Describing complexity

Brice Ménard

Johns Hopkins University

& collaboration with OpenAI

Interpretability

CNN InceptionV1

convolutional neural nets

complexity

Estimation of cosmological parameters

Ravanbakhsh et al. (2016)

Gaussian statistics

manifold learning

dictionary learning

convolutional neural nets

complexity

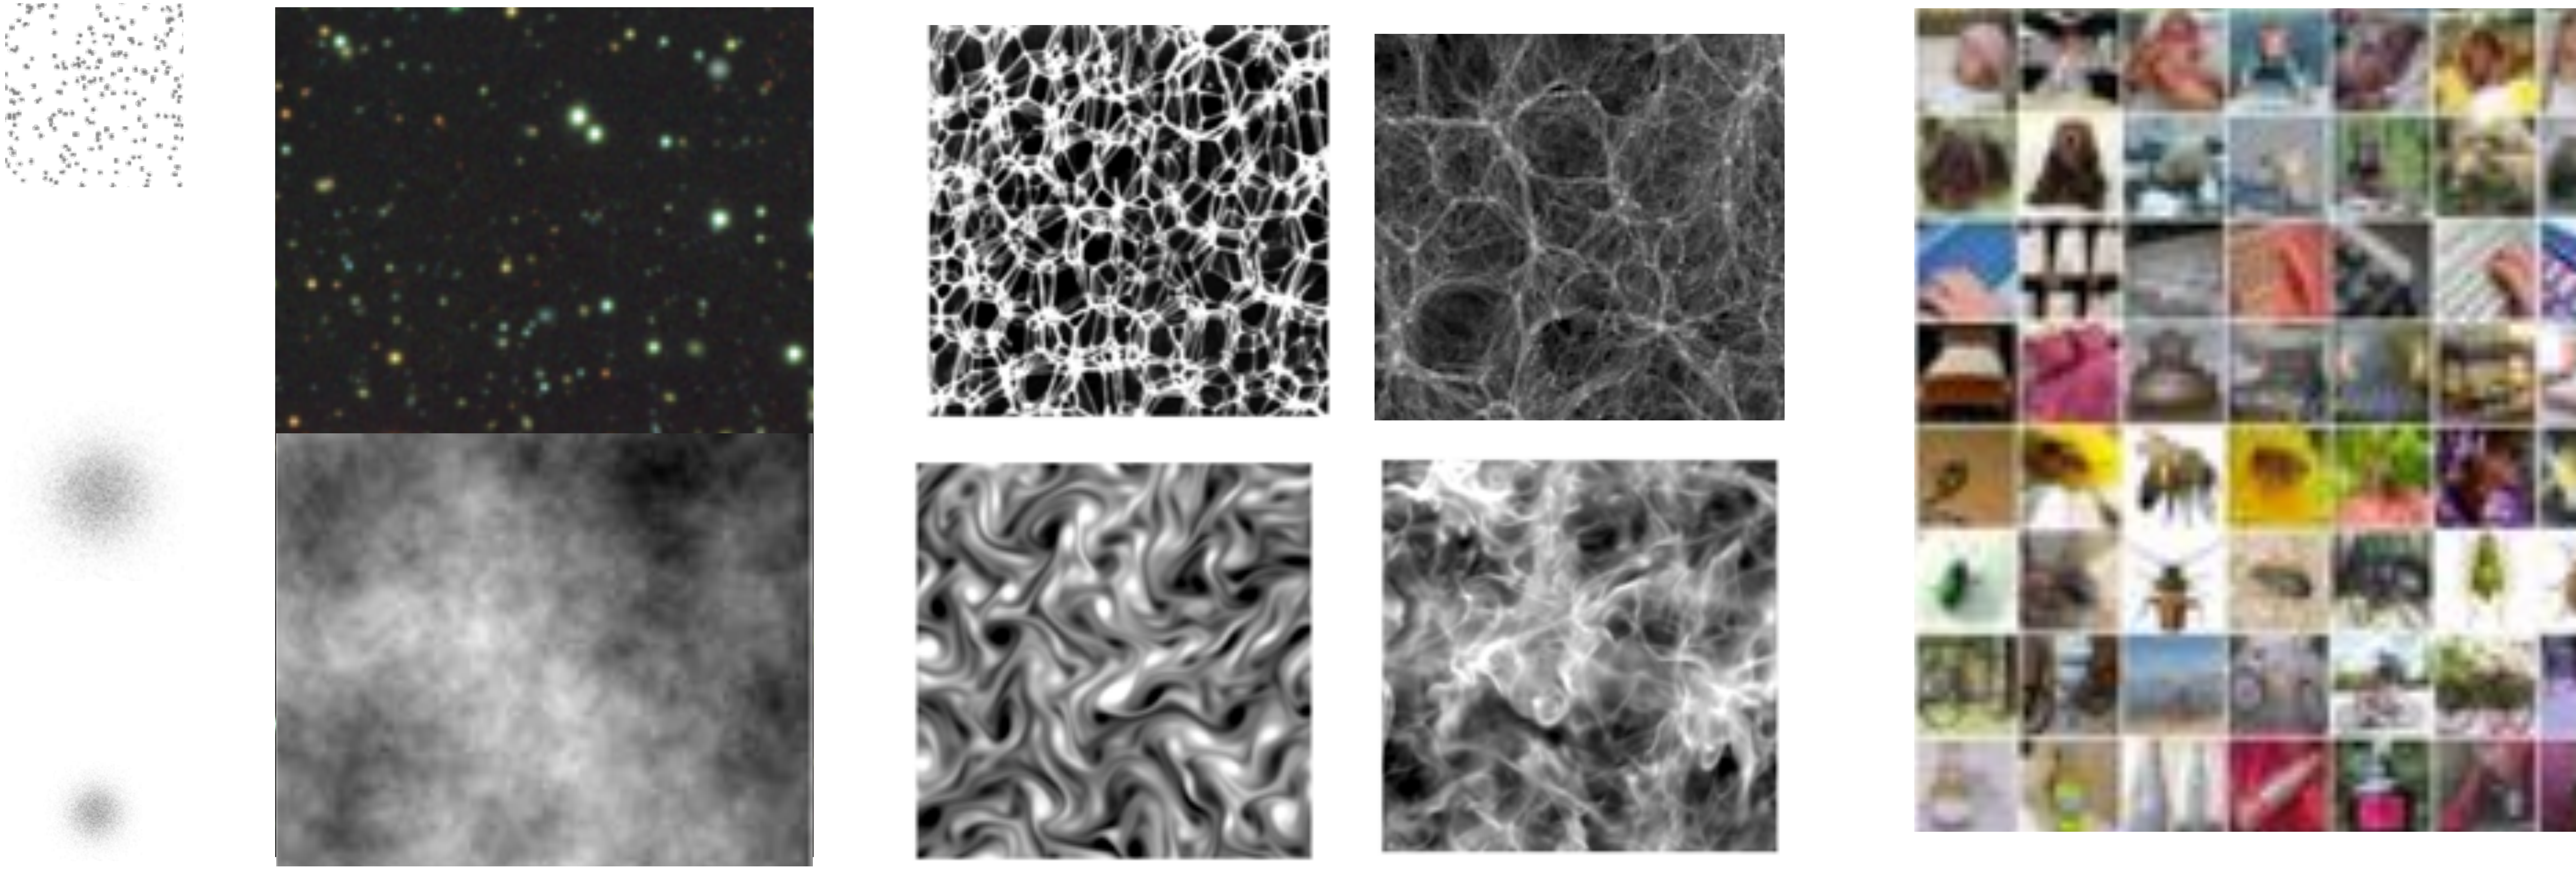

Gaussian statistics

manifold learning

