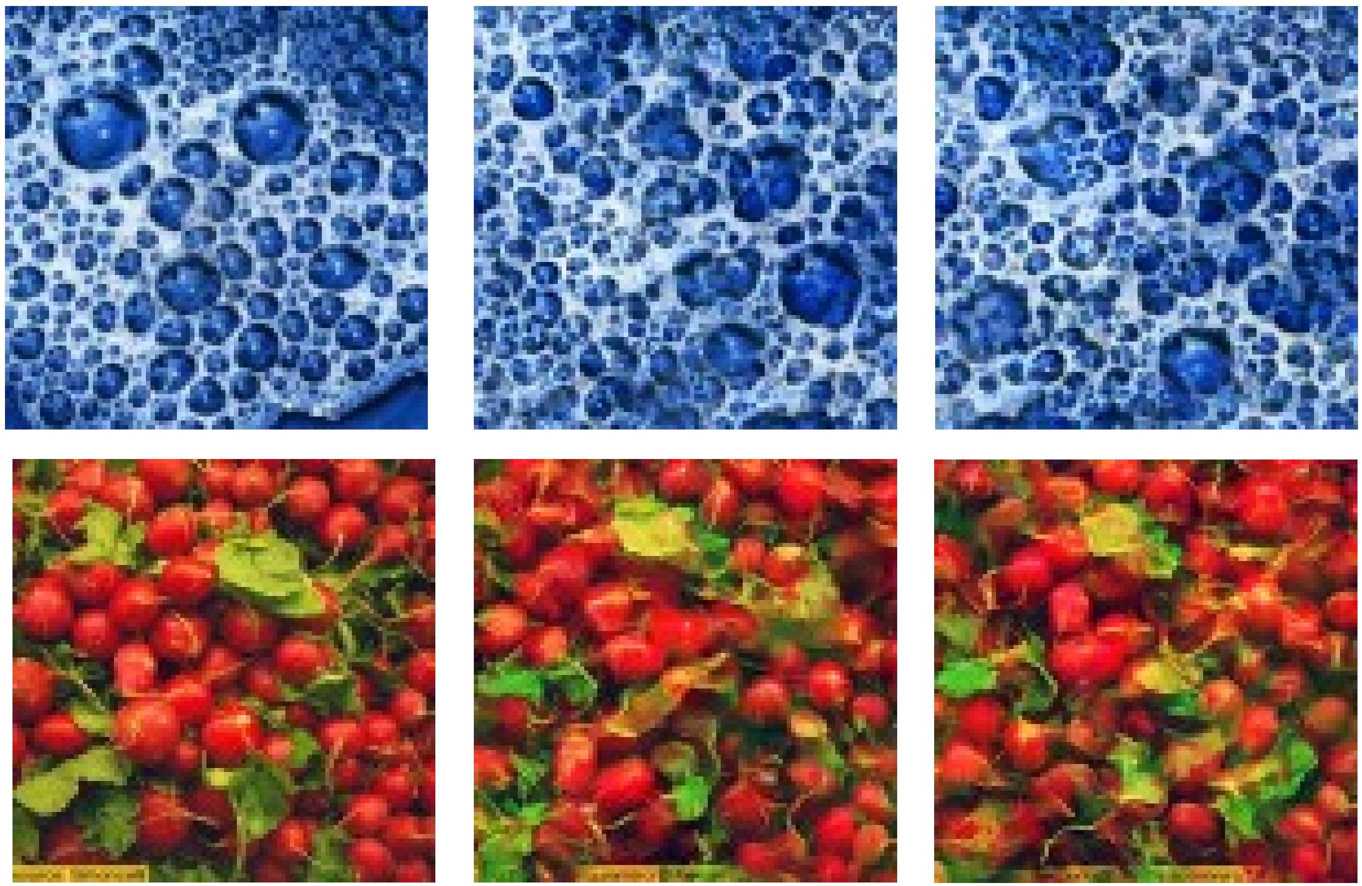

Open: understand the reduced covariance set

From texture to point process

- Point process: random measure $\mu = \sum_{1 \leq i \leq I} \delta_{x_i},\ I \in \bar{\mathbb{N}}$
- Samples of point processes of various geometries