dictionary learning

convolutional neural nets

complexity

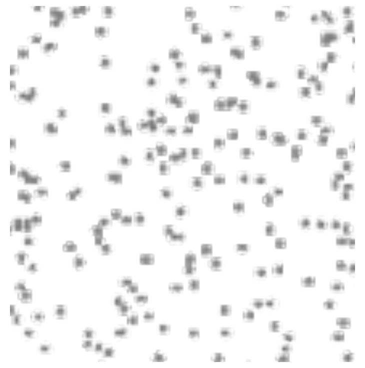

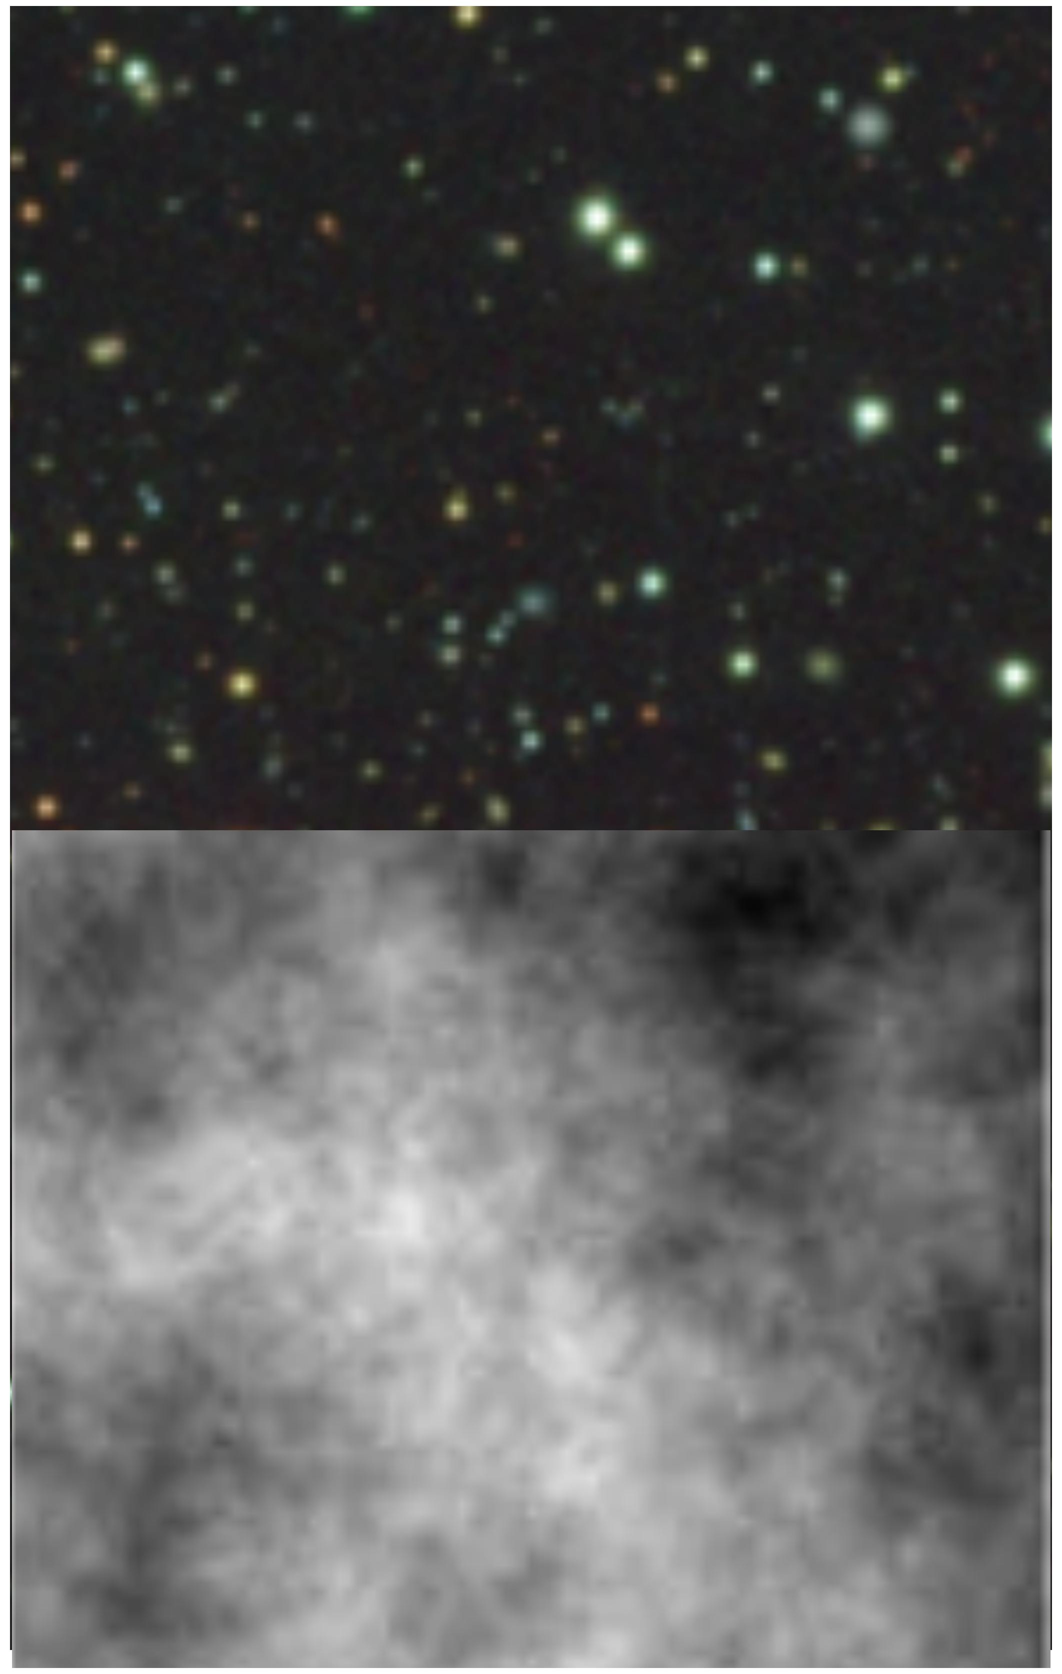

Gaussian statistics

manifold learning

dictionary learning

convolutional neural nets

complexity

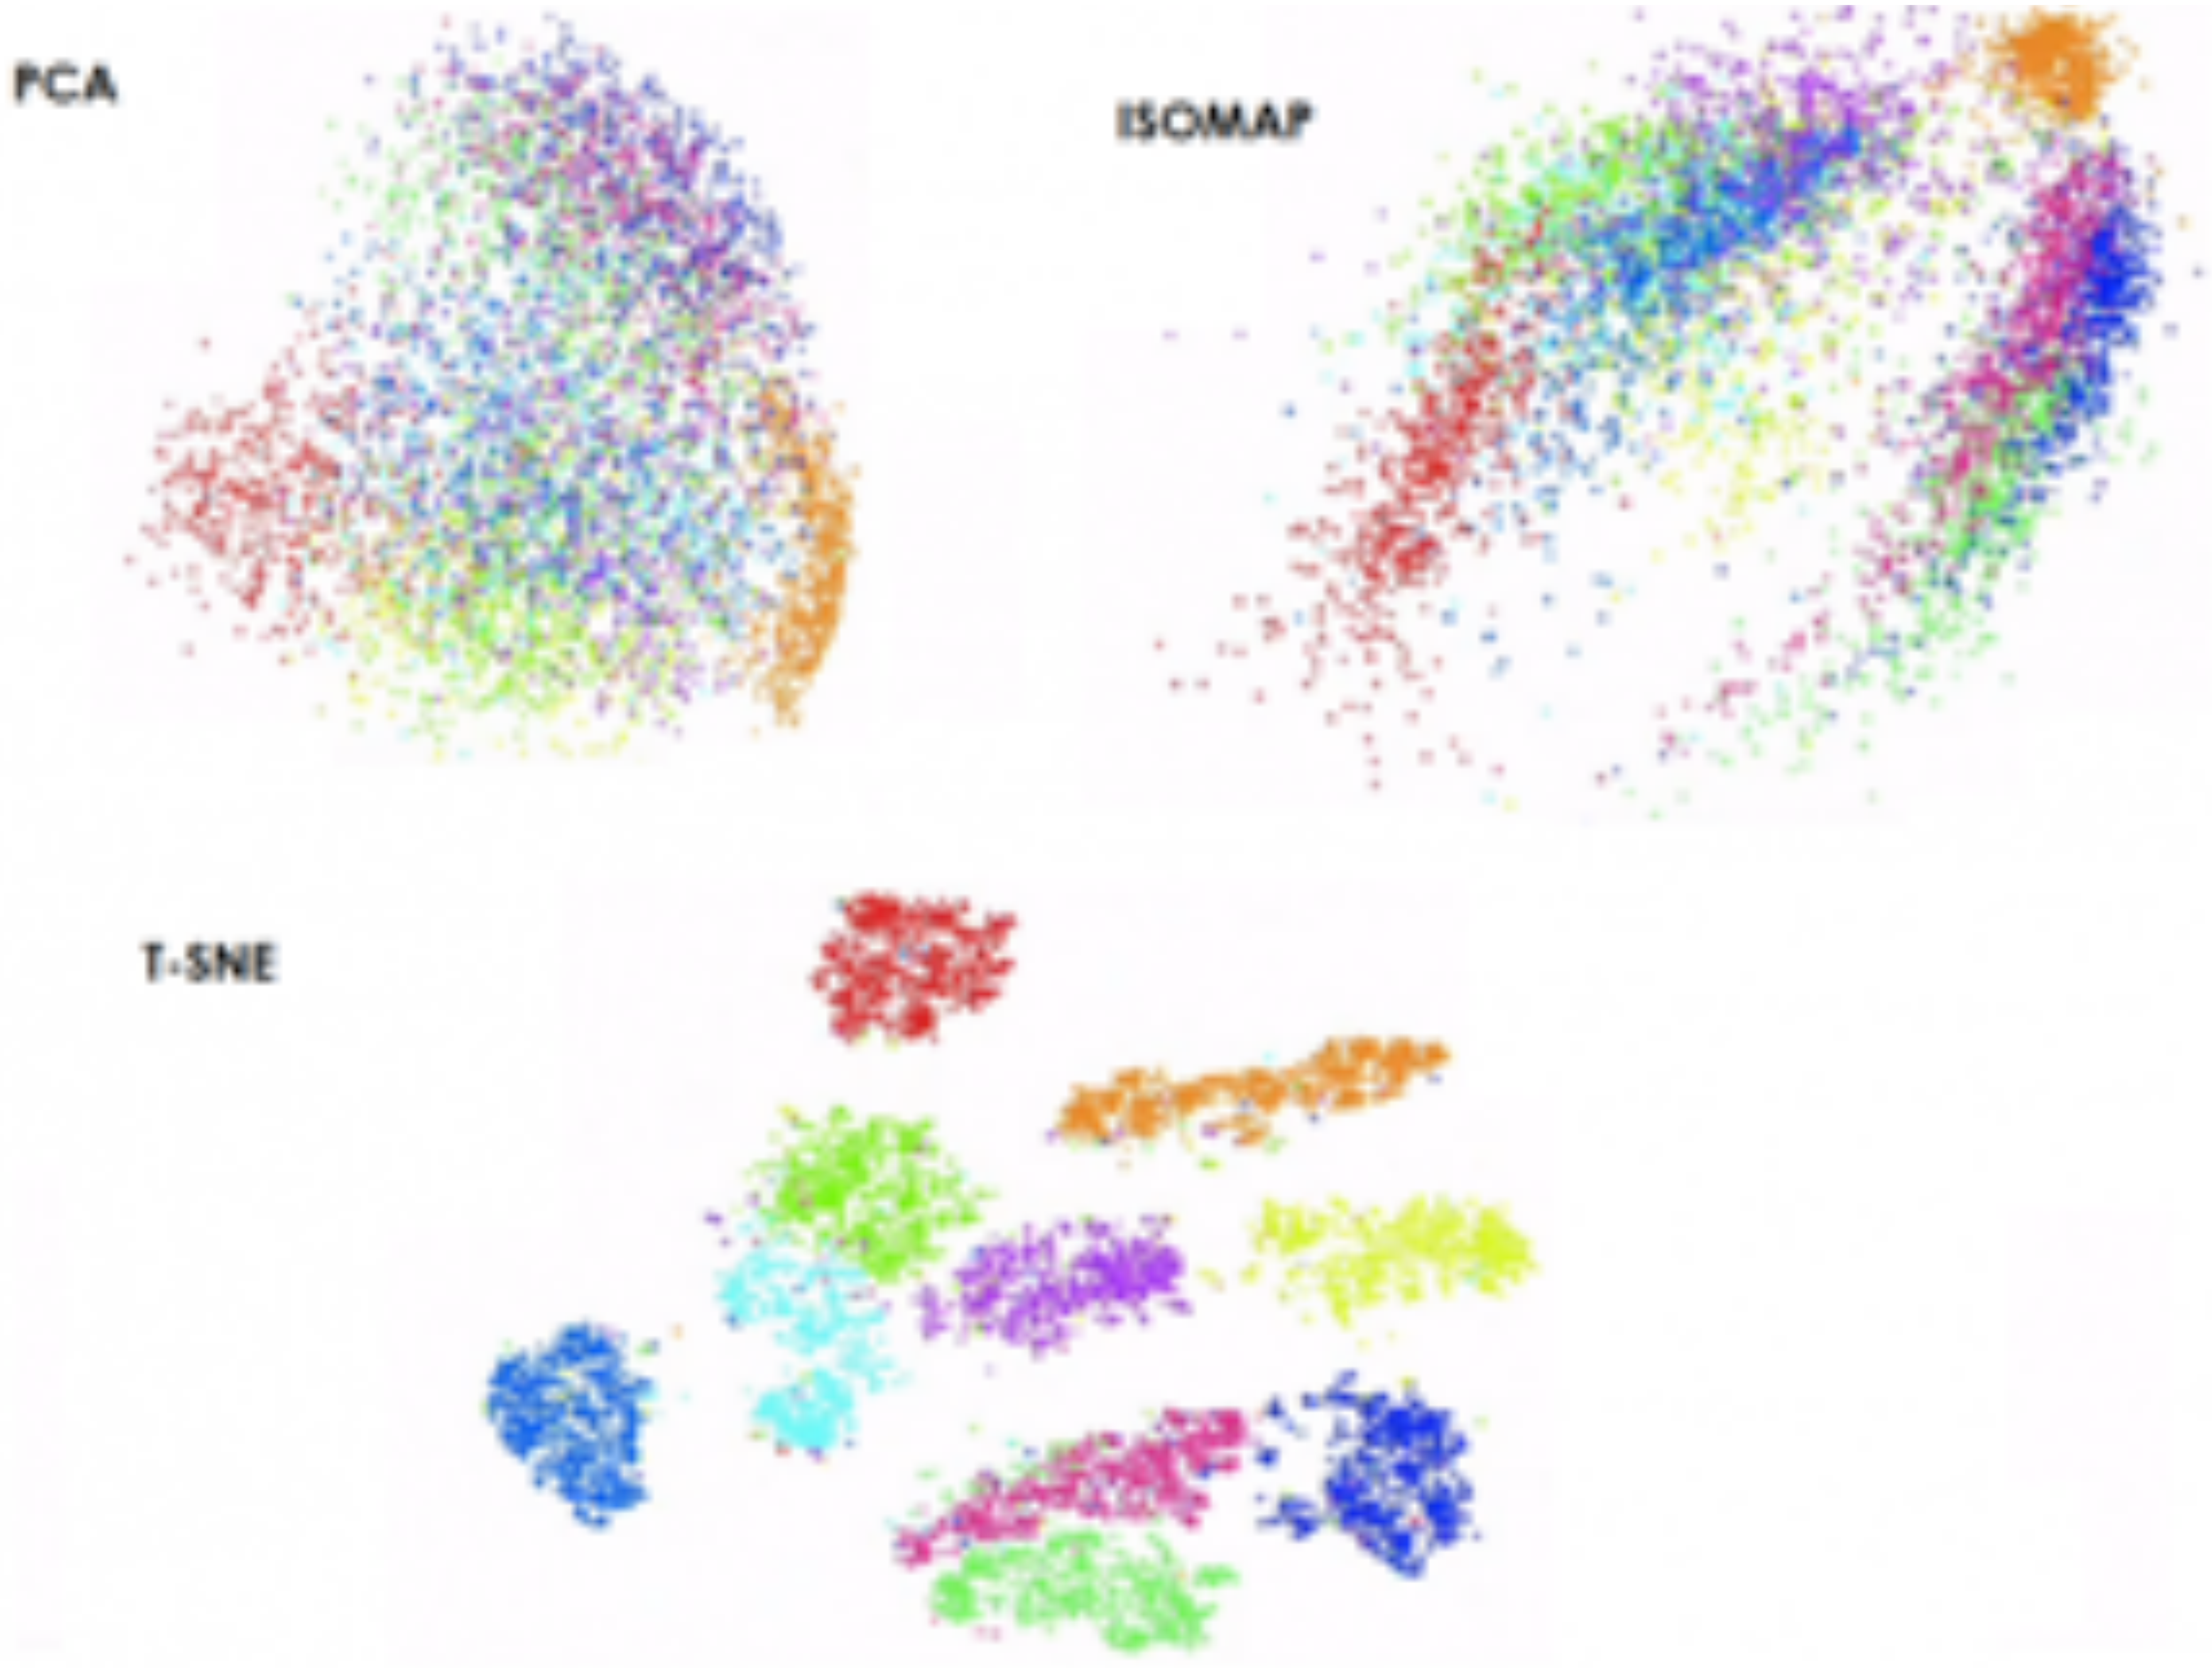

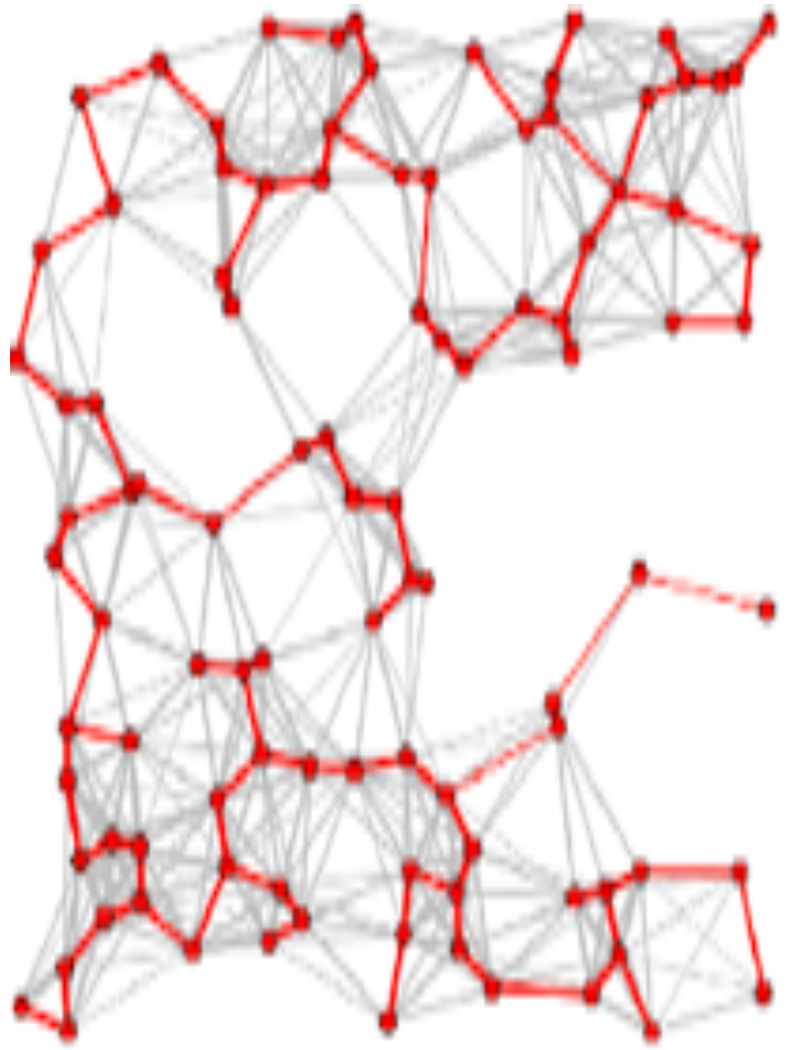

TheSequencer.org

Gaussian statistics

manifold learning

dictionary learning

convolutional neural nets

complexity

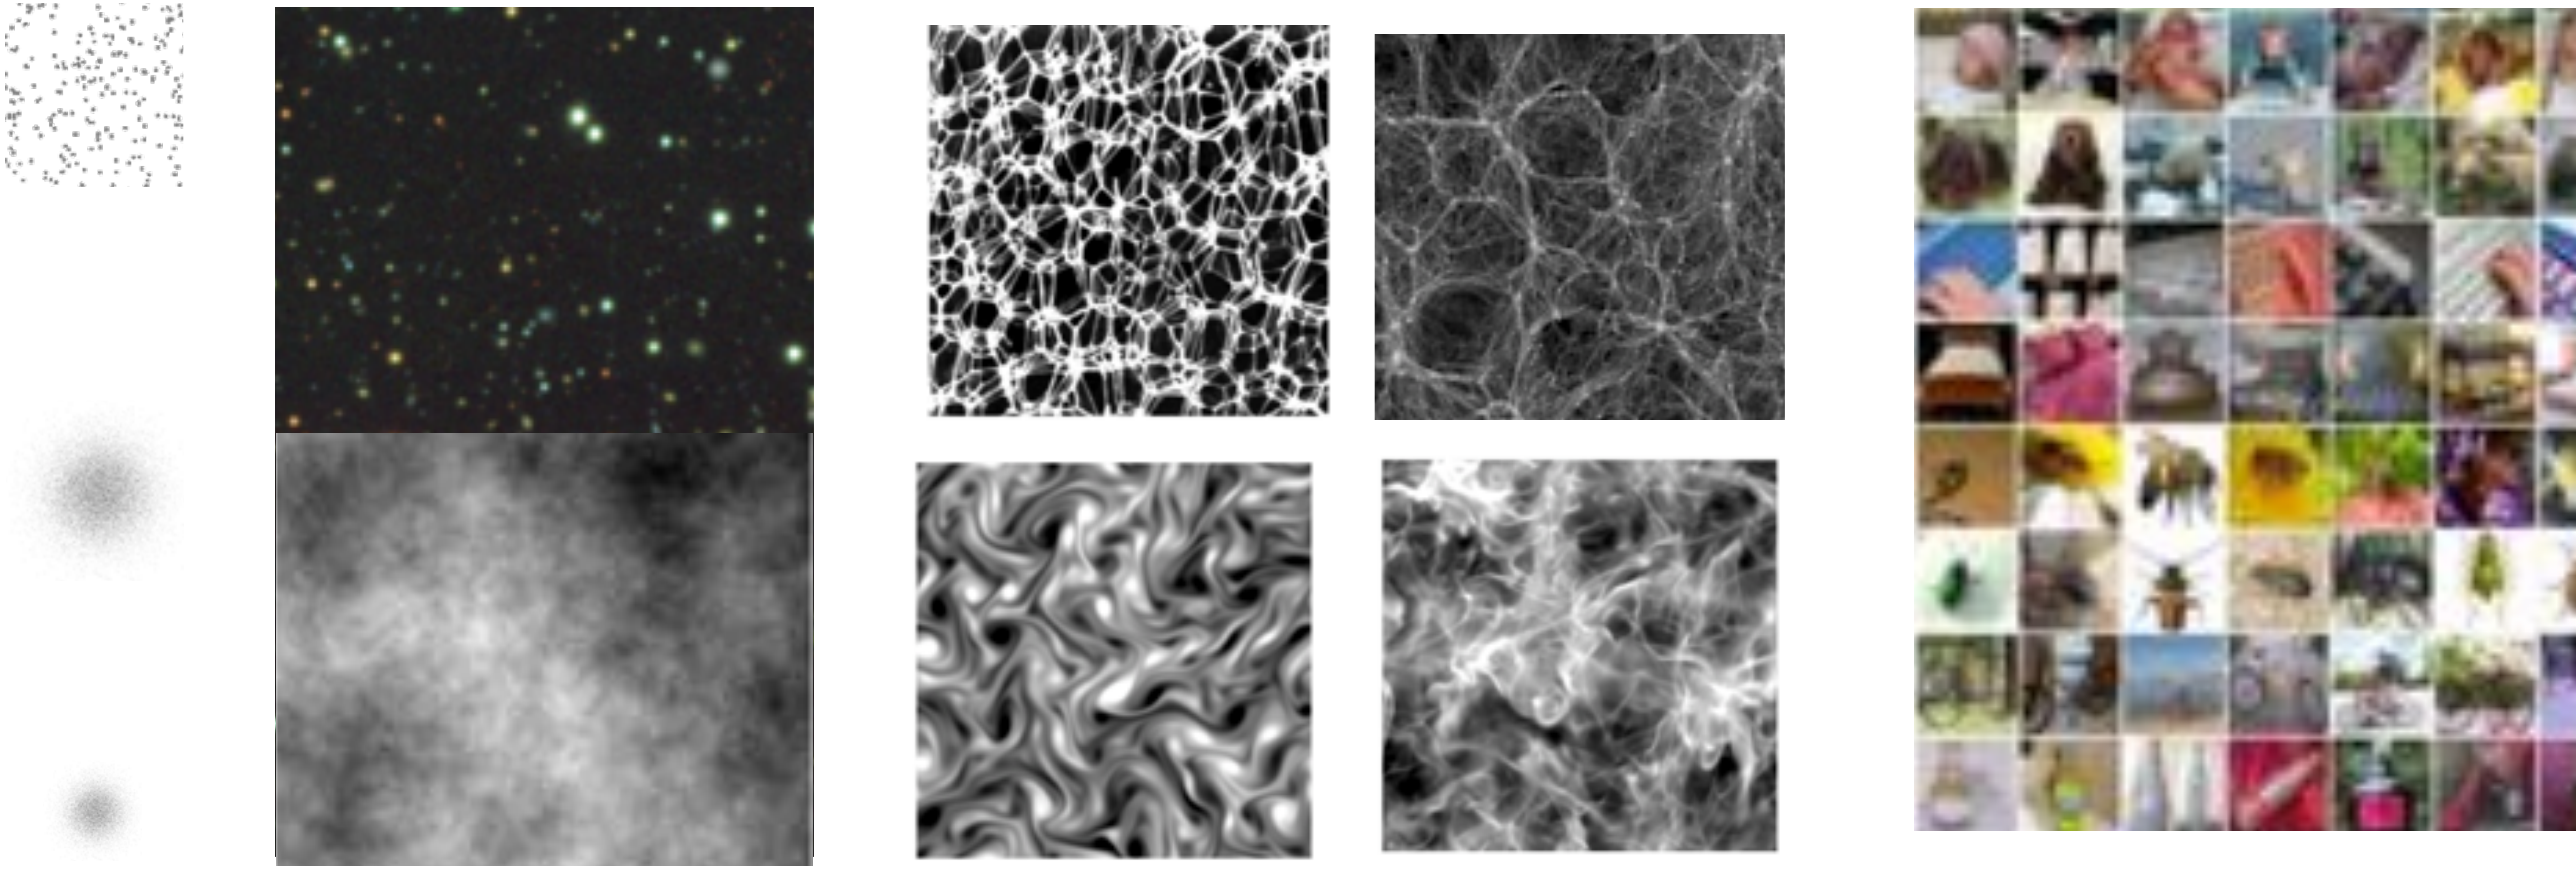

texture classification

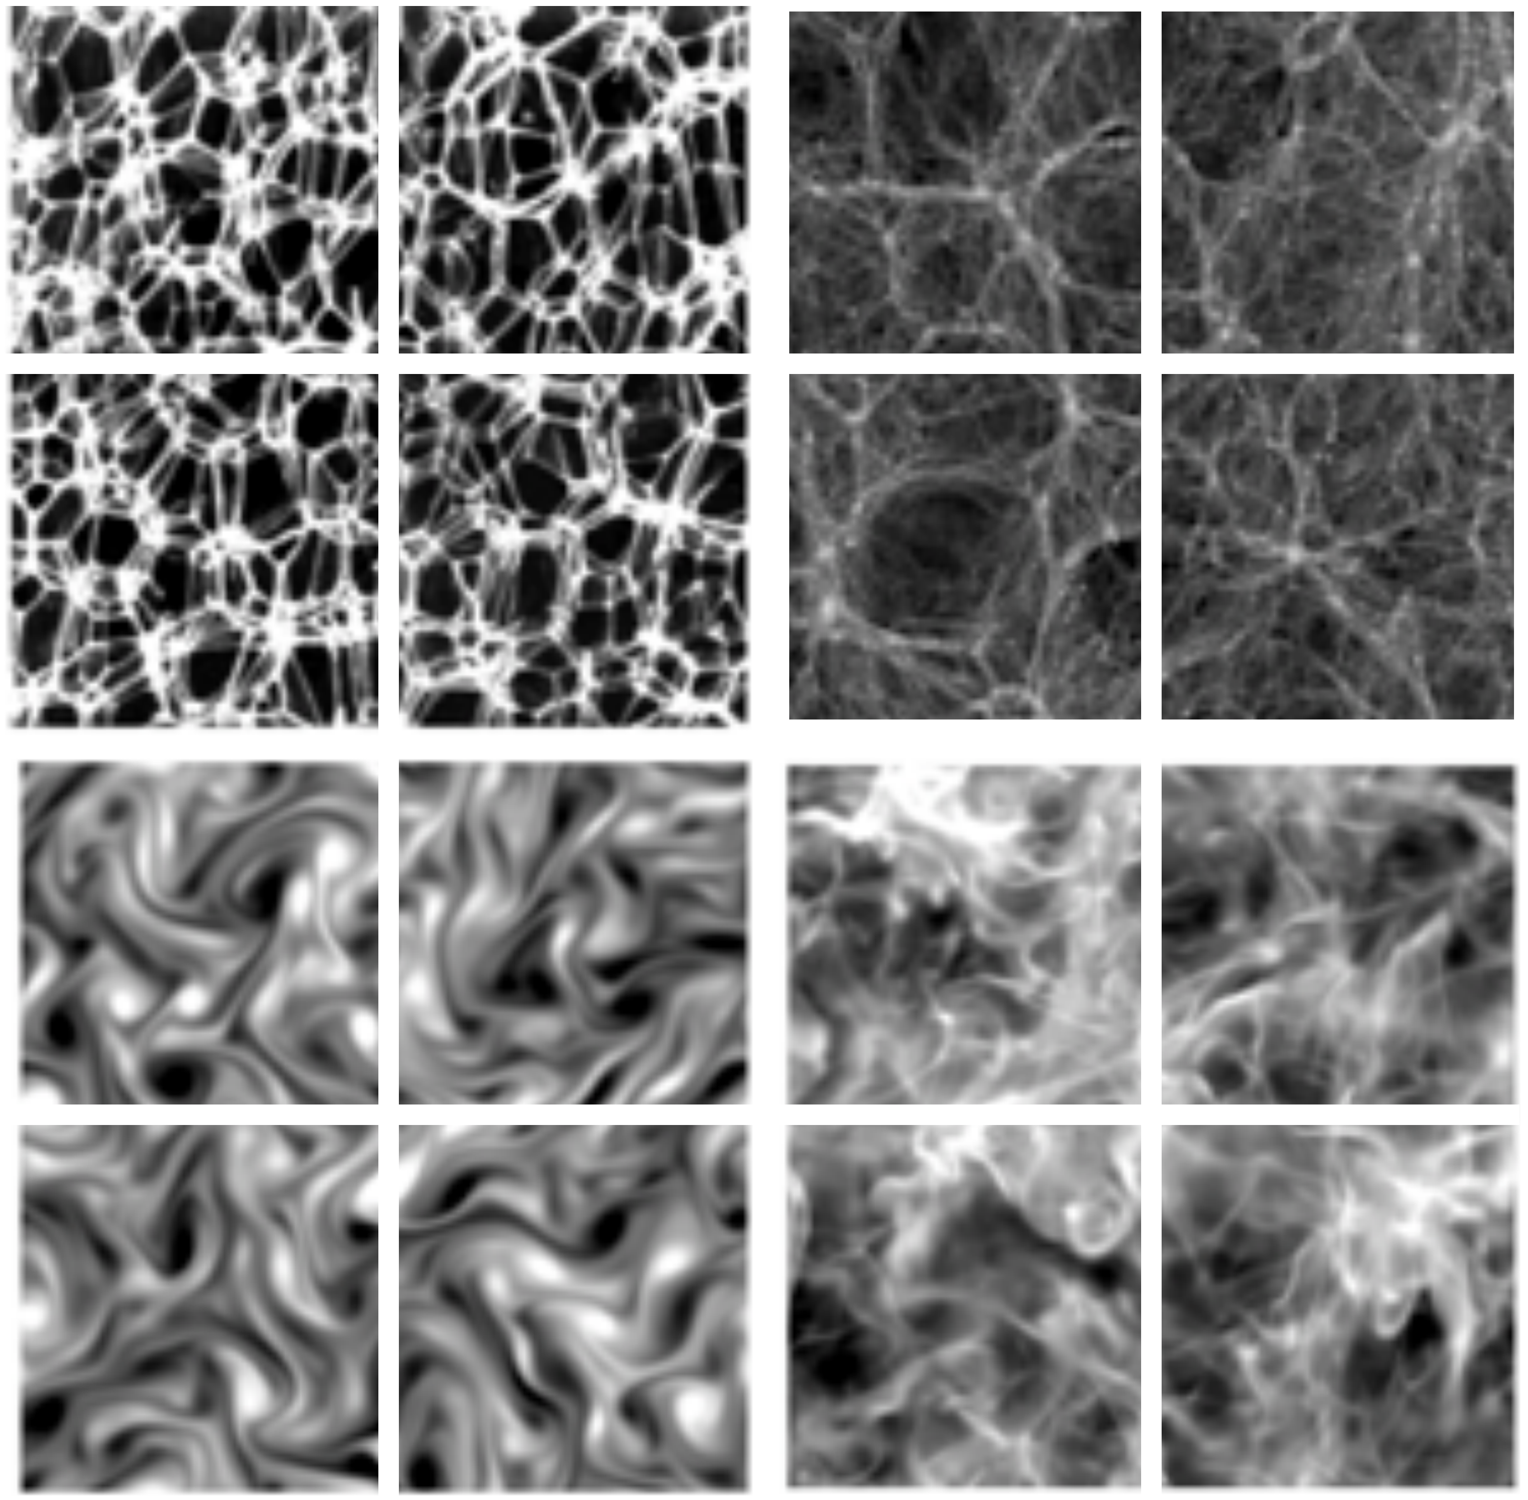

object classification

airplane

horse

cat

rabbit

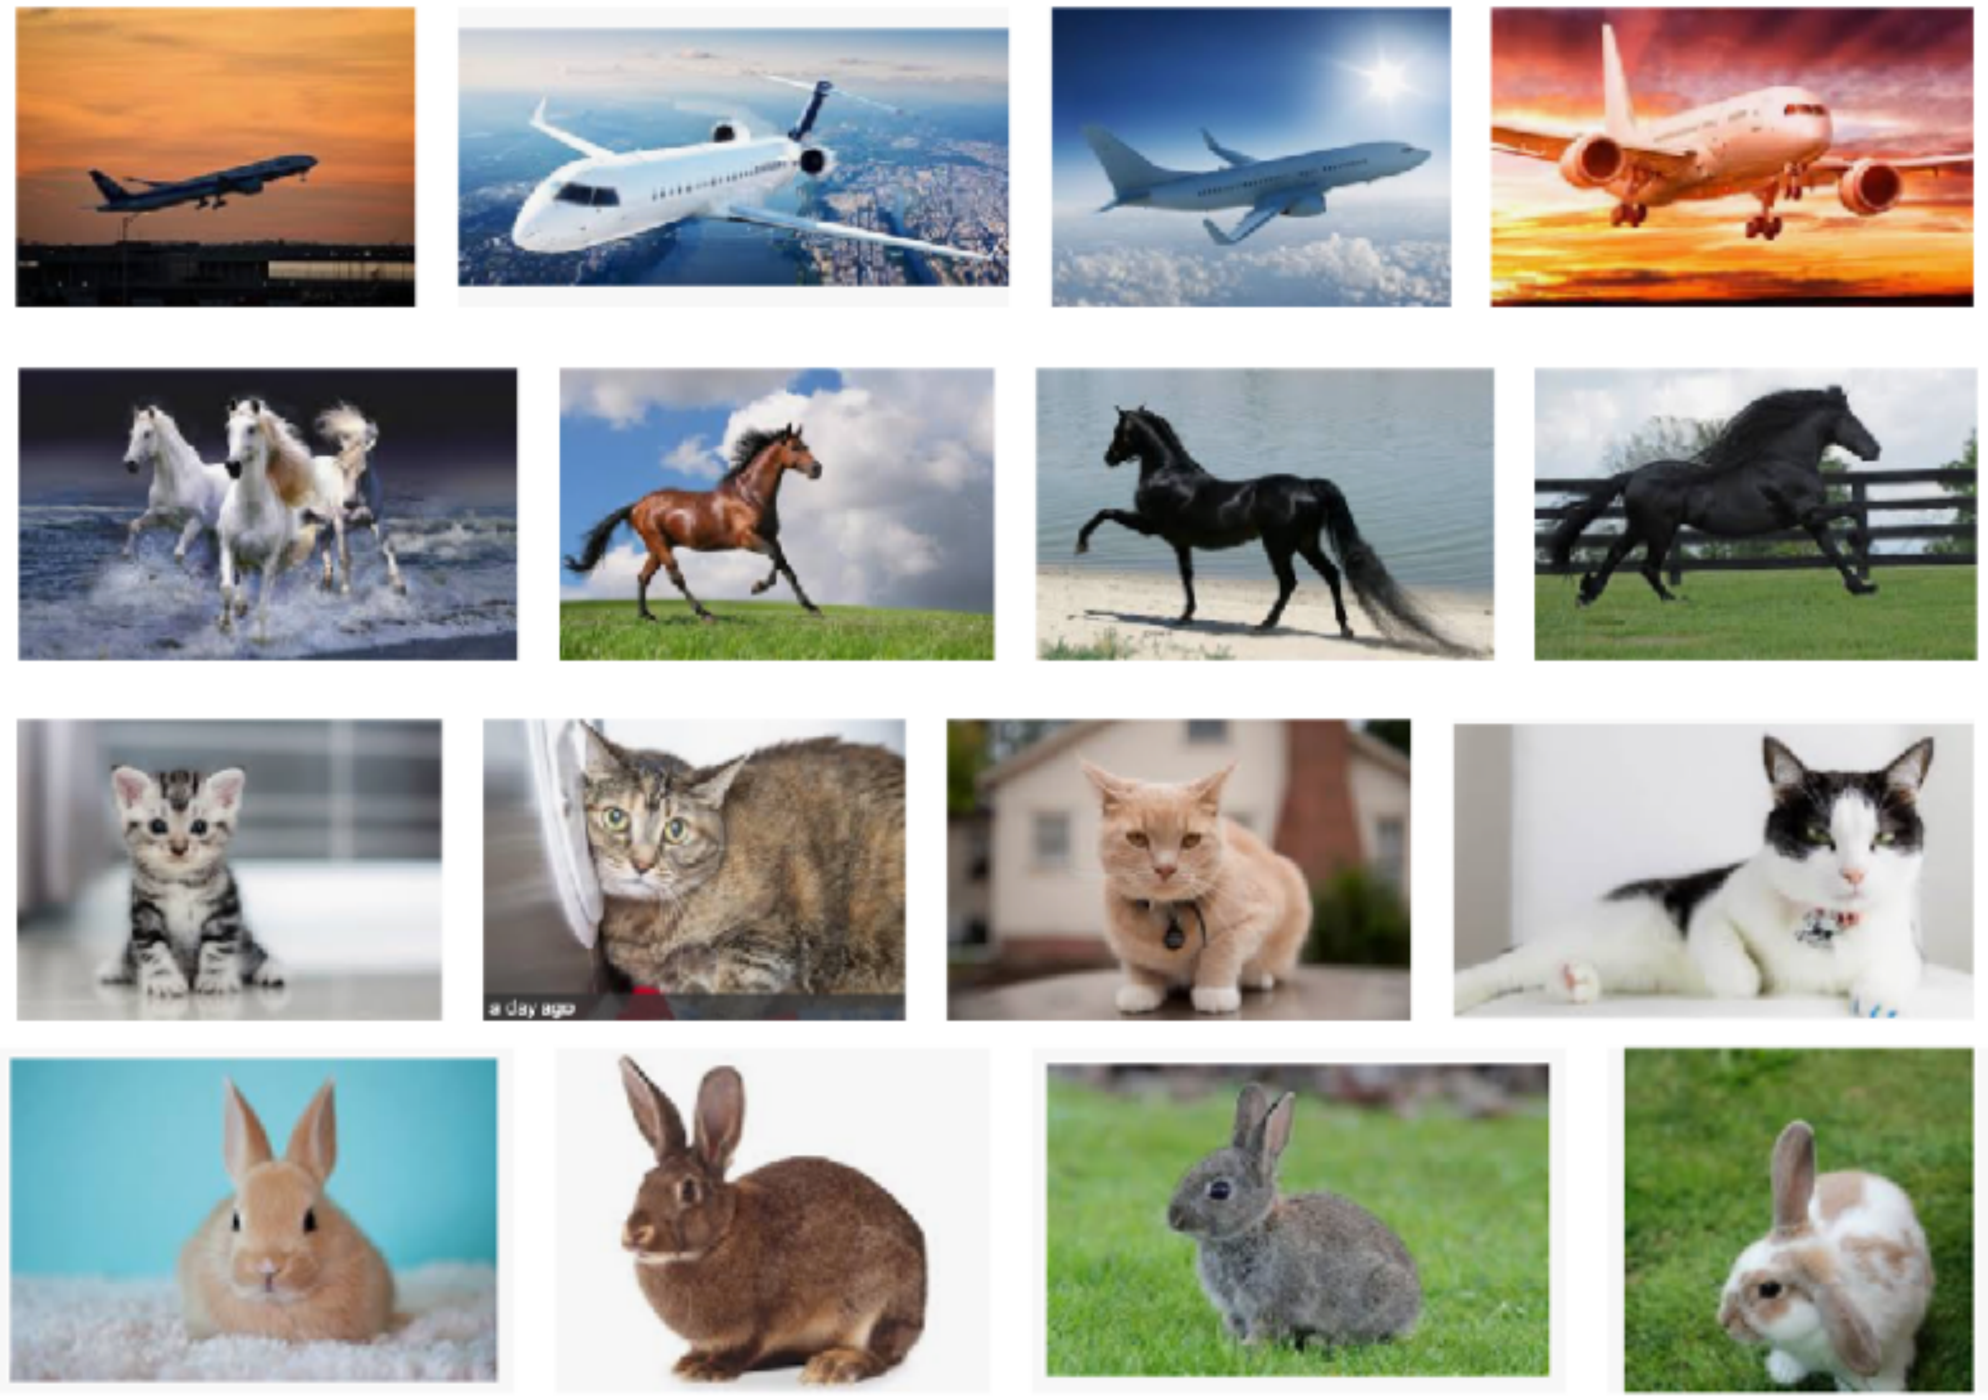

Data $x \in \mathbb{R}^d$

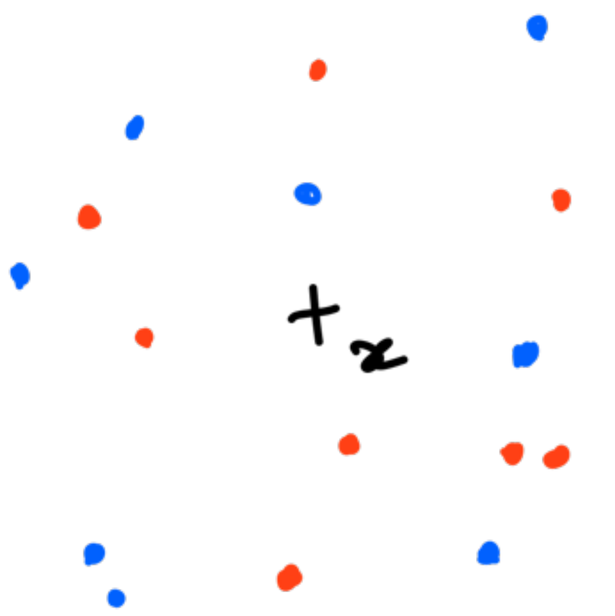

Representation

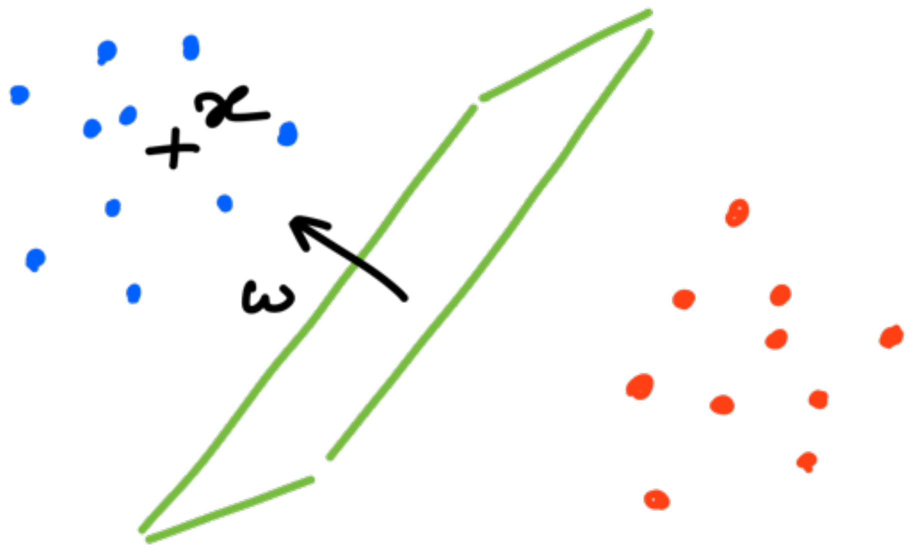

We need to make use of symmetries
This is similar to what we do in physics.