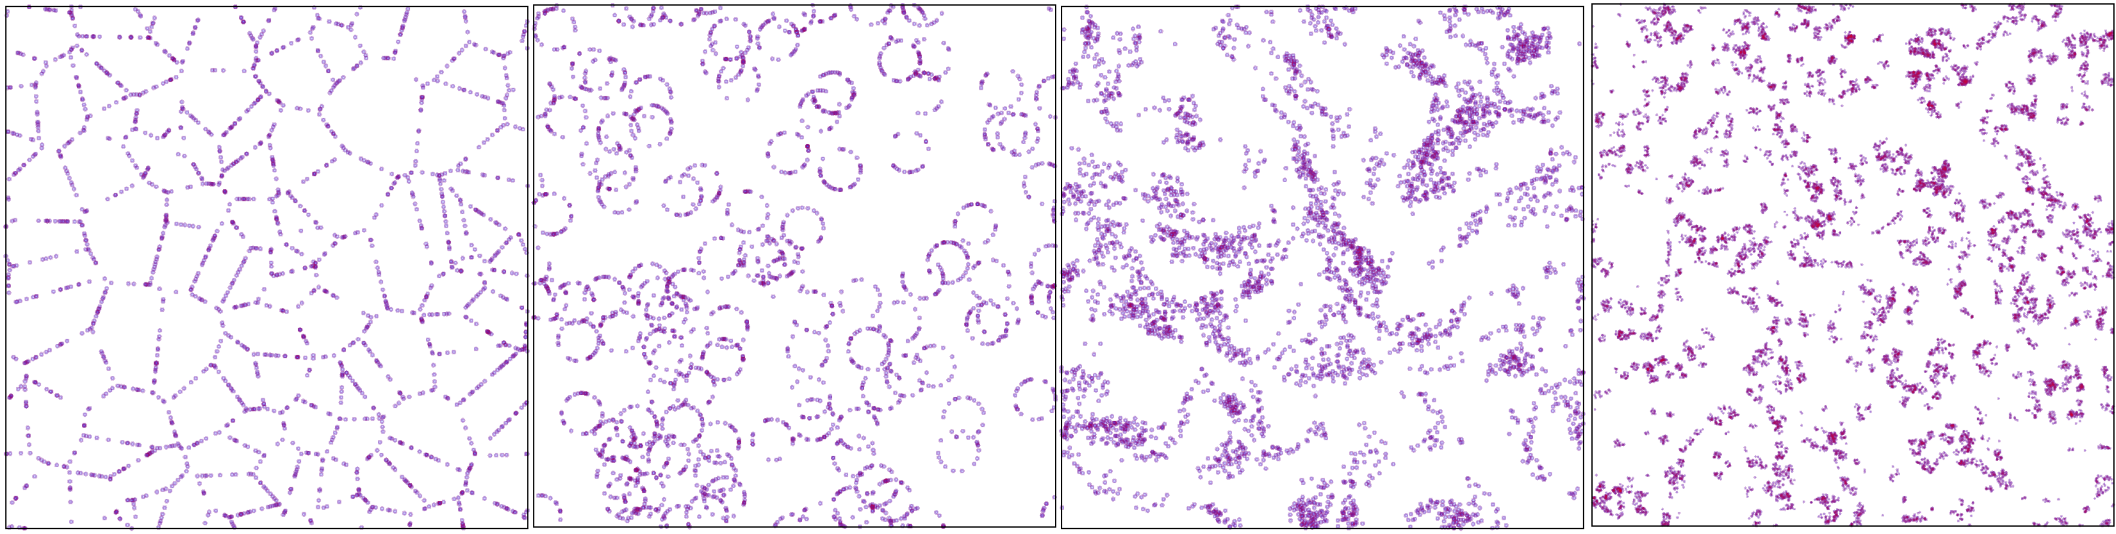

number of points I: 1000-13000

From texture to point process

- Point process: random measure $\mu = \sum_{1 \leq i \leq I} \delta_{x_i}, \; I \in \bar{\mathbb{N}}$
- Samples of point processes of various geometries

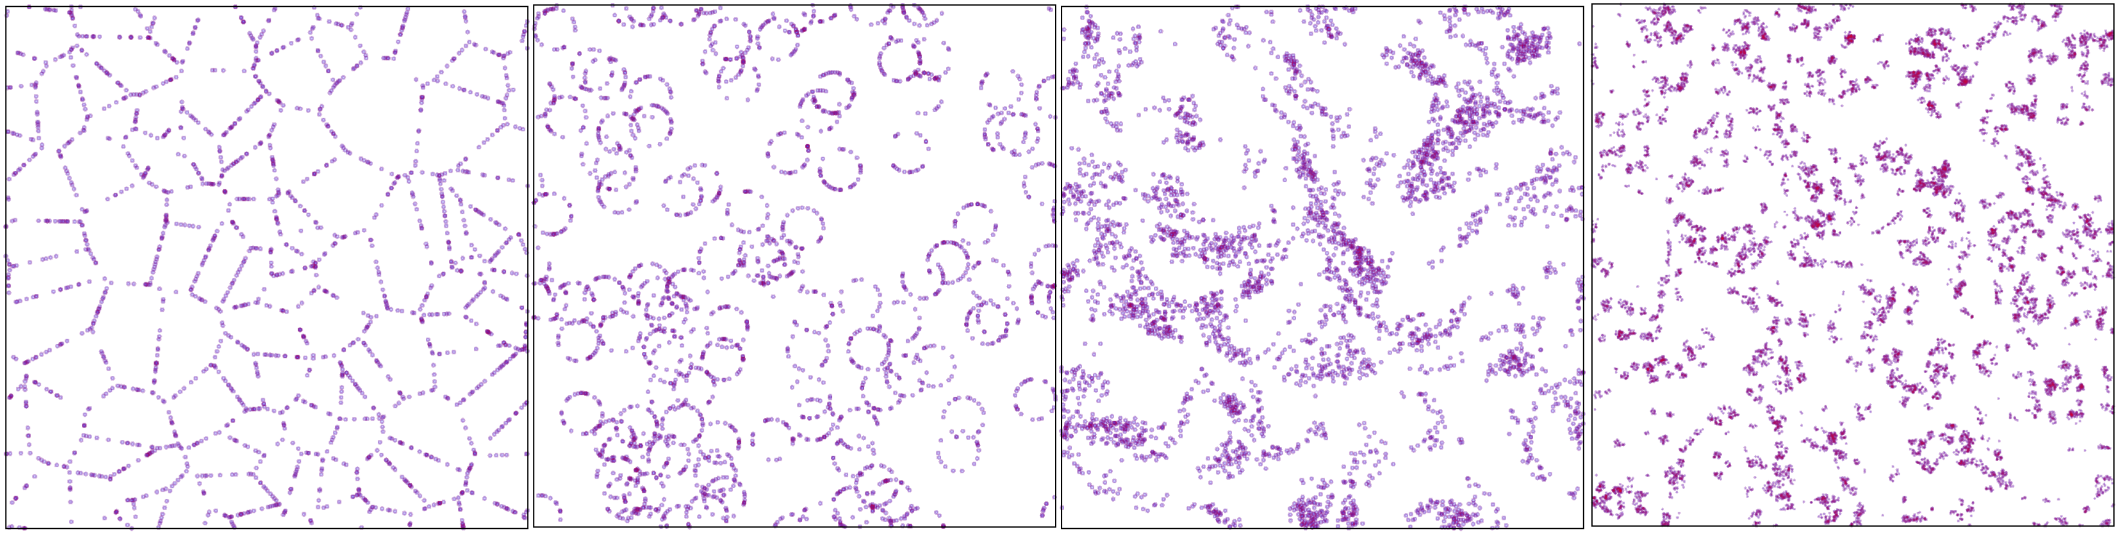

number of points I: 1000-13000

- With existing approaches in stochastic geometry, it is difficult to model point process with complex geometries formed by a large number of particles

Particle gradient-descent model

With A. Brochard, B. Blaszczyszyn, S. Mallat

- Micro-canonical synthesis method for textures
  - Given one observation x, synthesize $\tilde{x} \in \mathbb{R}^d$ such that

$$\tilde{x} \sim \text{Uniform}(\{\tilde{x} : \|\Phi(x) - \Phi(\tilde{x})\| < \epsilon\})$$

key (ergodic) assumption: $\Phi(x) \approx \mathbb{E}_{x \sim p} \Phi(x)$

Particle gradient-descent model

With A. Brochard, B. Blaszczyszyn, S. Mallat

- Micro-canonical synthesis method for textures
  - Given one observation x, synthesize $\tilde{x} \in \mathbb{R}^d$ such that

$$\tilde{x} \sim \text{Uniform}(\{\tilde{x} : \|\Phi(x) - \Phi(\tilde{x})\| < \epsilon\})$$

key (ergodic) assumption: $\Phi(x) \approx \mathbb{E}_{x \sim p} \Phi(x)$

- Idea: given μ, syntheize particles $\{\tilde{x}_i\}$ in $\tilde{\mu} = \sum_i \delta_{\tilde{x}_i}$ by $\min_{\tilde{\mu}} \|K(\mu) - K(\tilde{\mu})\|^2$ using gradient descent