- translations
- rotations
- scaling
- small deformations
- change in brightness / illumination
- additive constant

Using the power spectrum (invariant to translation)

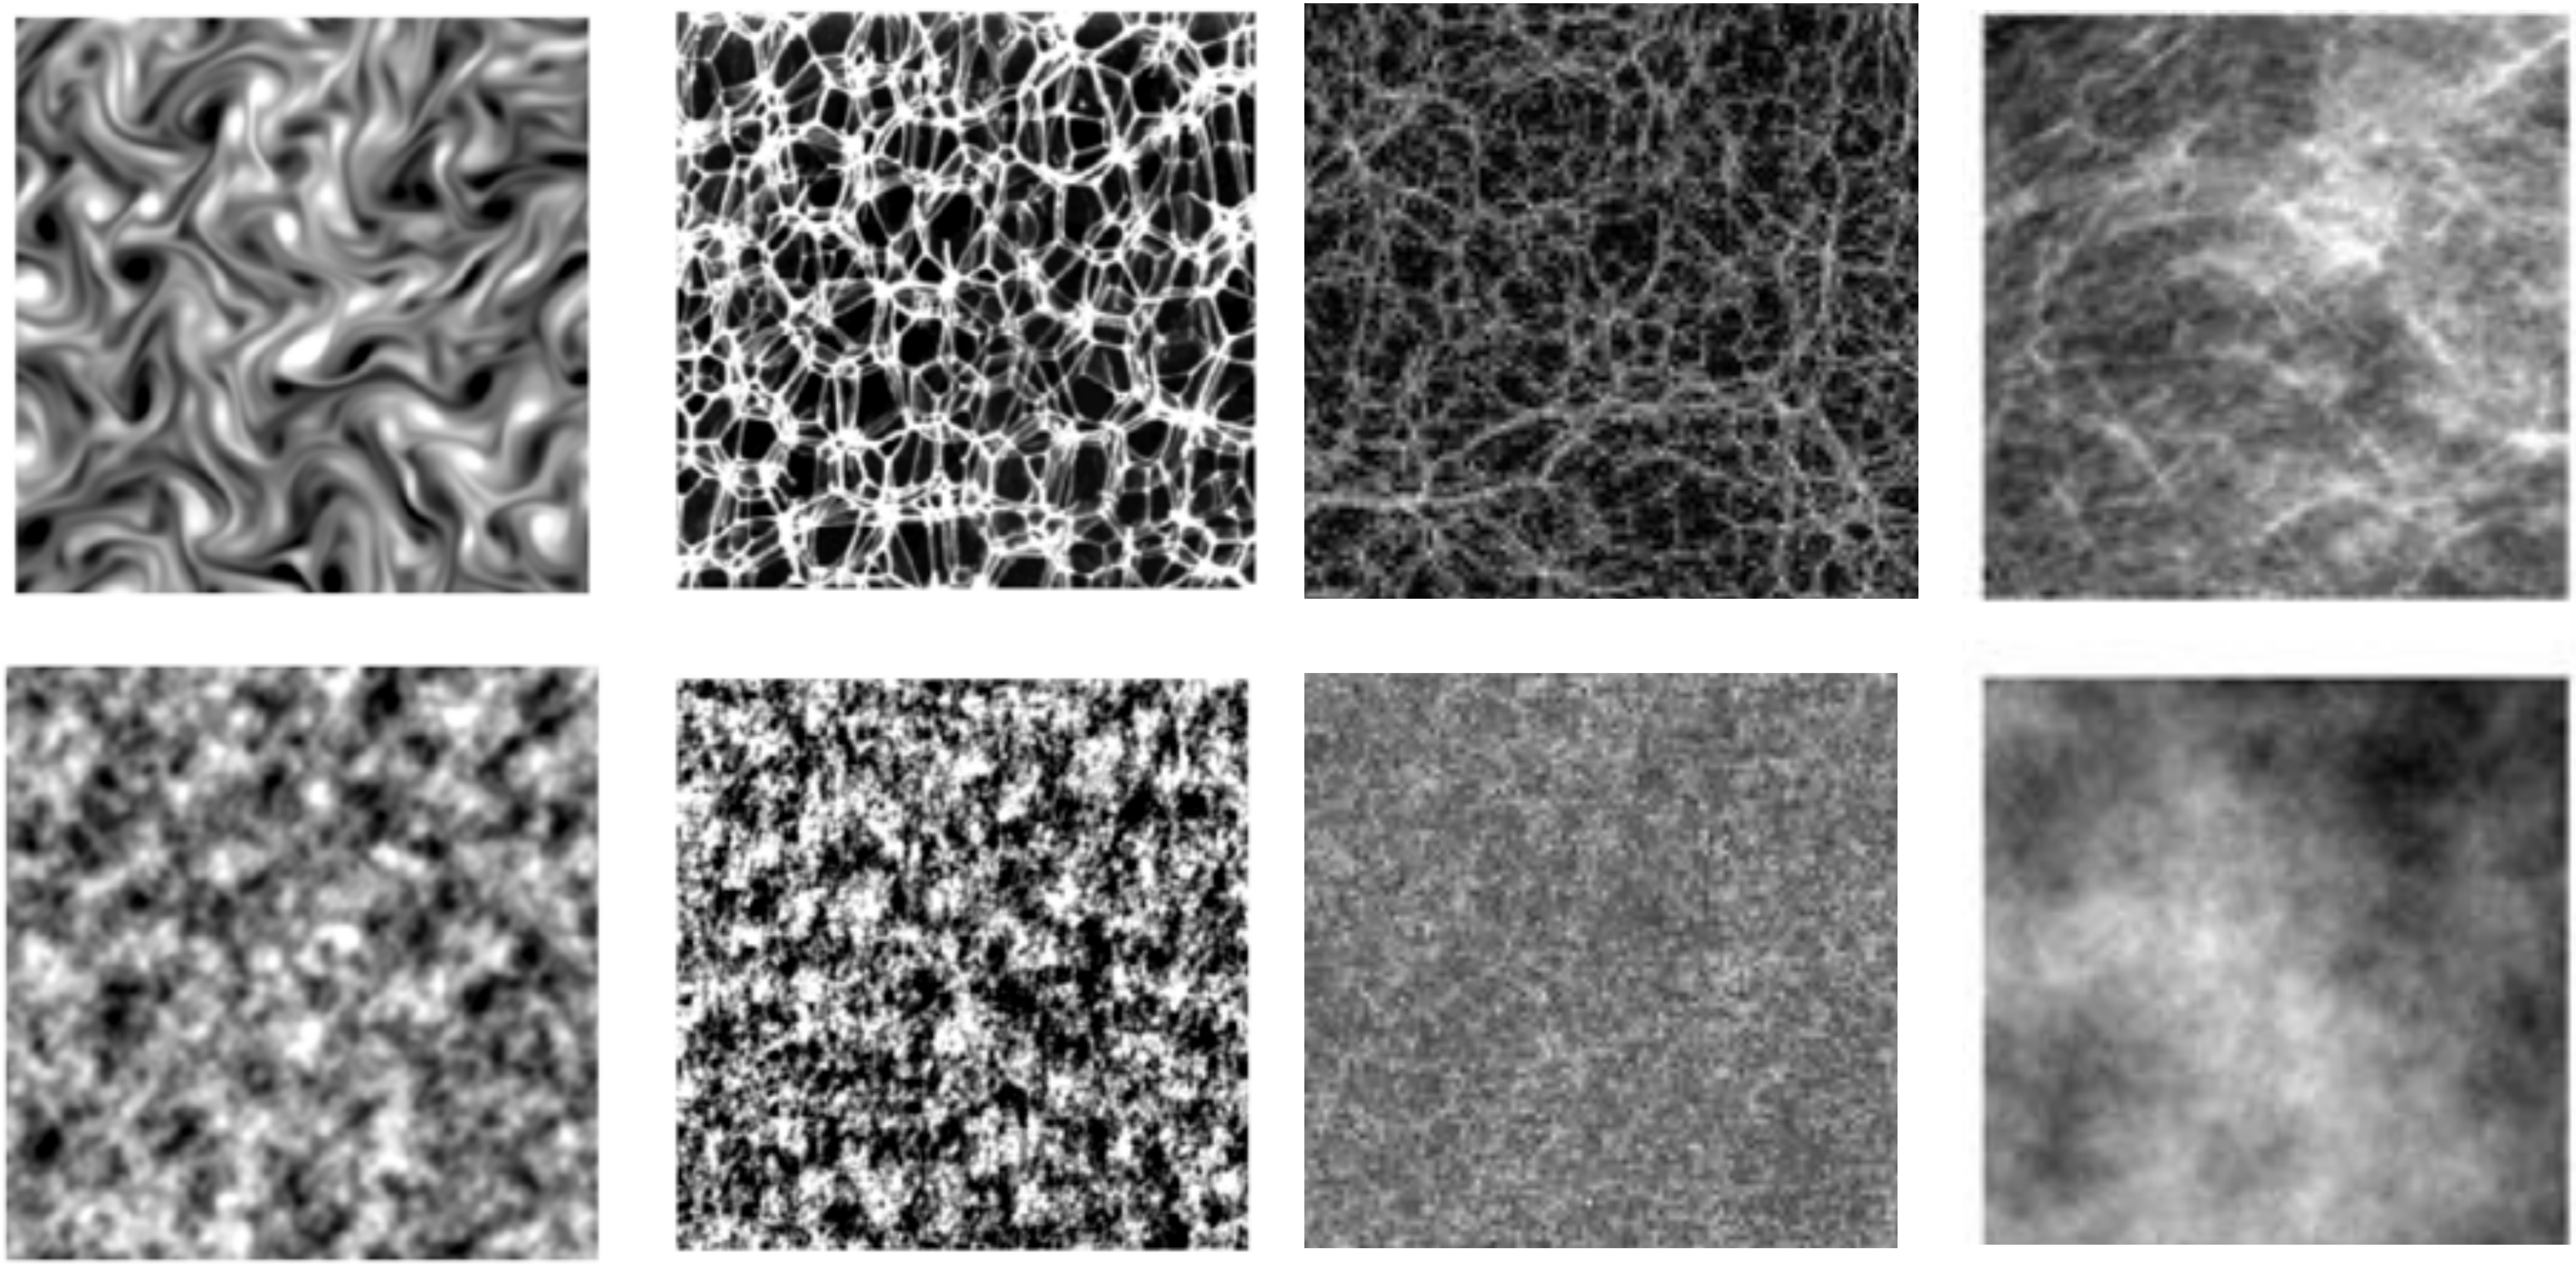

Higher-order moments?

Another approach: the Fourier transform

It is invariant to translations. $F_n(x) = \sum_{k=-n}^{n} c_n e^{ikx}$ with $c_n = \frac{1}{2\pi}\int_0^{2\pi} f(x)e^{-ikx}\mathrm{d}x$

- Translational invariance: if $x_c(t) = x(t-c)$ then $\left|\hat{x}_c(\omega)\right| = |\hat{x}(\omega)|$

- Expanding the finite series adds more information.

- Robust to small distortions (i.e. warping) ?

What happens in Fourier space?

original object: $x(t)$

warped object: $x_\tau(t) = x(t-\tau(t))$

$|\tilde{x}|$

frequency

stable unstable

the scattering transform

Stéphane Mallat & Joan Bruna (2012)

The Fourier transform is too delocalized → Let's use a wavelet ψ. $x \star \psi$

A wavelet is not translational invariant. → Let's use some averaging ϕ. $x \star \psi \star \phi$

Averaging will tend to give zero → Let's use an absolute value. $|x \star \psi| \star \phi$

original object
warped object

$|\tilde{x}|$

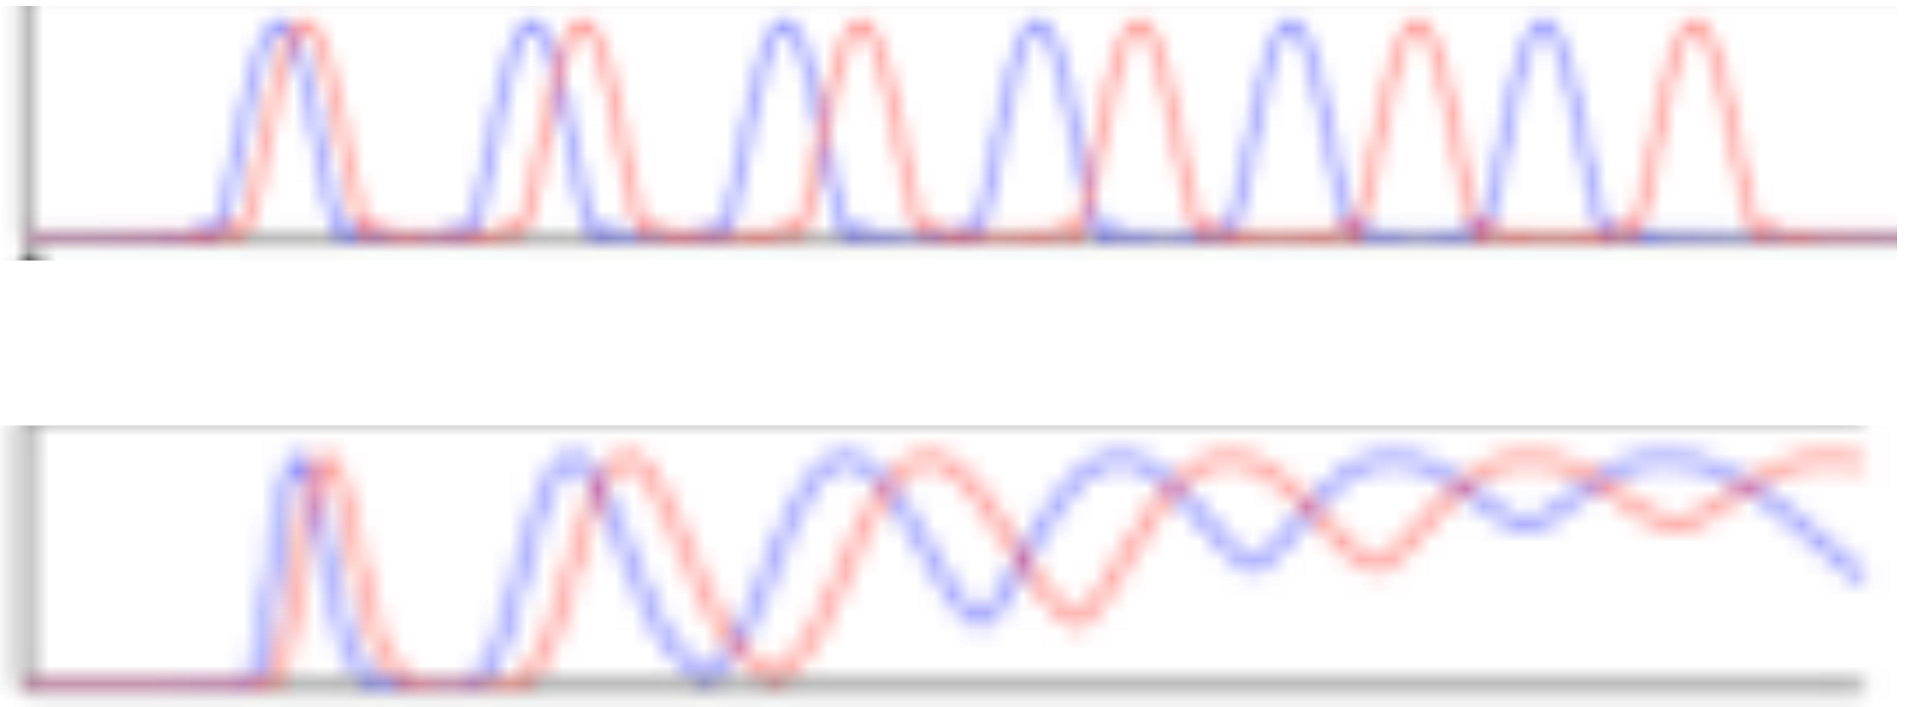

frequency

frequency

We now are invariant to small translations & deformations.
We can do this as a function of scale.

$$\left\{ \left| x \star \psi_\lambda \right| \star \phi, \right\}_\lambda$$

BUT… we have lost some information by taking an absolute value…

SOLUTION: do it hierarchically.

$$\left\{ \left| \left| x \star \psi_{\lambda_1} \right| \star \psi_{\lambda_2} \right| \ldots \star \phi, \right\}_{\lambda_1, \lambda_2}$$

A cascade of convolutions

$$\left\{ \left| \left| \left| x \star \psi_{\lambda_1} \right| \star \psi_{\lambda_2} \right| \dots \star \psi_{\lambda_i} \right| \star \phi, \right\}_{\lambda_1,\dots,\lambda_i}$$

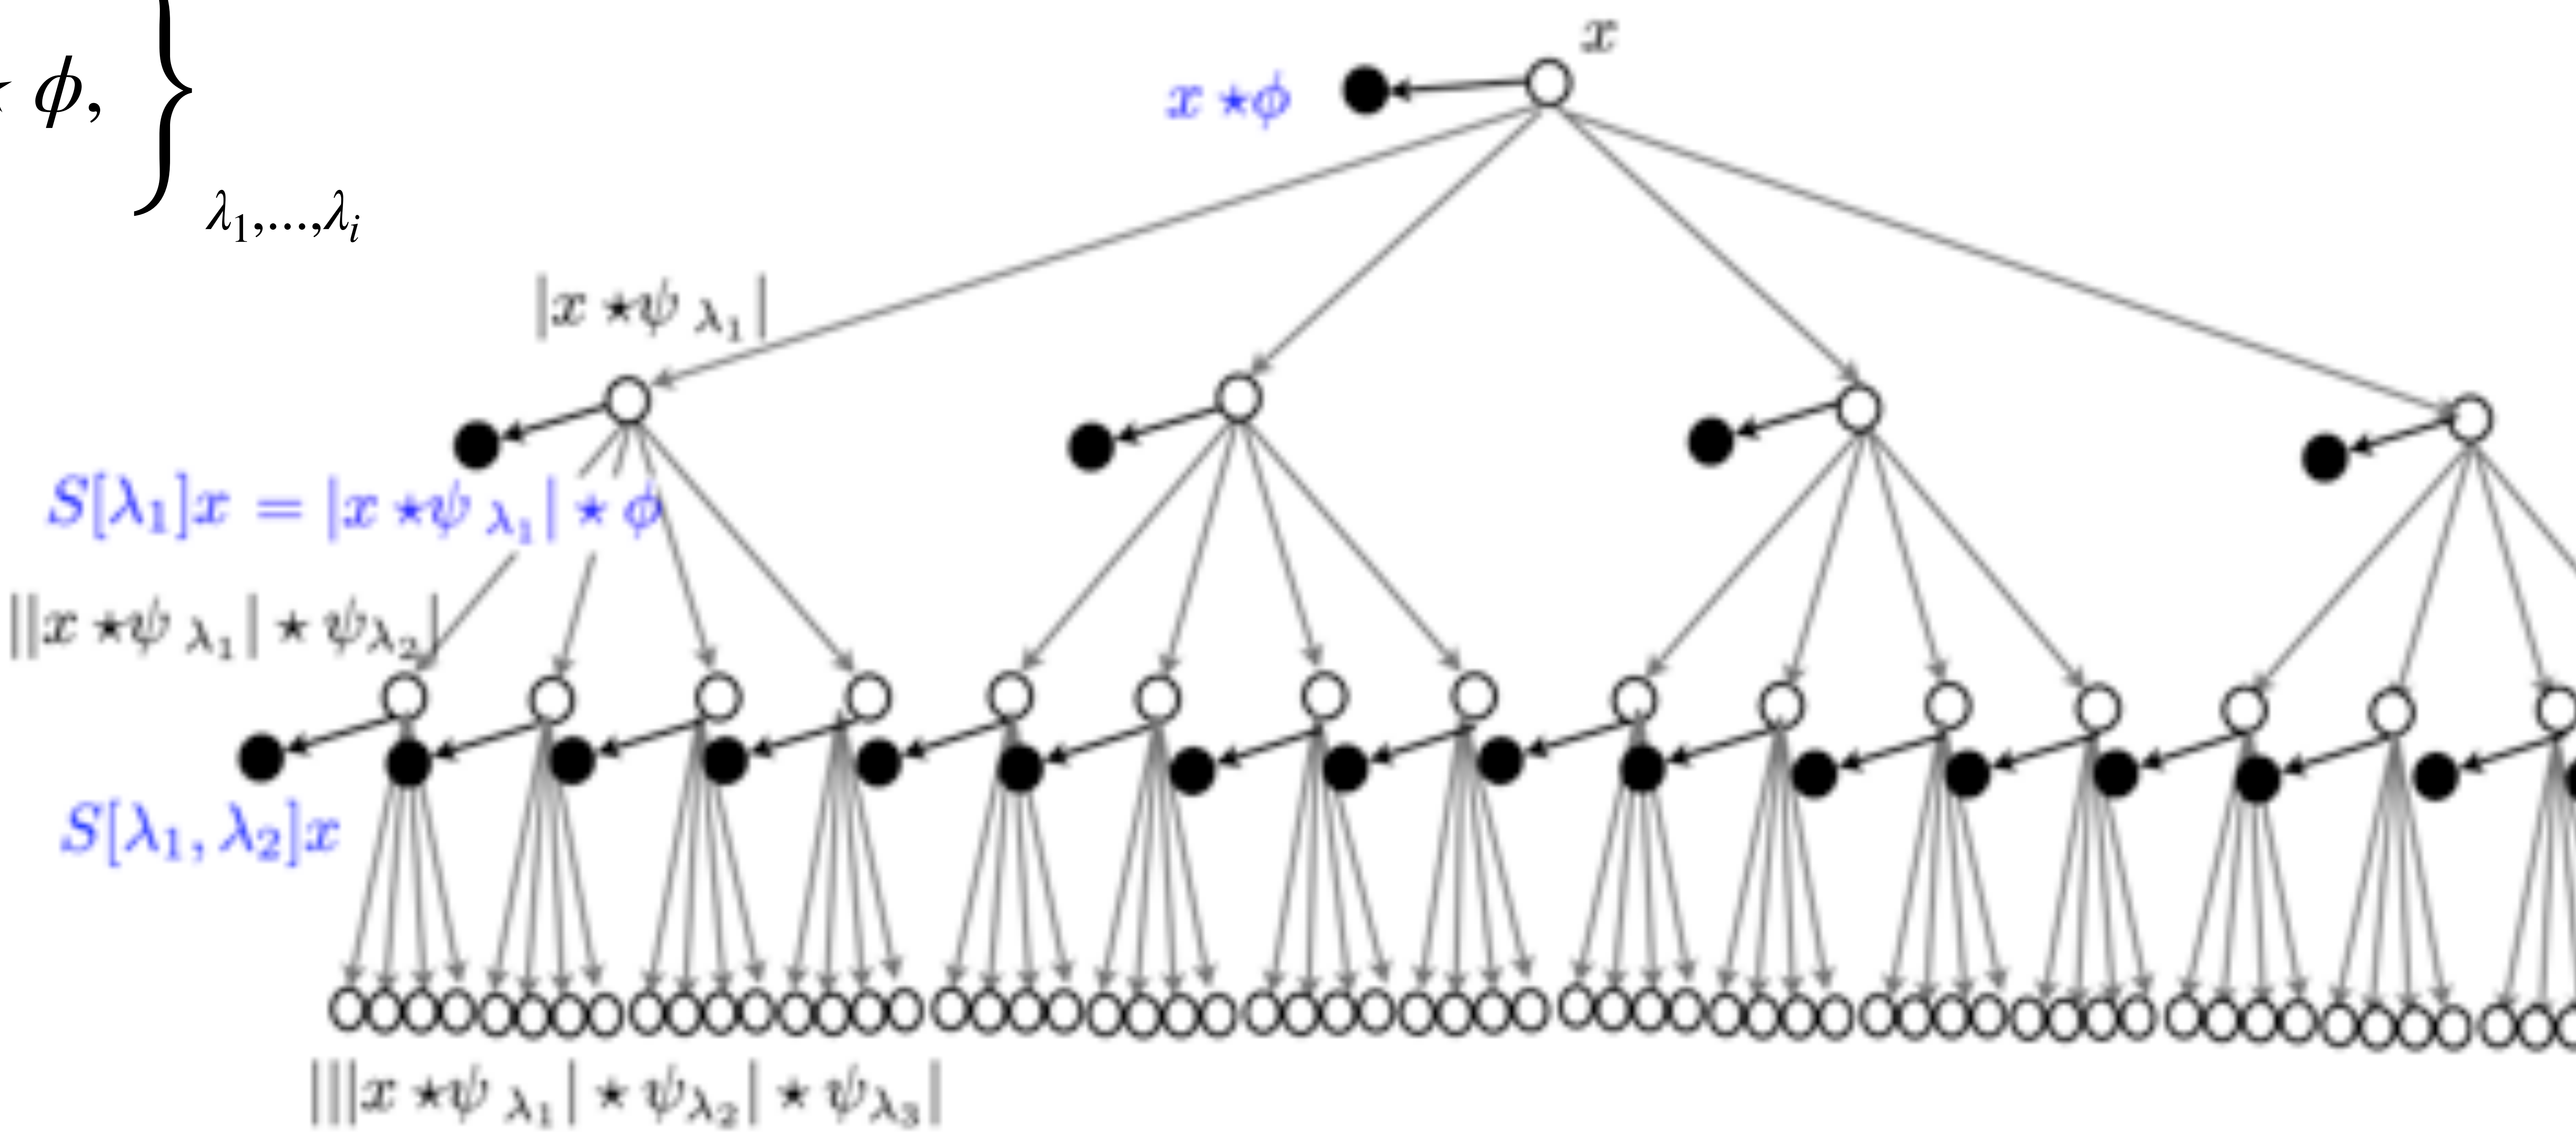

Data $x \in \mathbb{R}^d$

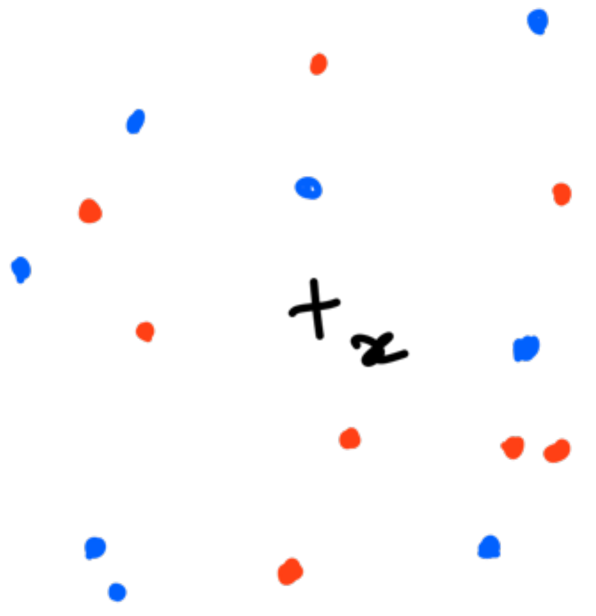

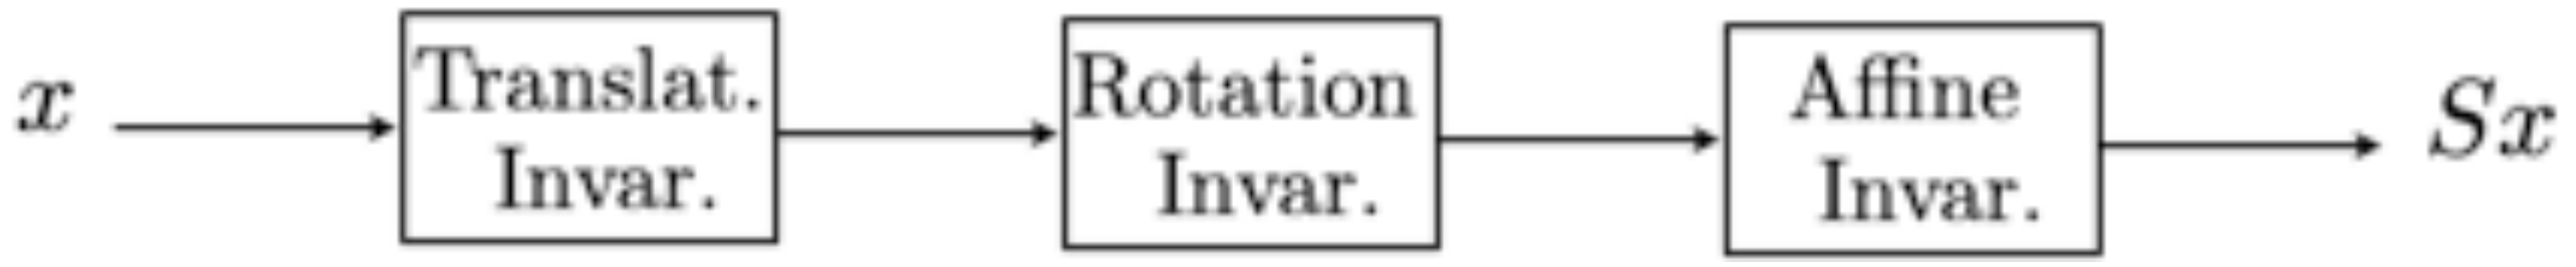

Representation

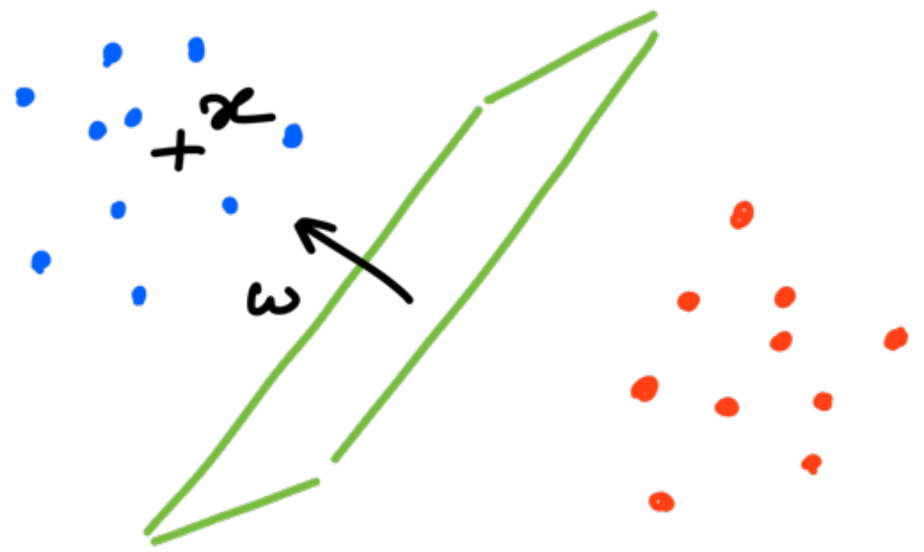

J. Bruna & S. Mallat

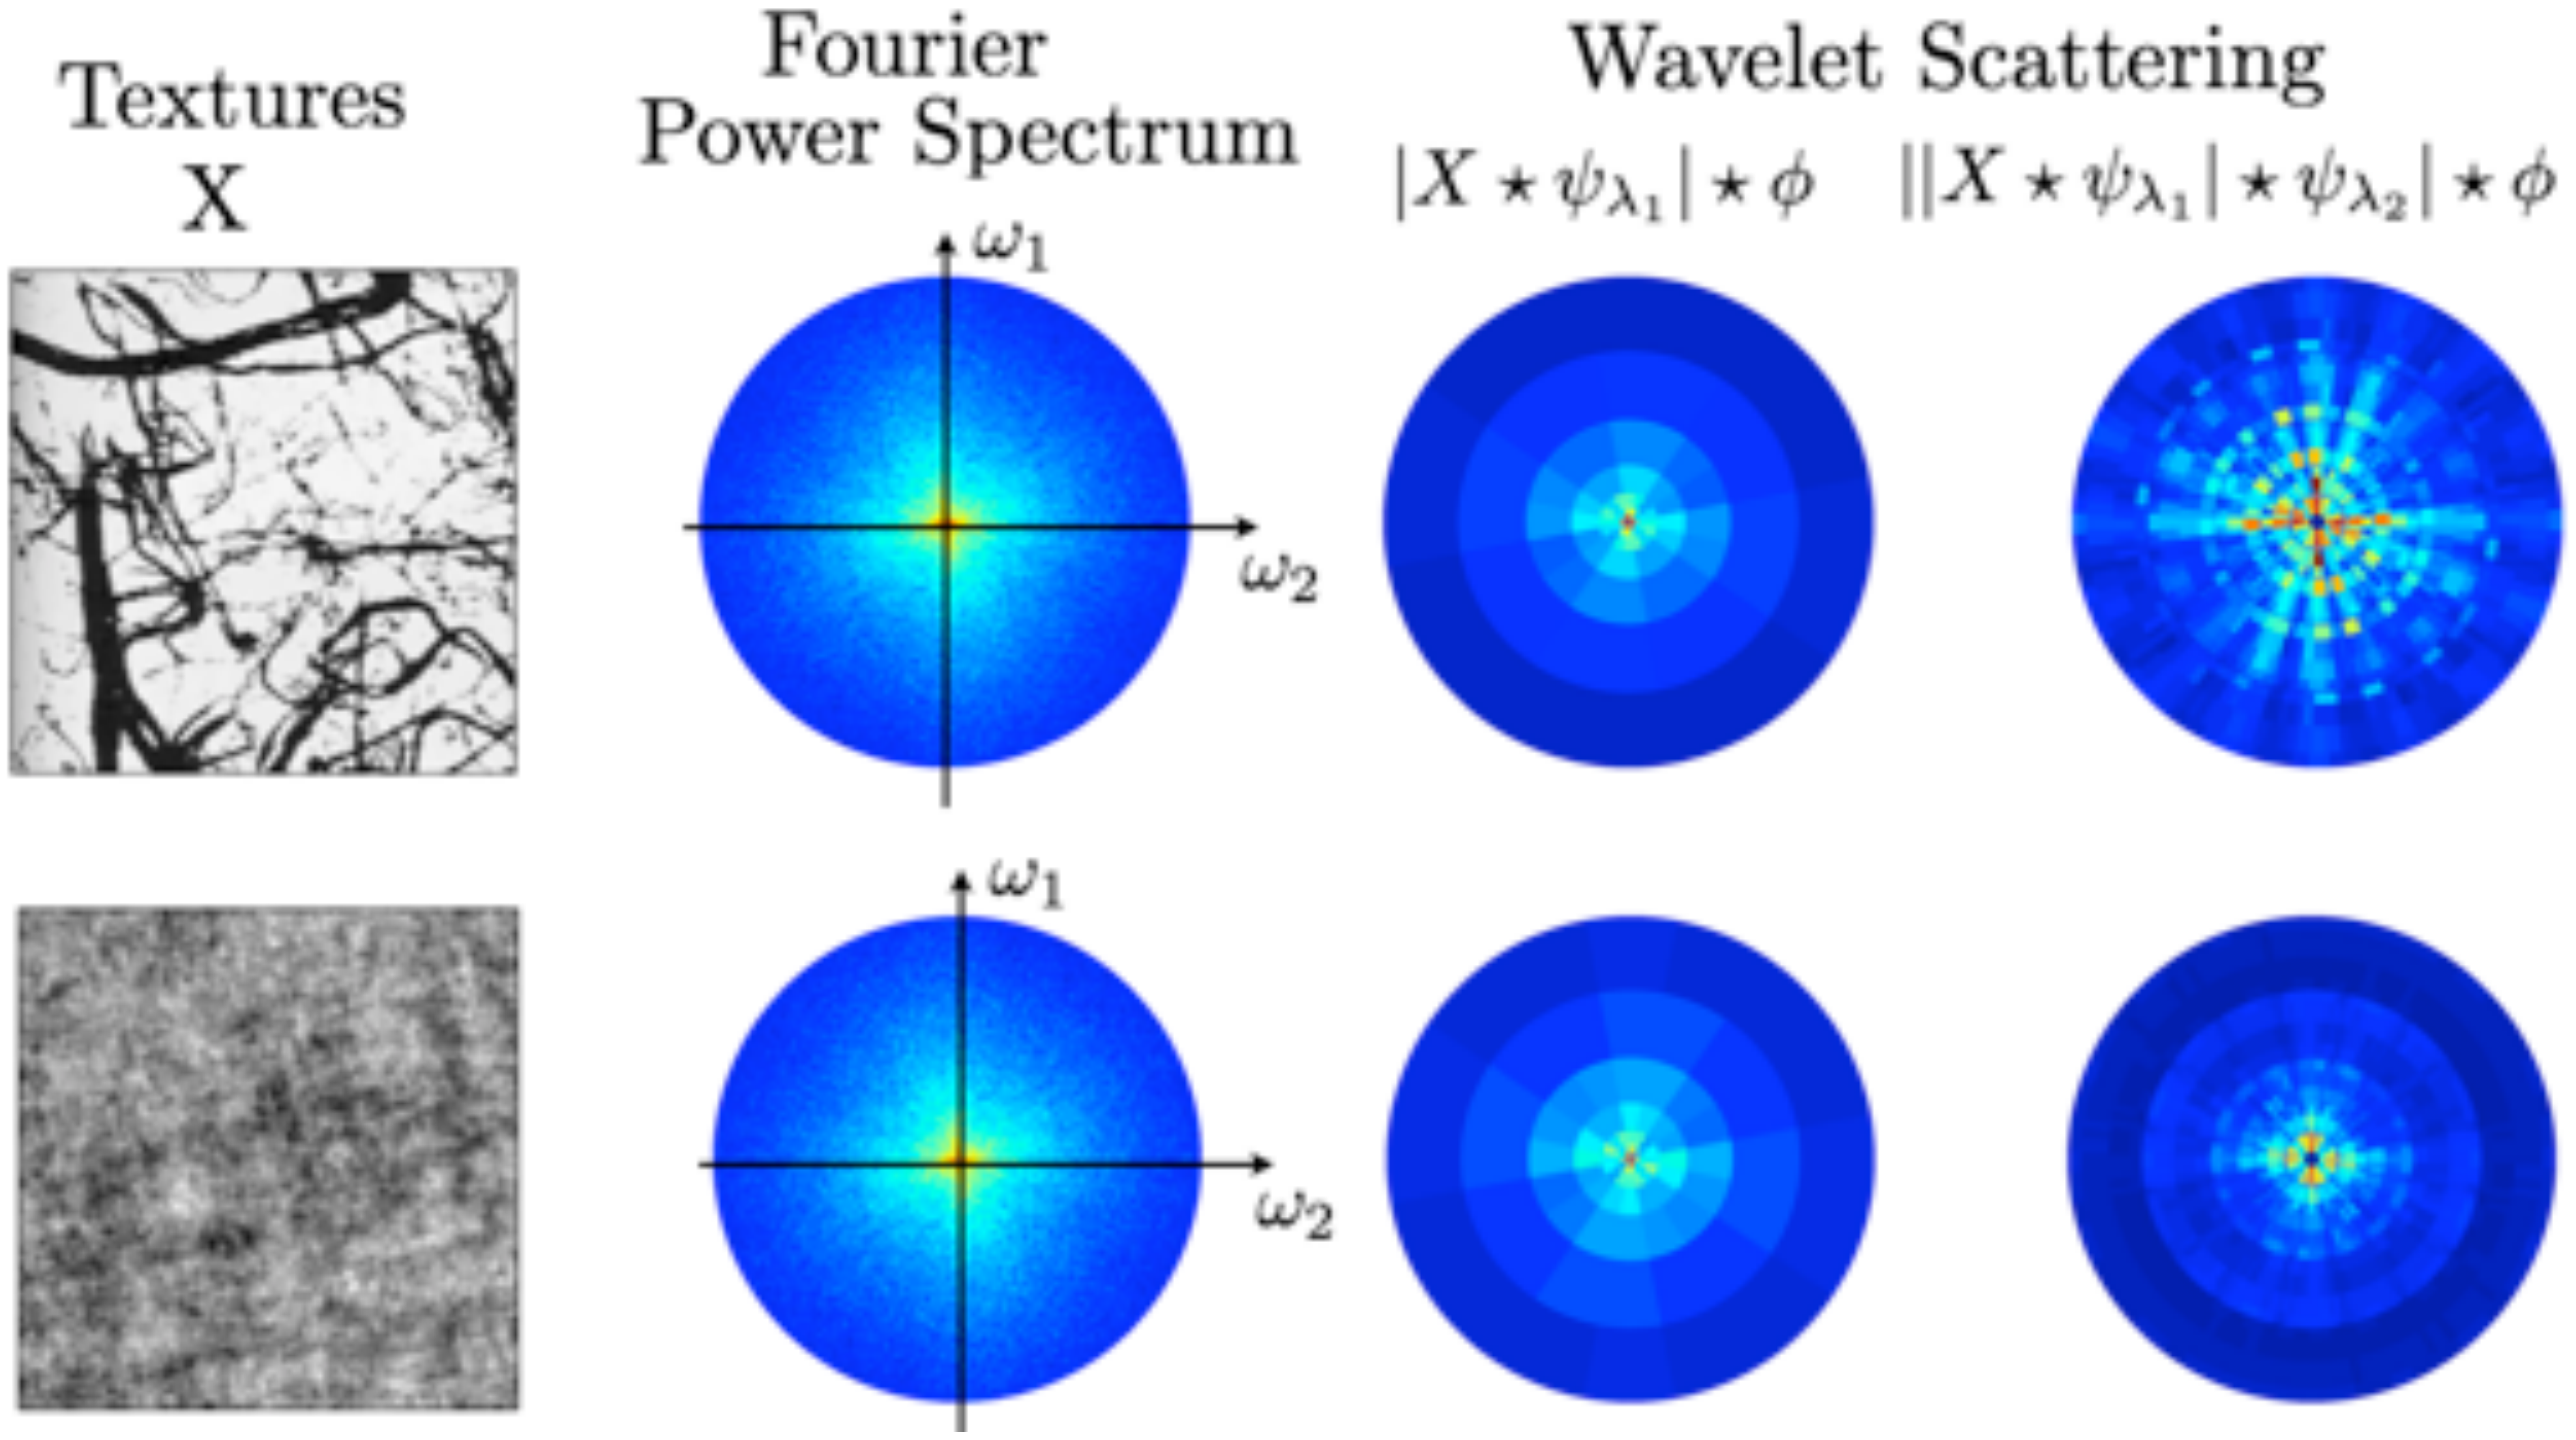

image from J. Bruna & S. Mallat

Texture classification

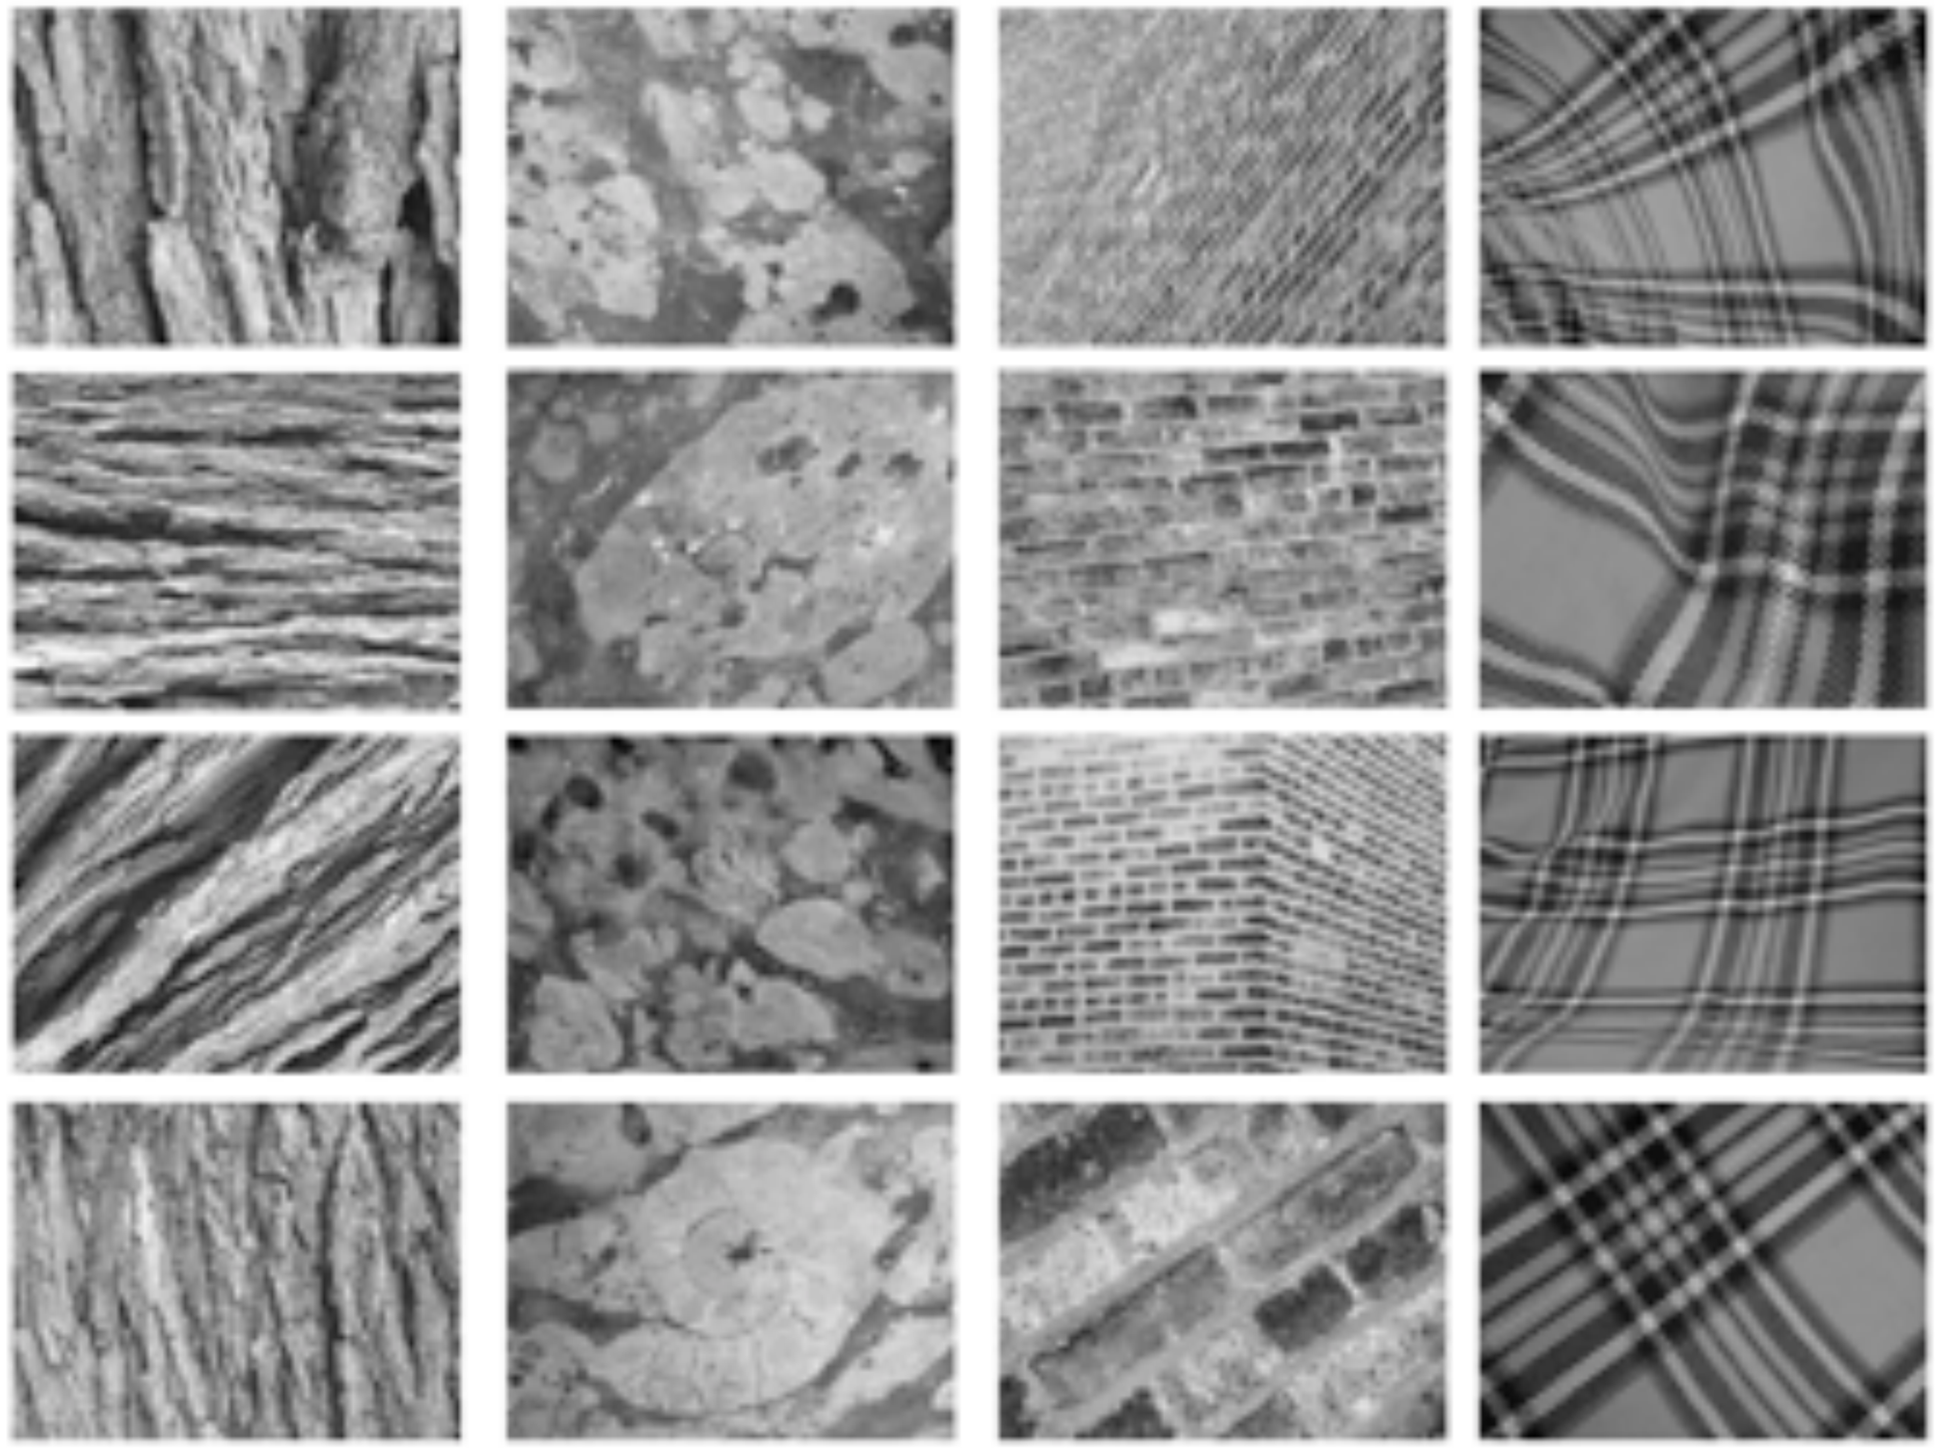