Particle gradient-descent model

With A. Brochard, B. Blaszczyszyn, S. Mallat

- Micro-canonical synthesis method for textures
  - Given one observation x, synthesize $\tilde{x} \in \mathbb{R}^d$ such that

$$\tilde{x} \sim \text{Uniform}(\{\tilde{x} : \|\Phi(x) - \Phi(\tilde{x})\| < \epsilon\})$$

key (ergodic) assumption: $\Phi(x) \approx \mathbb{E}_{x \sim p} \Phi(x)$

- Idea: given μ, syntheize particles $\{\tilde{x}_i\}$ in $\tilde{\mu} = \sum_i \delta_{\tilde{x}_i}$ by

  $\min_{\tilde{\mu}} \|K(\mu) - K(\tilde{\mu})\|^2$ using gradient descent

- Descriptor $K : [-s, s]^{2 \times I} \to \mathbb{R}^{d'}$ captures the geometry in μ
  - Capture geometry in the point process: $K(\mu) \approx \mathbb{E}_\mu(K(\mu))$
  - Convert μ to an image μ_σ^N to use Φ

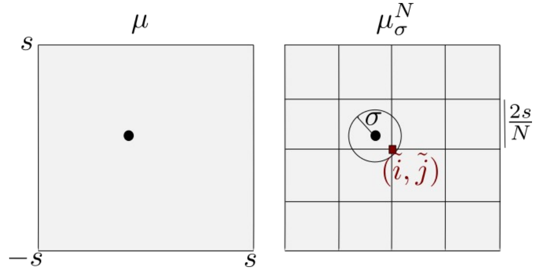

Wavelet phase harmonic descriptors

• Capture multi-scale interactions between the particles, while controlling explicitly the number of moments by the scales of the structures to model.

$$K(\mu) = \left(\frac{1}{4s^2} \int_{[-s,s]^2} \mu_{\lambda,k}(x)\, \mu_{\lambda',k'}(x-\tau')^* dx \right)_{(\lambda,k,\lambda',k',\tau')}$$

Wavelet phase harmonic descriptors

• Capture multi-scale interactions between the particles, while controlling explicitly the number of moments by the scales of the structures to model.

$$K(\mu) = \left(\frac{1}{4s^2} \int_{[-s,s]^2} \mu_{\lambda,k}(x)\, \mu_{\lambda',k'}(x-\tau')^* dx \right)_{(\lambda,k,\lambda',k',\tau')}$$

- For $x \in [-s,s]^2$, $\mu_{\lambda,k}(x) := [\mu \star \psi_\lambda(x)]^k - \mathbb{E}([\mu \star \psi_\lambda(x)]^k)$
- The shift $\tau' \in [-s,s]^2$ captures correlations along nearby edges in μ
- $\lambda = (j,\theta)$, j for the size of structures, θ orientation

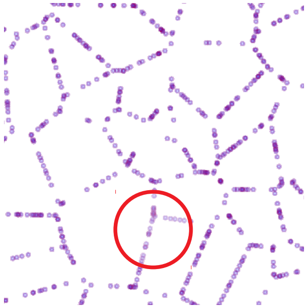

μ

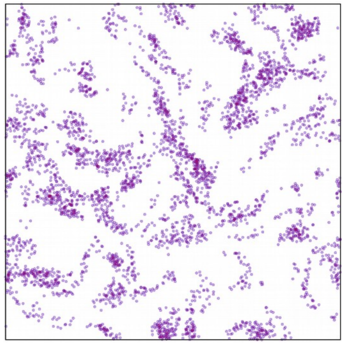

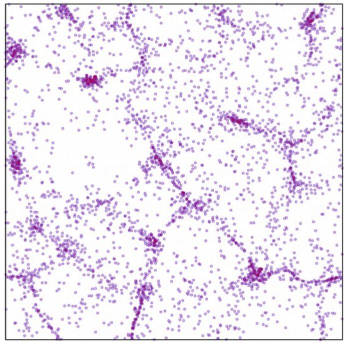

$\tilde{\mu}$ without phase harmonics in K (k=k'=1)

$\Rightarrow$ Importance of $k, k' \neq 1$?

Proposed model

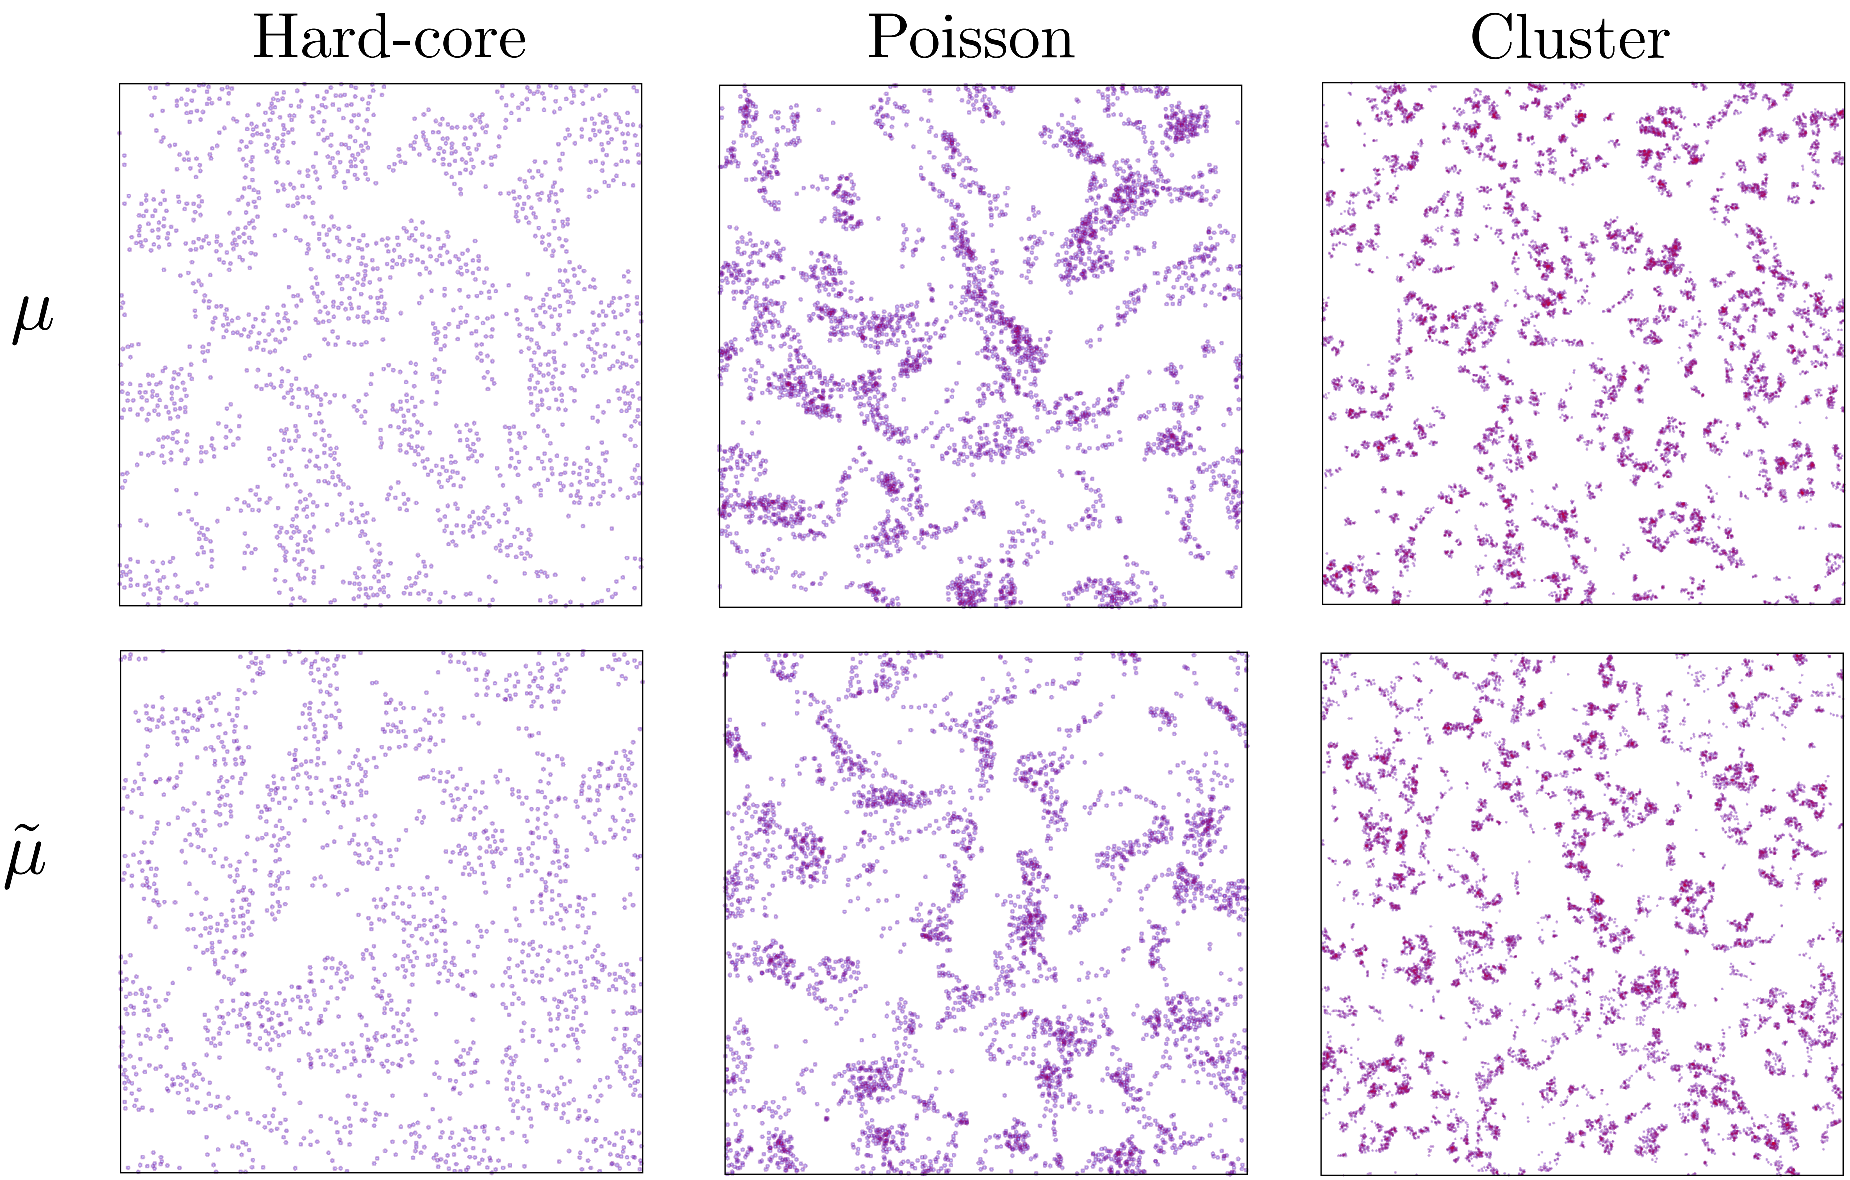

Proposed model

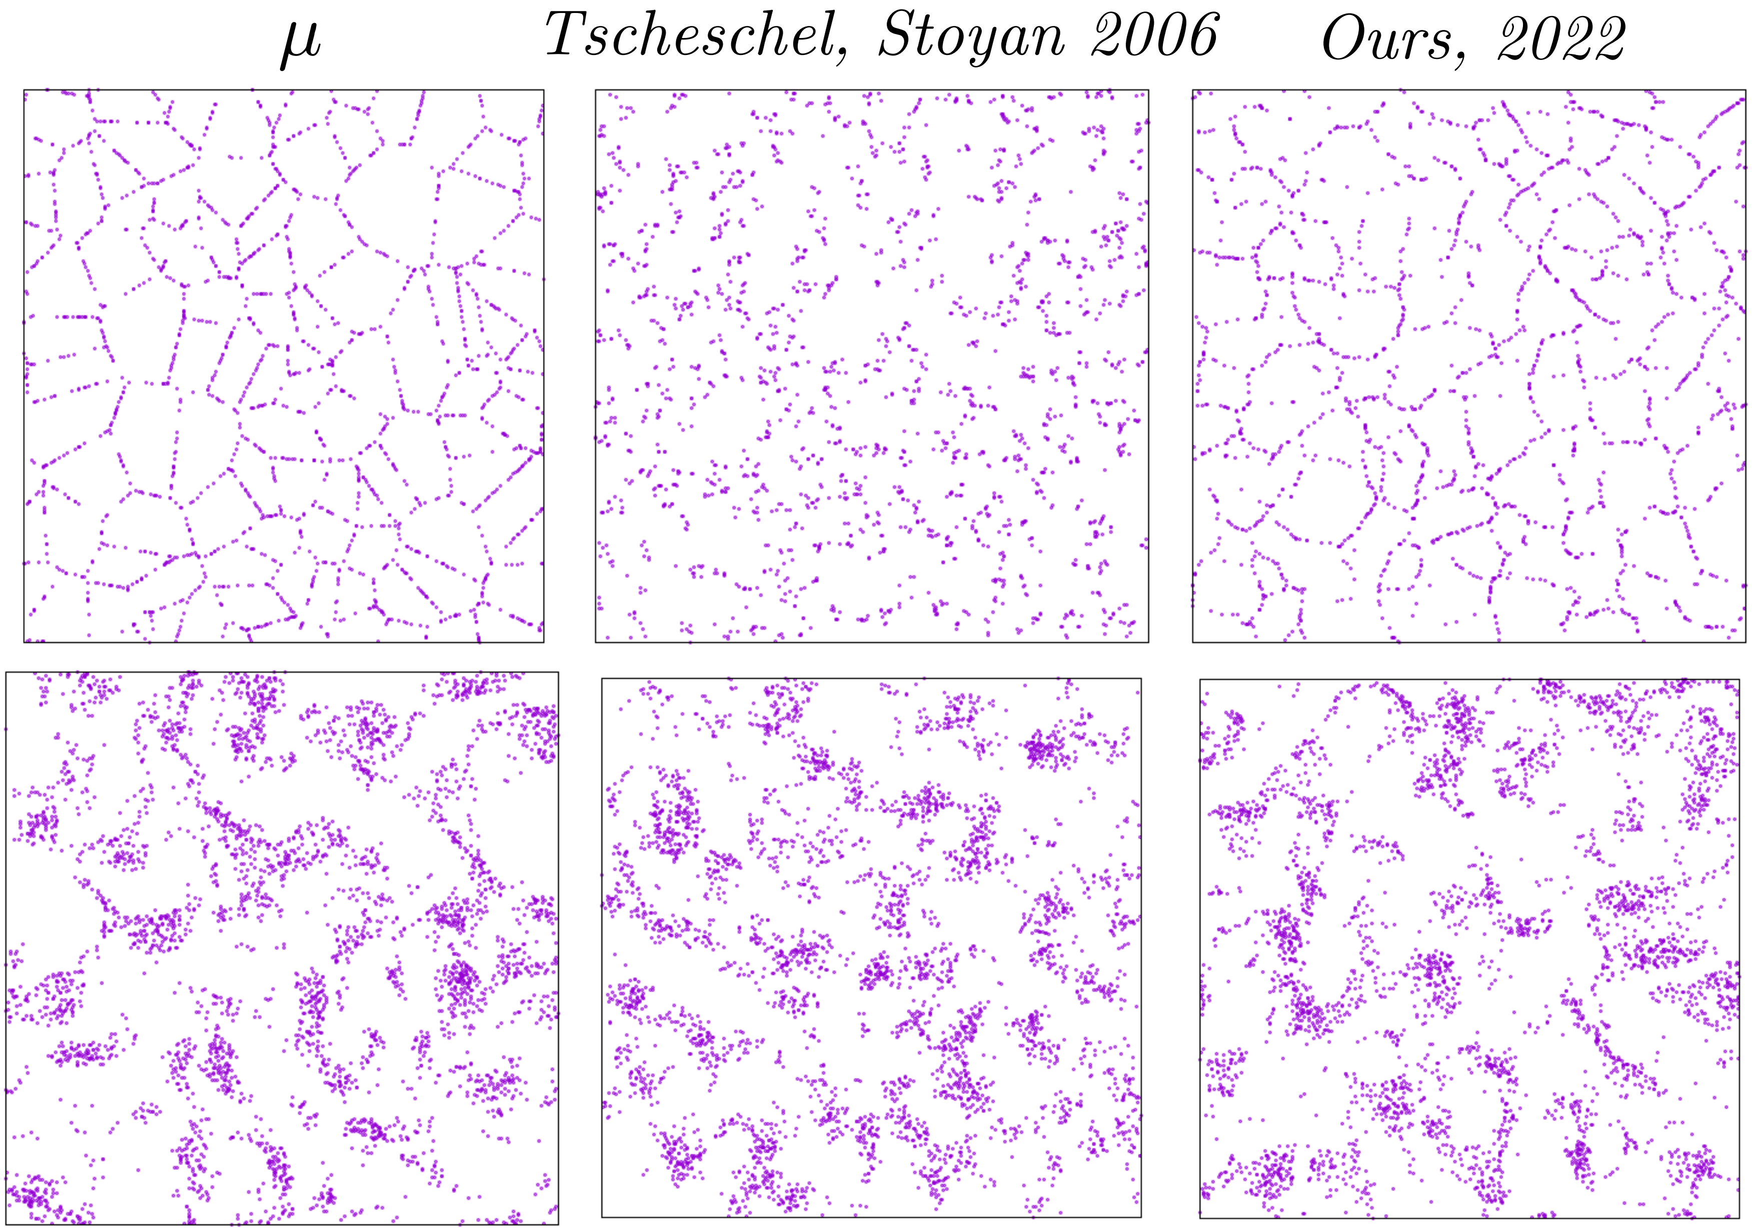

State-of-the-art: Tscheschel, Stoyan 2006

Topological data analysis

- Measure similarity between μ and $\tilde{\mu}$ based on topology
- Count connected components and holes in neighbor graphes
- Compute W_2 distance matrix between persistent diagrams