$$||X \star \psi_{\lambda_1}|| \star \phi \qquad\qquad ||X \star \psi_{\lambda_1}| \star \psi_{\lambda_2}| \star \phi$$

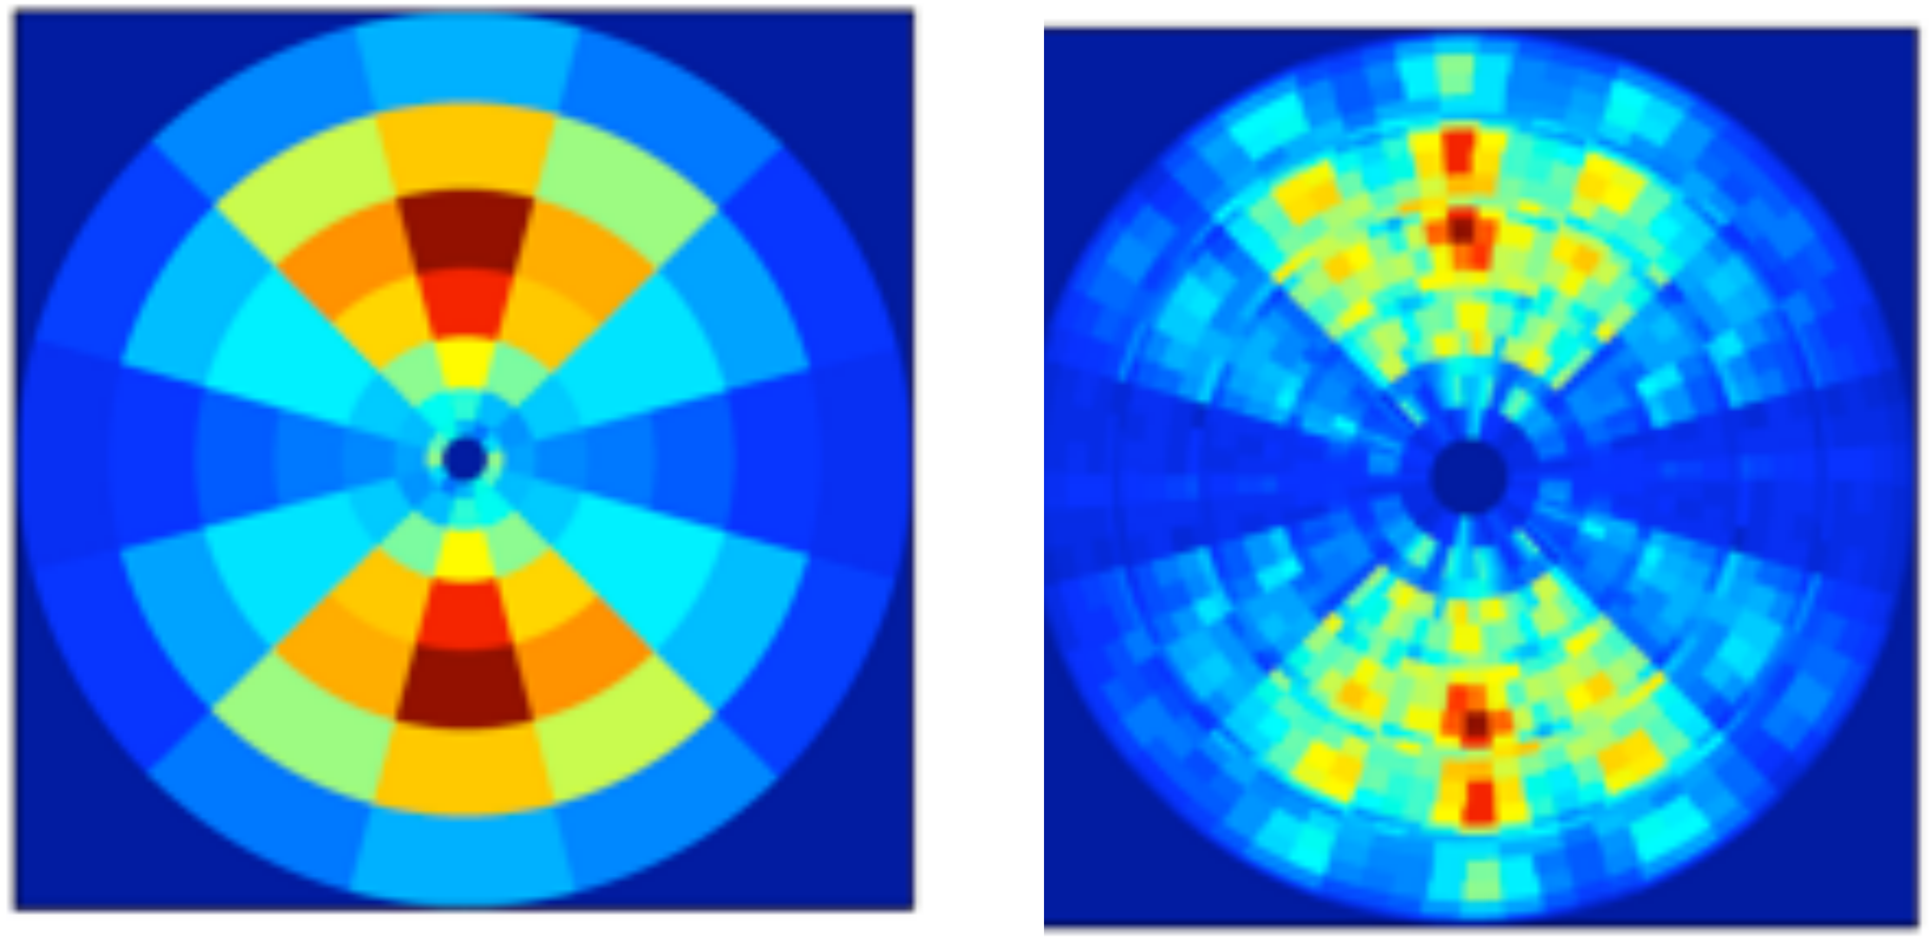

Data $x \in \mathbb{R}^d$

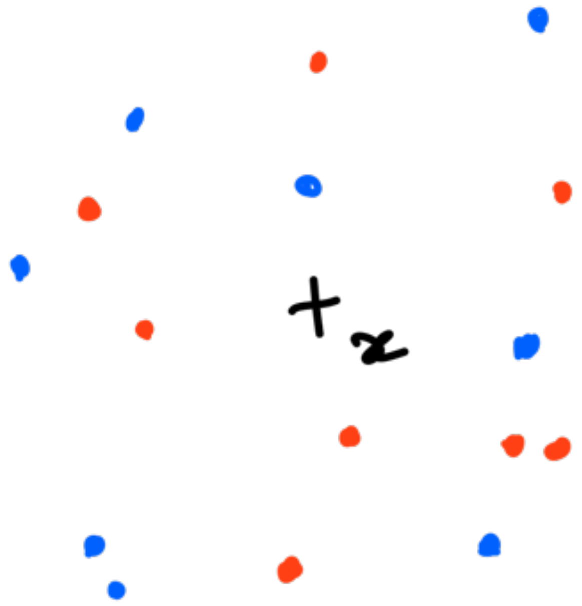

$$x \longrightarrow \boxed{\text{Translat. Invar.}} \longrightarrow \boxed{\text{Rotation Invar.}} \longrightarrow \boxed{\text{Affine Invar.}} \longrightarrow Sx$$