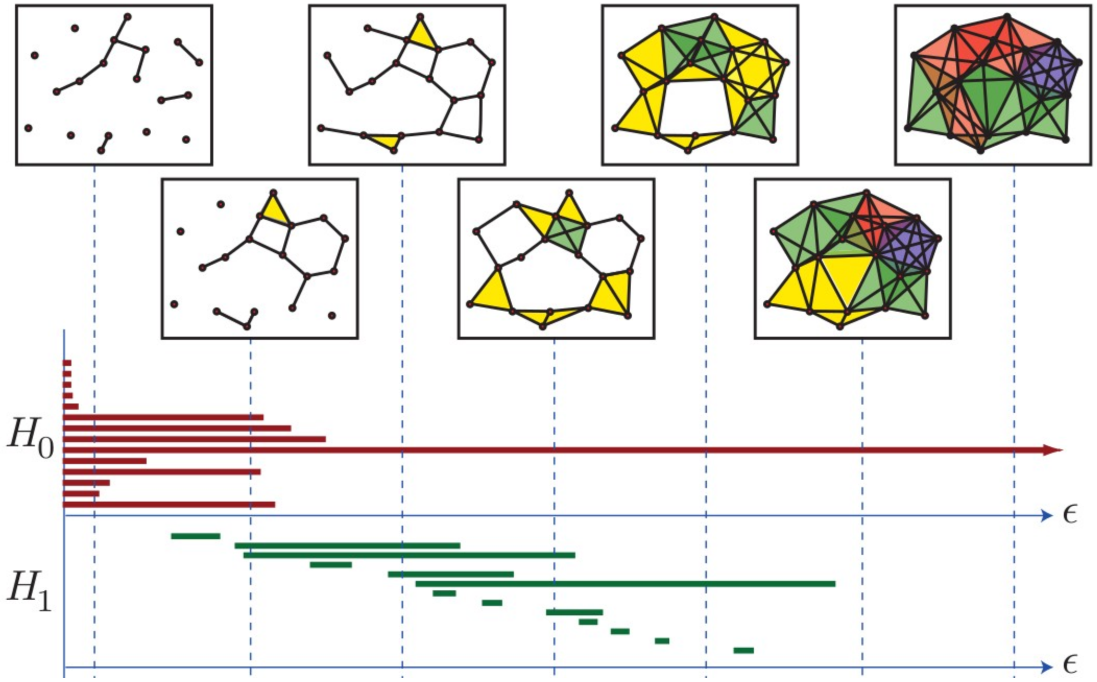

Image from BARCODES:
THE PERSISTENT TOPOLOGY OF DATA, 2007

Topological data analysis

- Measure similarity between μ and $\tilde{\mu}$ based on topology
- Count connected components and holes in neighbor graphes
- Compute W_2 distance matrix between persistent diagrams

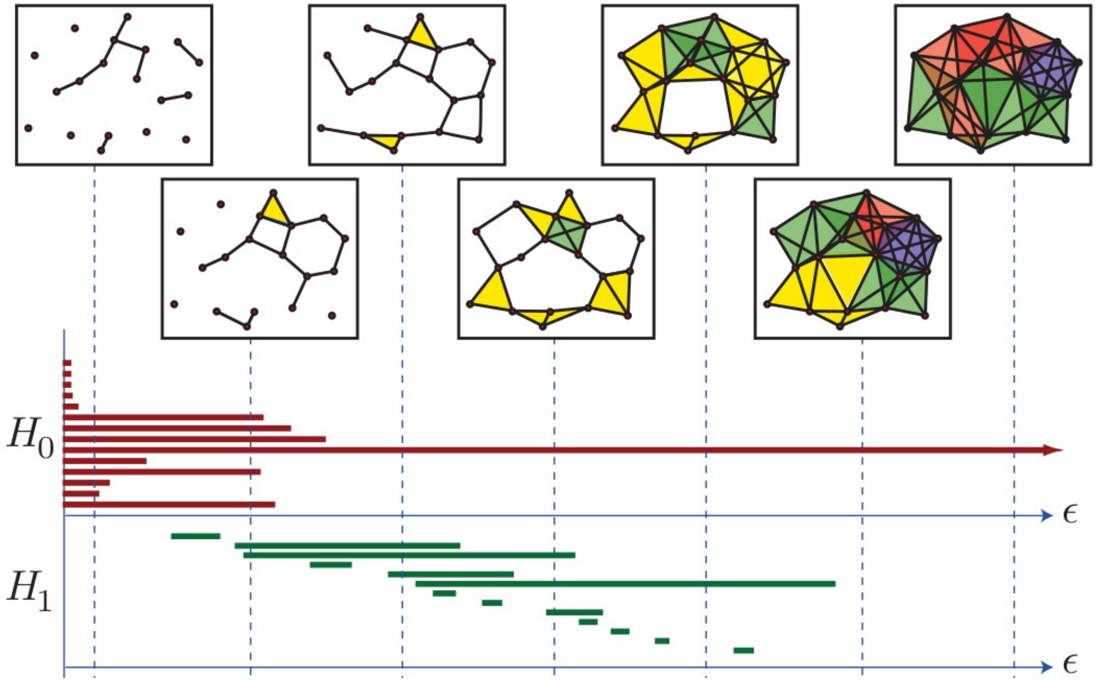

Image from BARCODES:
THE PERSISTENT TOPOLOGY OF DATA, 2007

MDS of distance matrix

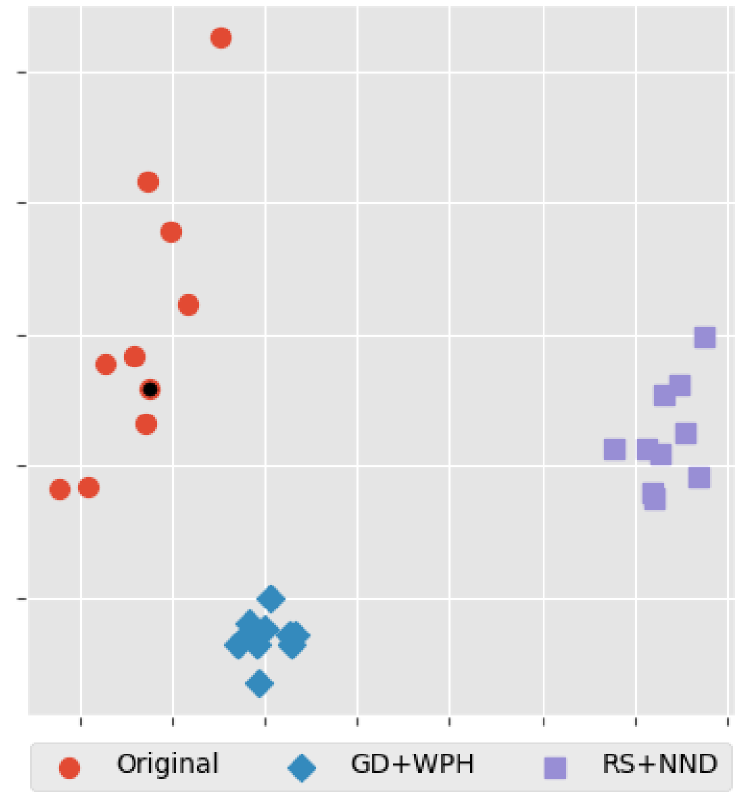

Texture models and cosmology

With E. Allys, F. Levrier, et al.

- Motivation: quantitative comparison between observational and simulated data requires to characterise non-Gaussian processes due to complex nonlinear physics

Texture models and cosmology

With E. Allys, F. Levrier, et al.

• Motivation: quantitative comparison between observational and simulated data requires to characterise non-Gaussian processes due to complex nonlinear physics

• Synthesis problem: Given one observation x, synthesize $\tilde{x}$ such that $\Phi(x) \approx \Phi(\tilde{x})$. If $\tilde{x}$ similar to various $x \sim p$, then Φ characterises the random process p.

Texture models and cosmology

With E. Allys, F. Levrier, et al.

- Motivation: quantitative comparison between observational and simulated data requires to characterise non-Gaussian processes due to complex nonlinear physics

- Synthesis problem: Given one observation x, synthesize $\tilde{x}$ such that $\Phi(x) \approx \Phi(\tilde{x})$. If $\tilde{x}$ similar to various $x \sim p$, then Φ characterises the random process p.

- Reduced wavelet scattering moments: Reduce the dimension of Φ using physical intuitions (seperate scales and angles)

$$\text{e.g. } \Phi_{j_1,j_2,\theta_1,\theta_2}(x) = \log \| |x \star \psi_{j_1,\theta_1}| \star \psi_{j_2,\theta_2} \| \approx \sum_p \Phi'_{j_1,j_2,p}(x) f_{\theta_1,\theta_2,p}$$

The dimension of Φ is reduced using Φ' with small set for p

Reduced wavelet scattering

- On column density map from magnetohydrodynamics

$$\log \| |x \star \psi_{j_1,\theta_1}| \star \psi_{j_2,\theta_2} \| \approx$$

$$\Phi'_{j_1,j_2,iso1}(x) + \Phi'_{j_1,j_2,iso2}(x) \cos(\frac{2\pi}{\Theta}(\theta_1 - \theta_2)) +$$

$$\Phi'_{j_1,j_2,aniso1}(x) \cos(\frac{2\pi}{\Theta}(\theta_1 - \theta^{(ref)}_{j_1,j_2})) + \Phi'_{j_1,j_2,aniso2}(x) \cos(\frac{2\pi}{\Theta}(\theta_2 - \theta^{(ref)}_{j_1,j_2}))$$

Reduced wavelet scattering

- On column density map from magnetohydrodynamics

$$\log \| |x \star \psi_{j_1,\theta_1}| \star \psi_{j_2,\theta_2} \| \approx$$

$$\Phi'_{j_1,j_2,iso1}(x) + \Phi'_{j_1,j_2,iso2}(x) \cos(\frac{2\pi}{\Theta}(\theta_1 - \theta_2)) +$$

$$\Phi'_{j_1,j_2,aniso1}(x) \cos(\frac{2\pi}{\Theta}(\theta_1 - \theta^{(ref)}_{j_1,j_2})) + \Phi'_{j_1,j_2,aniso2}(x) \cos(\frac{2\pi}{\Theta}(\theta_2 - \theta^{(ref)}_{j_1,j_2}))$$

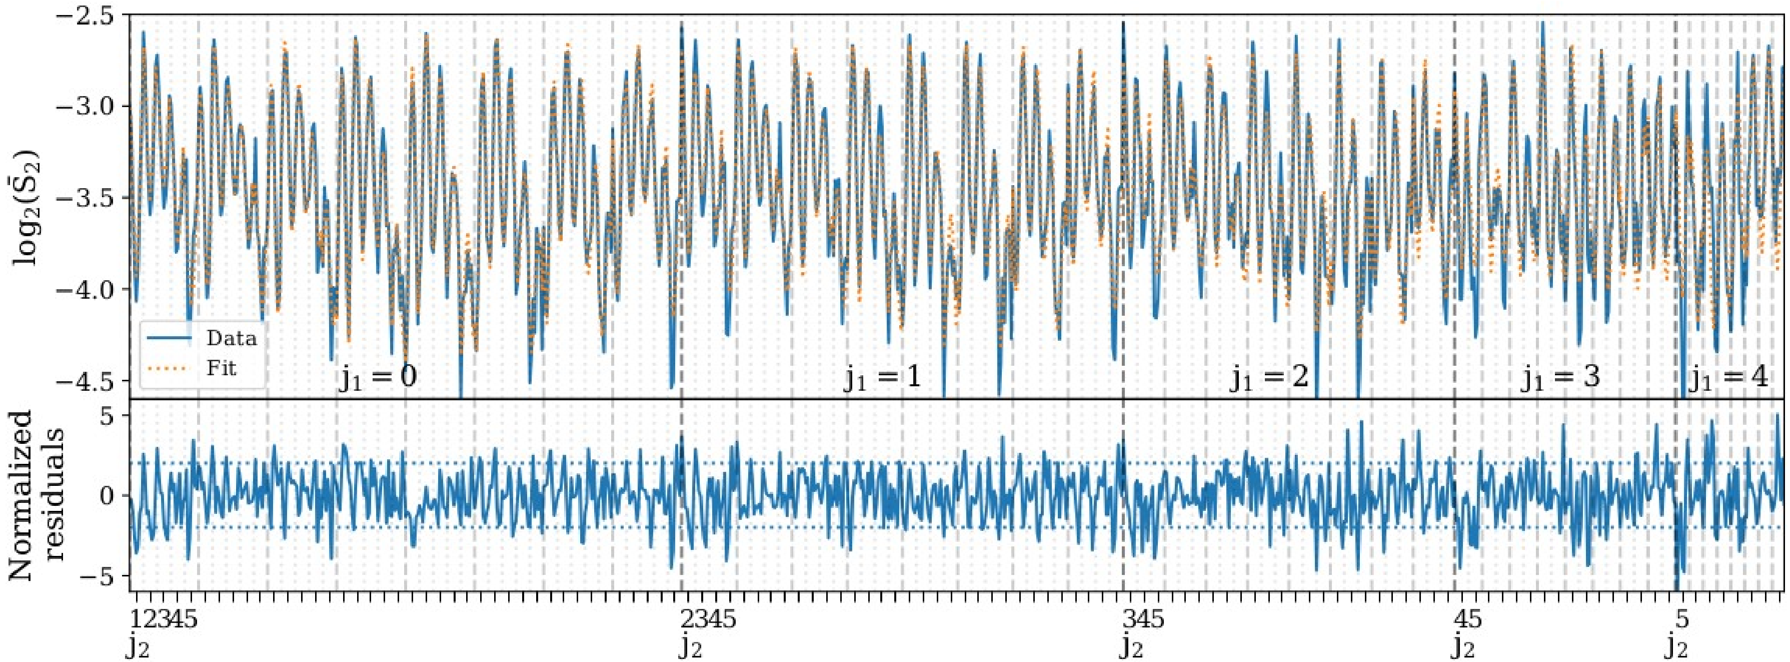

960 coefficients in Φ are fitted with 75 coefficients: $d' << d$

Texture models and cosmology

- Denoising problem: let $y = x^* + \epsilon^*$ with unknown noise ϵ^*, can we recover x^* from y?