Representation

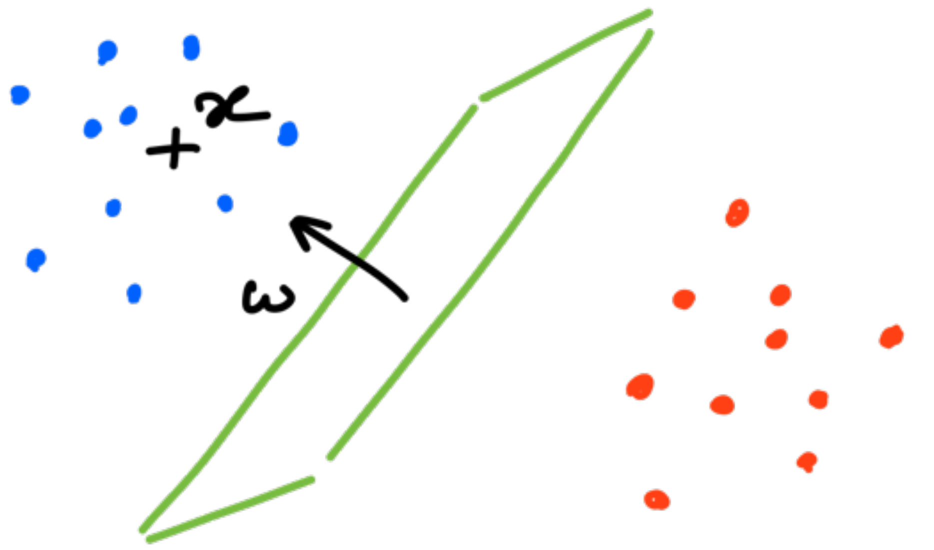

from Laurent Sifre & S. Mallat

Object classification

To construct a representation $\Phi(f)$ useful for classification, we need invariants:

- translations
- rotations
- scaling
- small deformations
- change in brightness / illumination
- additive constant

the scattering transform (Mallat & Bruna 2012) can achieve this with the following requirements:

- filters must be wavelets (convolution)
- complex modulus (non-linearity)
- averaging (pooling)
- hierarchical (multilayer)

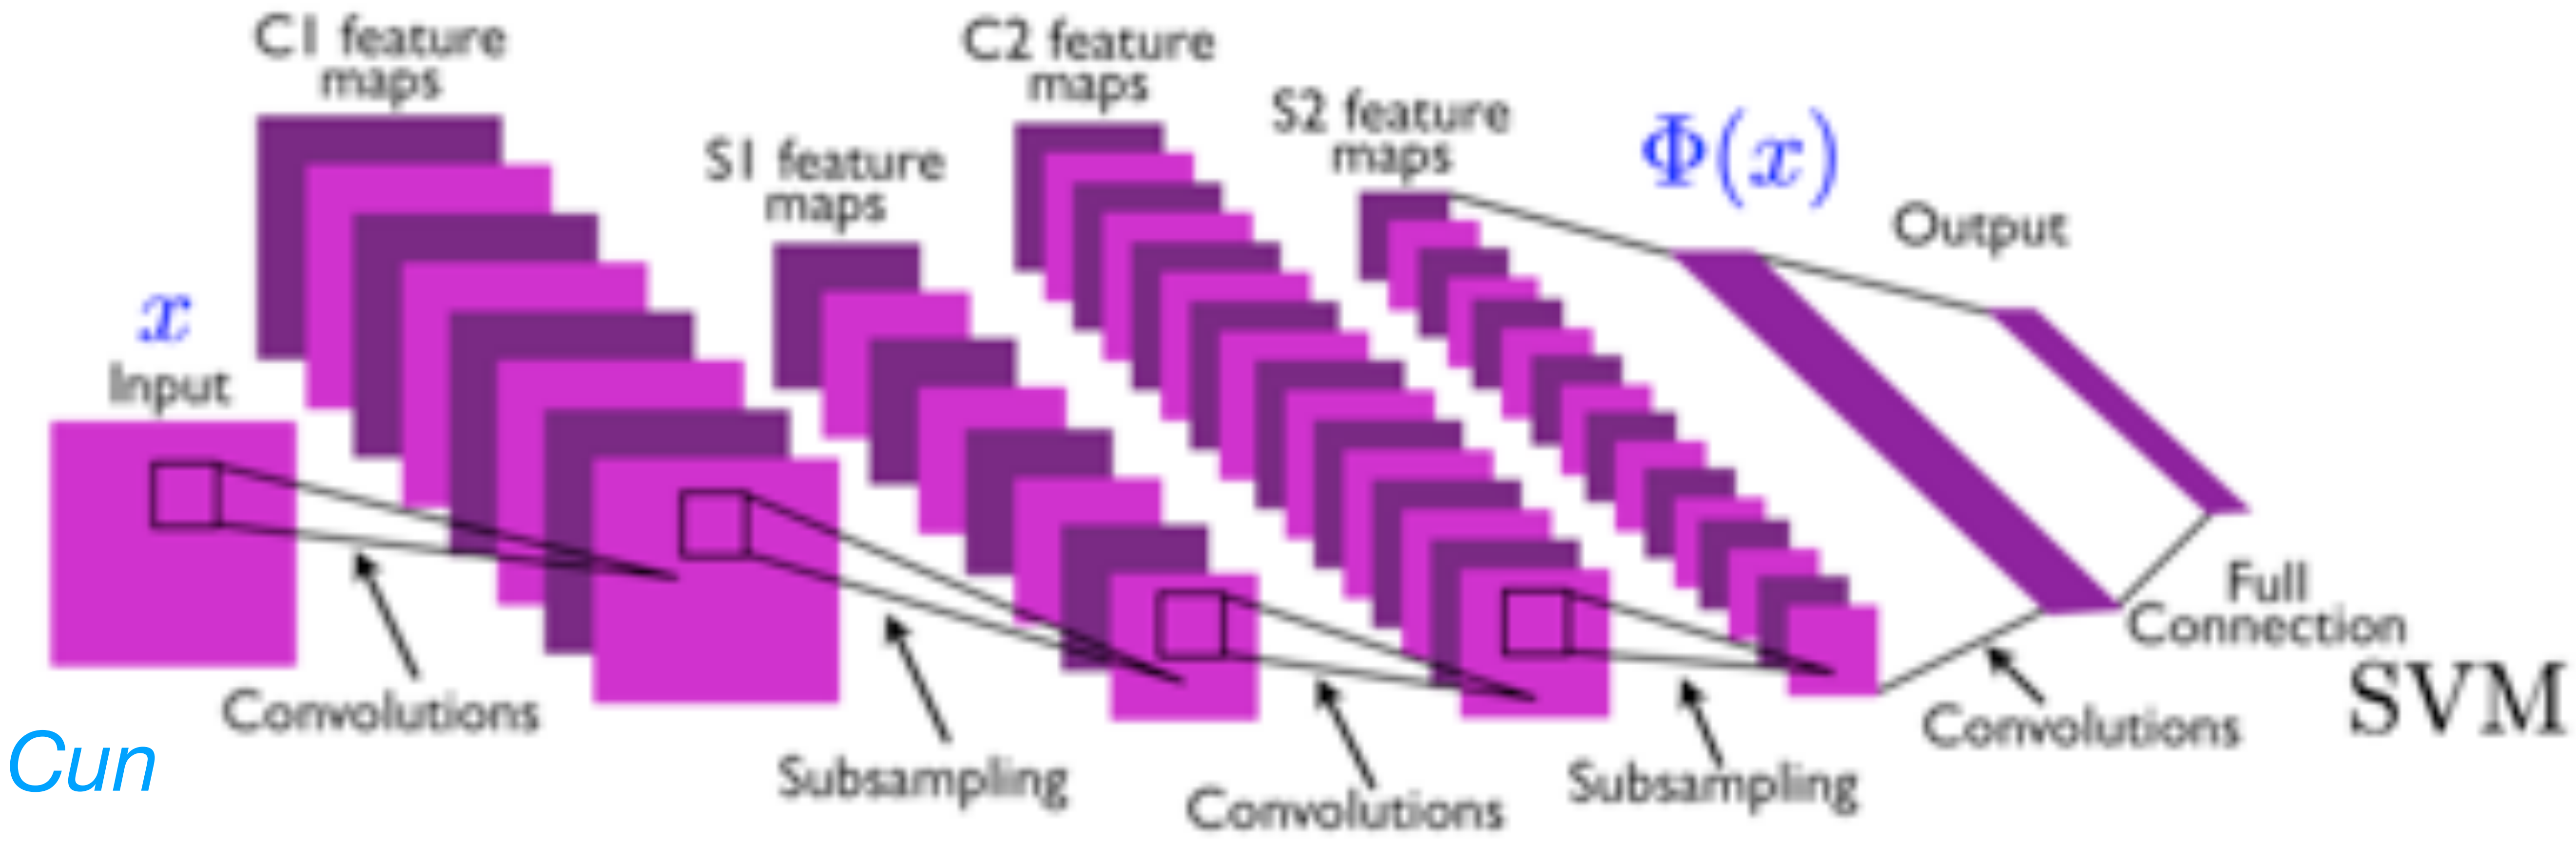

Yann Le Cun

Gaussian statistics, PCA | manifold learning | **dictionary learning** | (convolutional) neural nets

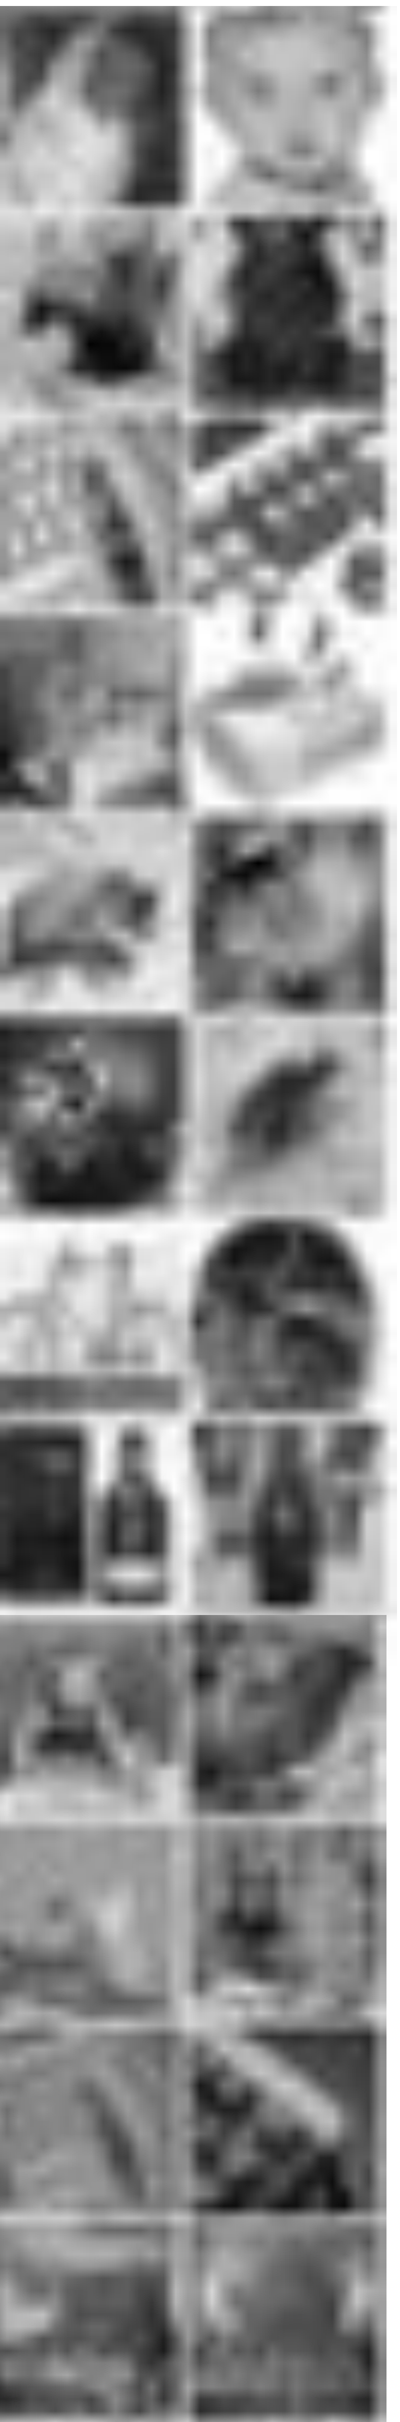

Principal Component Analysis

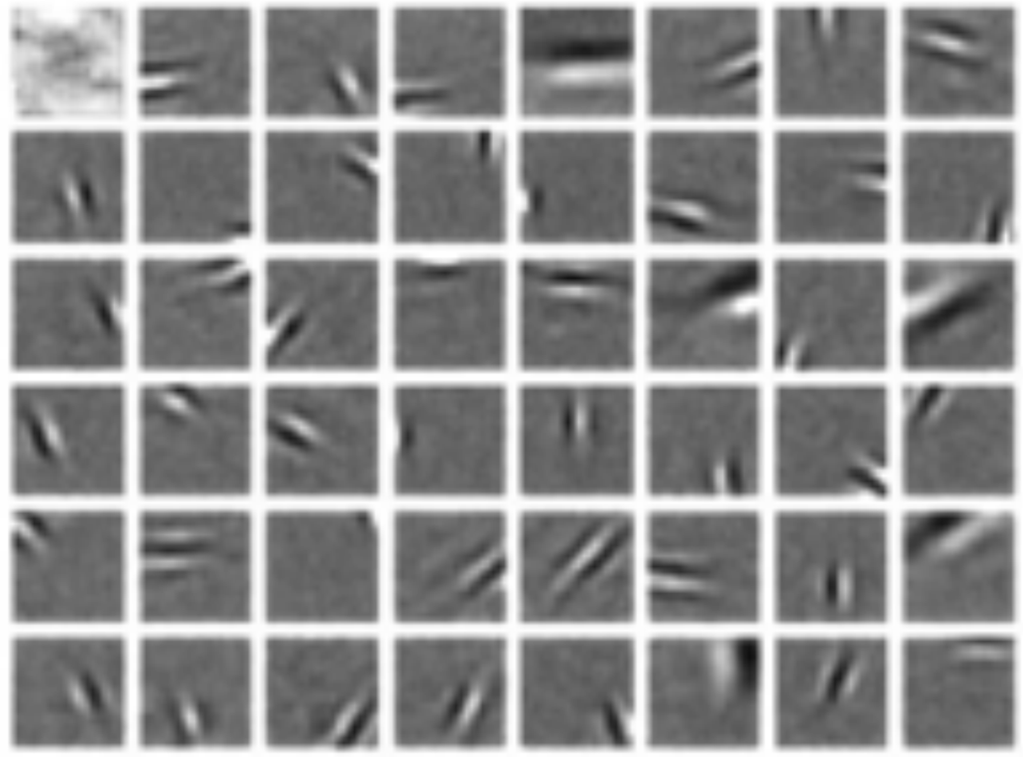

Emergence of simple-cell receptive field properties by learning a sparse code for natural images

Bruno A. Olshausen* & David J. Field

Department of Psychology, Uris Hall, Cornell University, Ithaca, New York 14853, USA

kernels learned in the first layer of a deep convolutional neural net with ImageNet

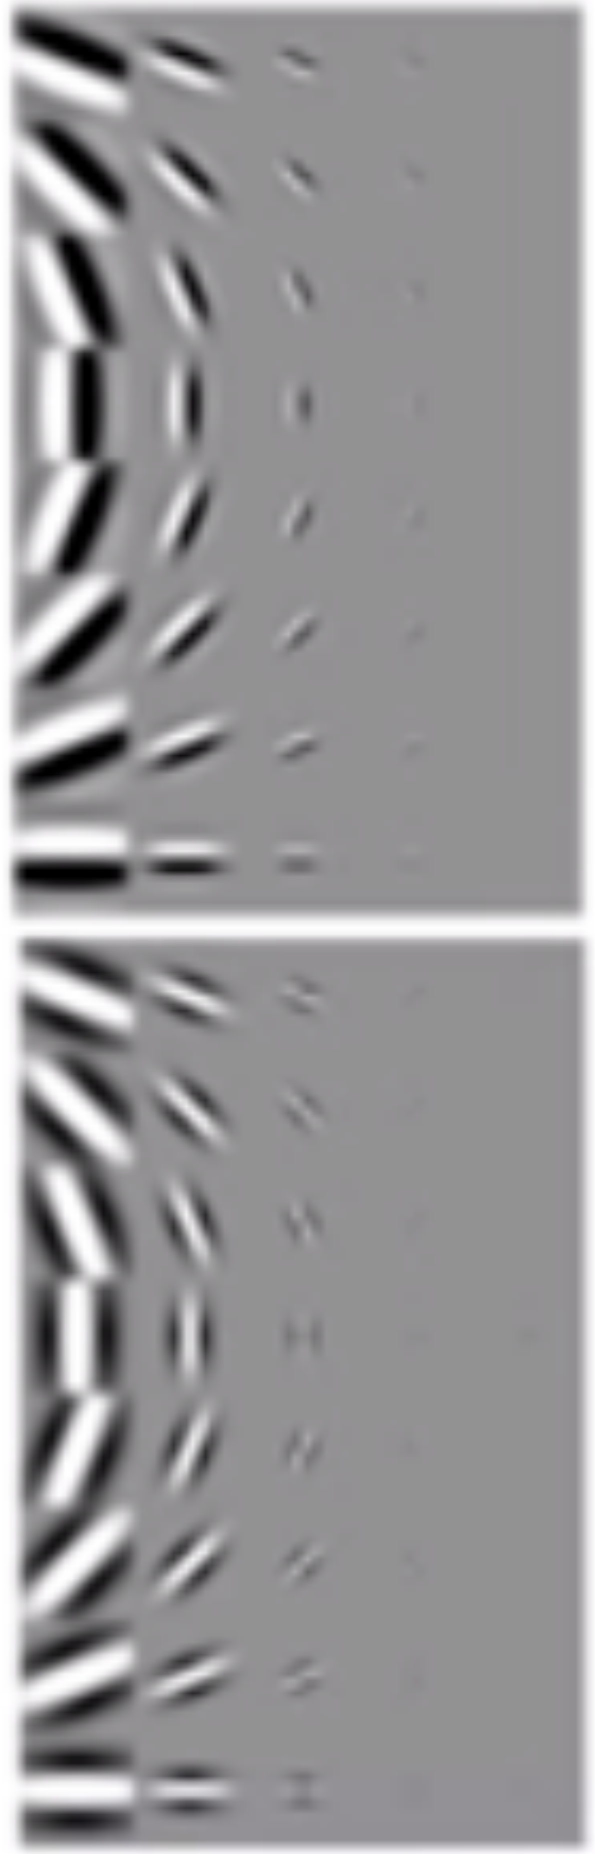

wavelets (Gabor filters)