Texture models and cosmology

- Denoising problem: let $y = x^* + \epsilon^*$ with unknown noise ϵ^*, can we recover x^* from y?

B. Blancard, E. Allys et al. 2021

- Idea: Find x such that y and $x + \epsilon_i$ is similar for random ϵ_i

$$\min_{x \in \mathbb{R}^d} \frac{1}{M} \sum_{i=1}^{M} \|\Phi(x + \epsilon_i) - \Phi(y)\|^2$$

Texture models and cosmology

- Denoising problem: let $y = x^* + \epsilon^*$ with unknown noise ϵ^*, can we recover x^* from y?

B. Blancard, E. Allys et al. 2021

- Idea: Find x such that y and $x + \epsilon_i$ is similar for random ϵ_i

$$\min_{x \in \mathbb{R}^d} \frac{1}{M} \sum_{i=1}^{M} \|\Phi(x + \epsilon_i) - \Phi(y)\|^2$$

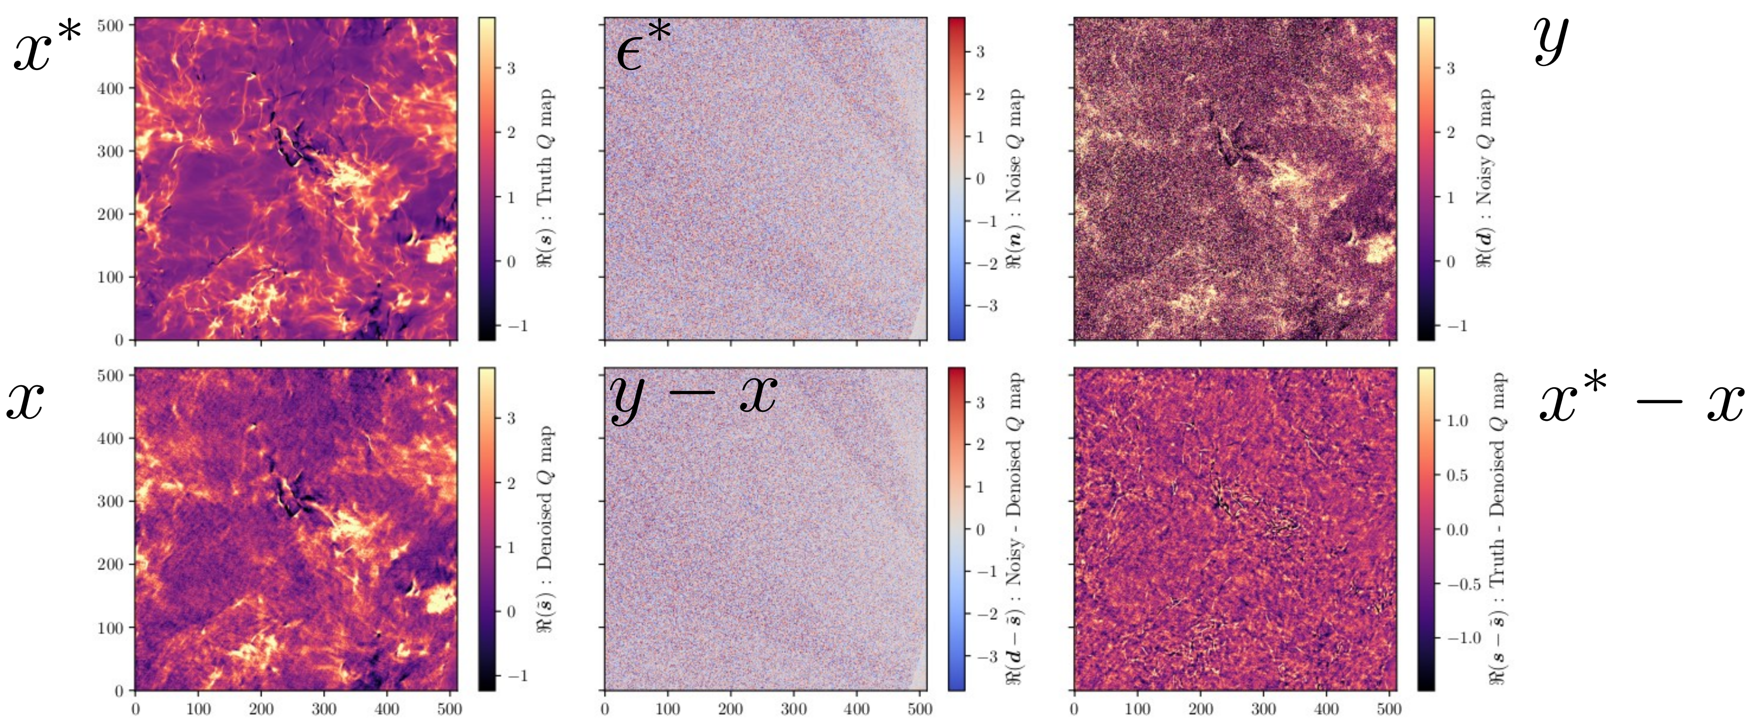

Conclusion

- Propose **wavelet phase harmonic covariance** for texture synthesis (generalizable to turbulence, point process)
  - Bridge the gap between wavelet-based models and CNN based models for high-quality and diverse synthesis
  - Show quality and diversity trade-off through covariance sets
  - Topological data analysis reveals strong connections with geometry
- We use texture models to analyze low-dimensional structures of data and to solve denoising problem in cosmology.

Conclusion

- Propose **wavelet phase harmonic covariance** for texture synthesis (generalizable to turbulence, point process)
  - Bridge the gap between wavelet-based models and CNN based models for high-quality and diverse synthesis
  - Show quality and diversity trade-off through covariance sets
  - Topological data analysis reveals strong connections with geometry
- We use texture models to analyze low-dimensional structures of data and to solve denoising problem in cosmology.
- Future directions:
  - Understand better the mathematics behind these models (Math)
  - Relate to physical models of turbulence (Math/Physics)
  - Relate to generative models of natural images (ML/SP)
  - Build large-scale models using parallel computing (ML/CS)