Application to cosmology: the texture of the density field

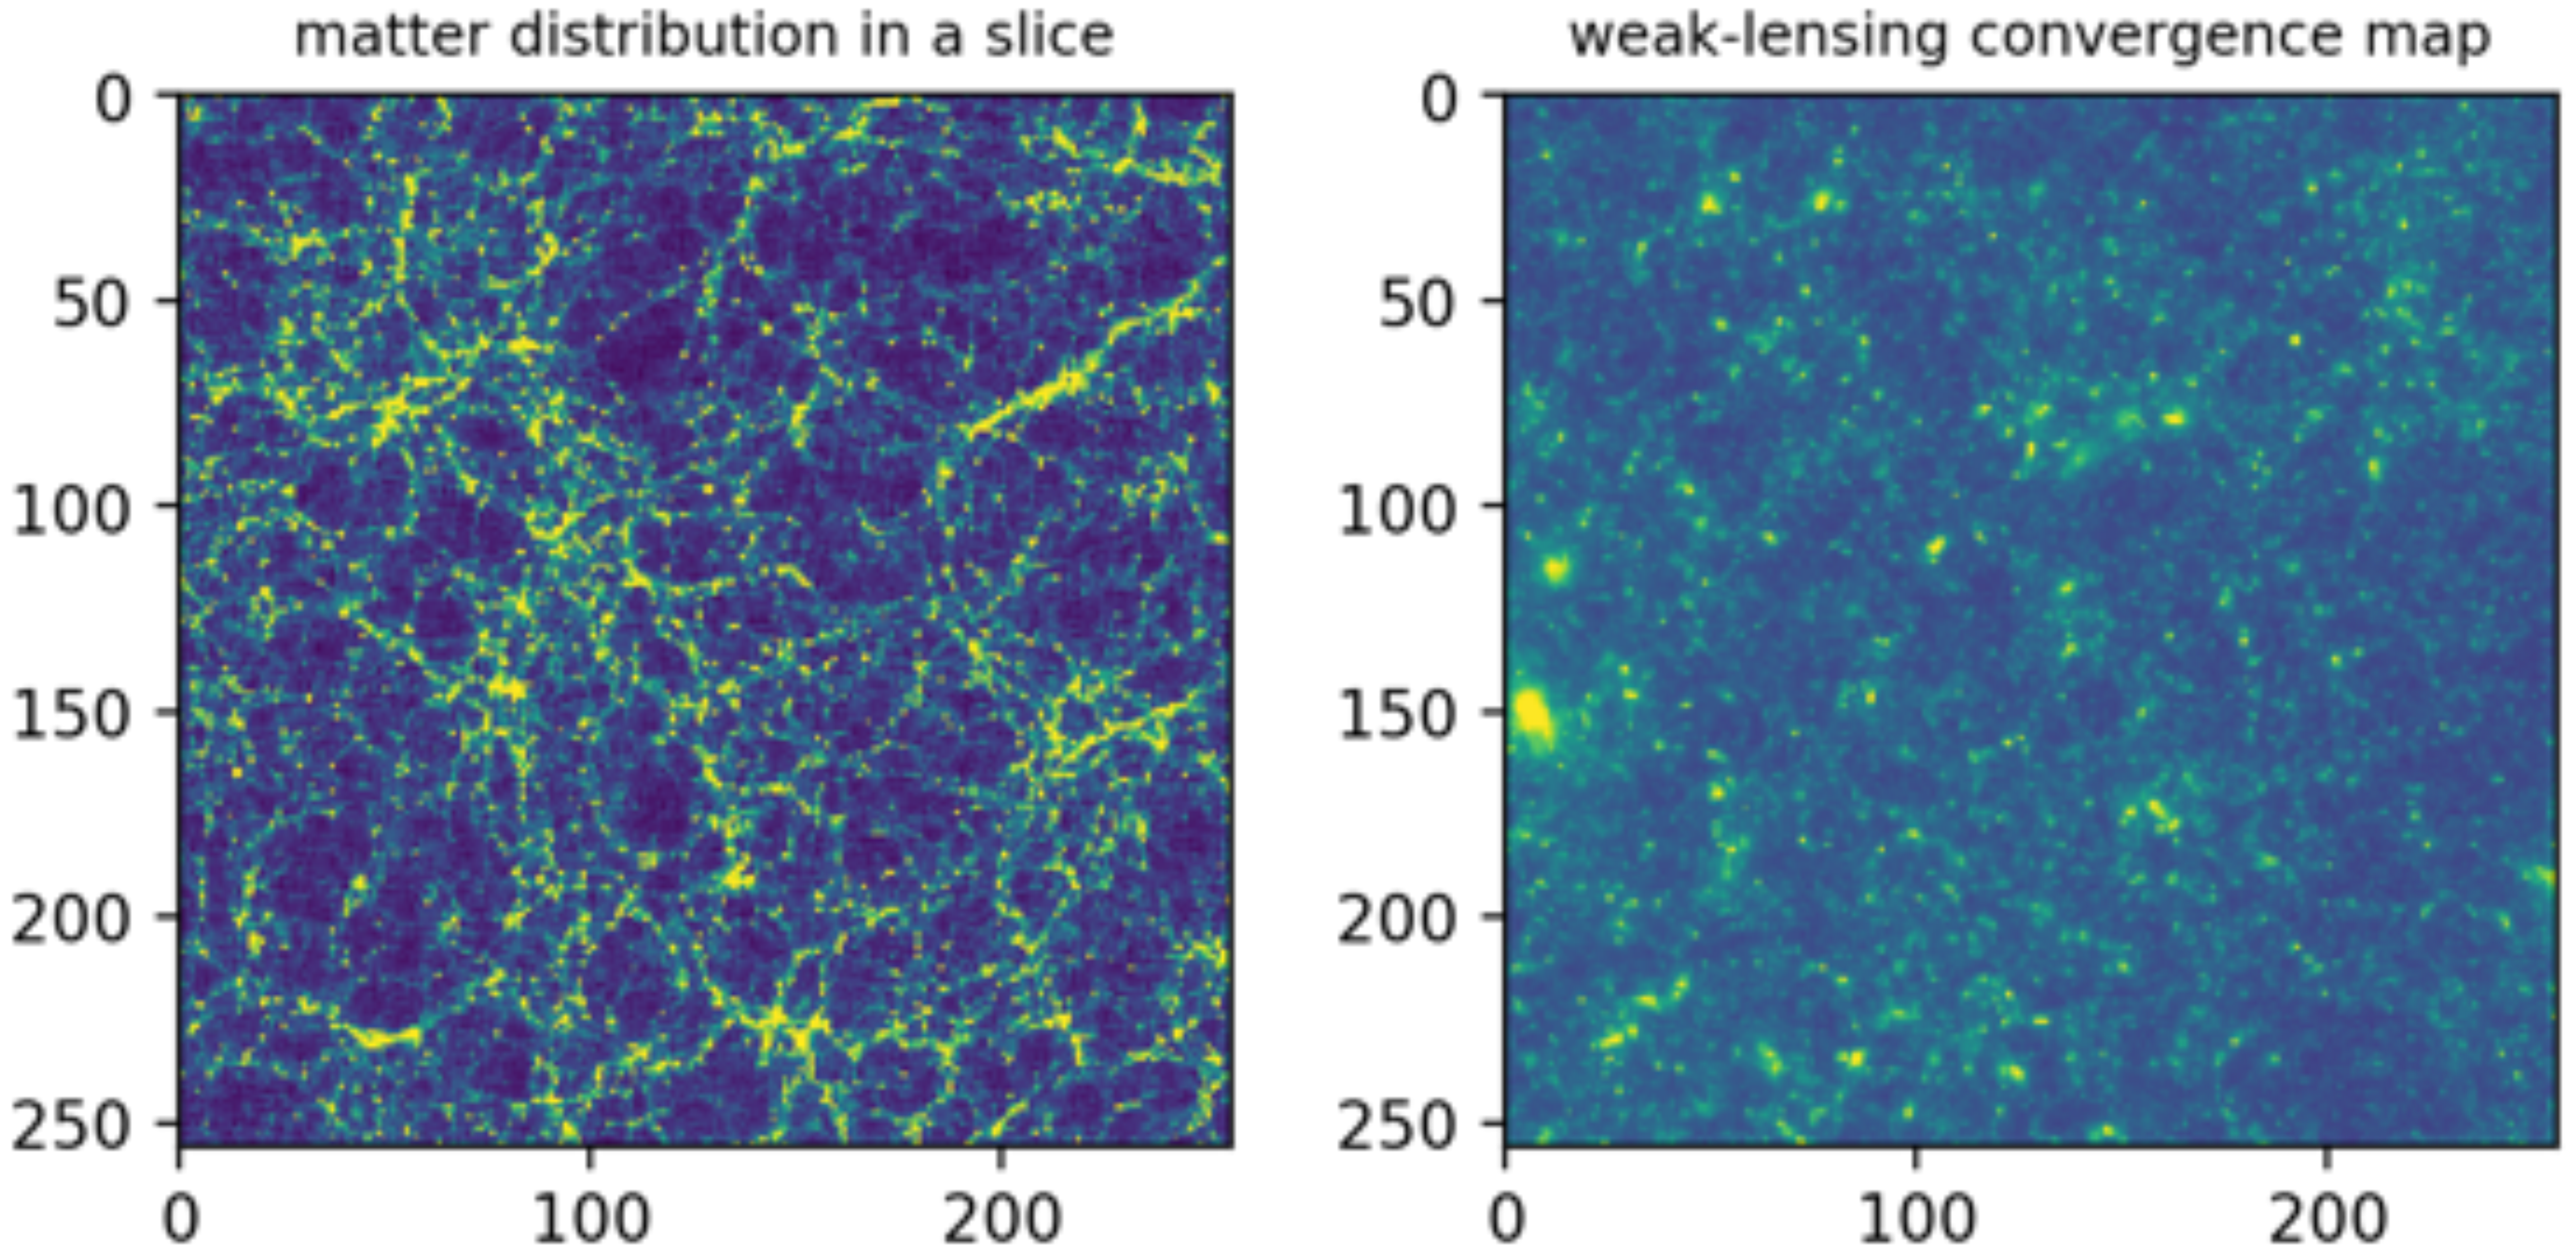

Gaussian statistics | manifold learning | scattering transform | (convolutional) neural nets

complexity

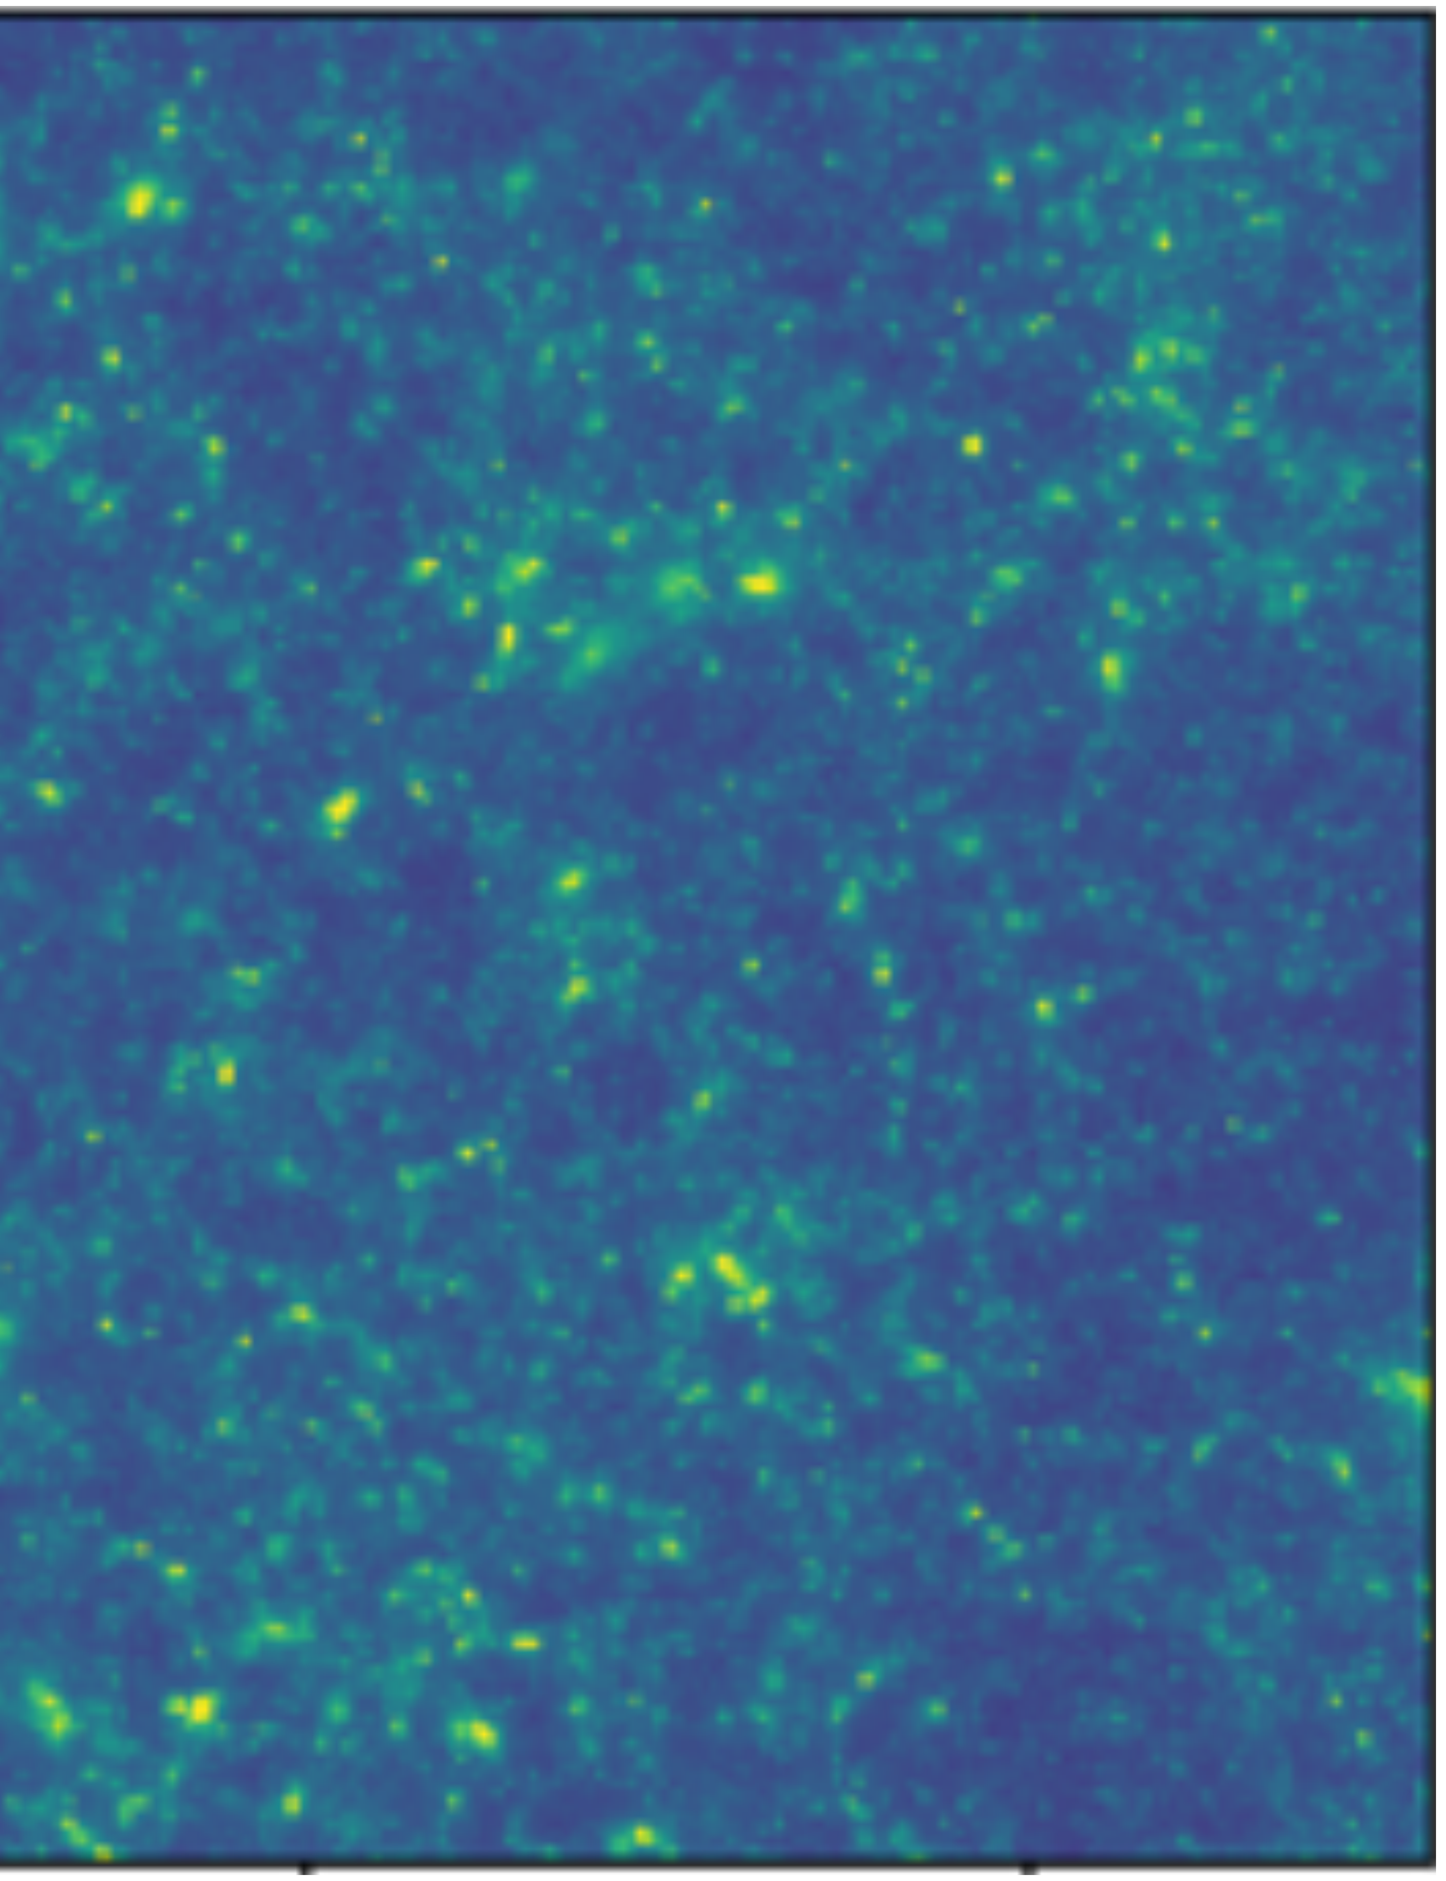

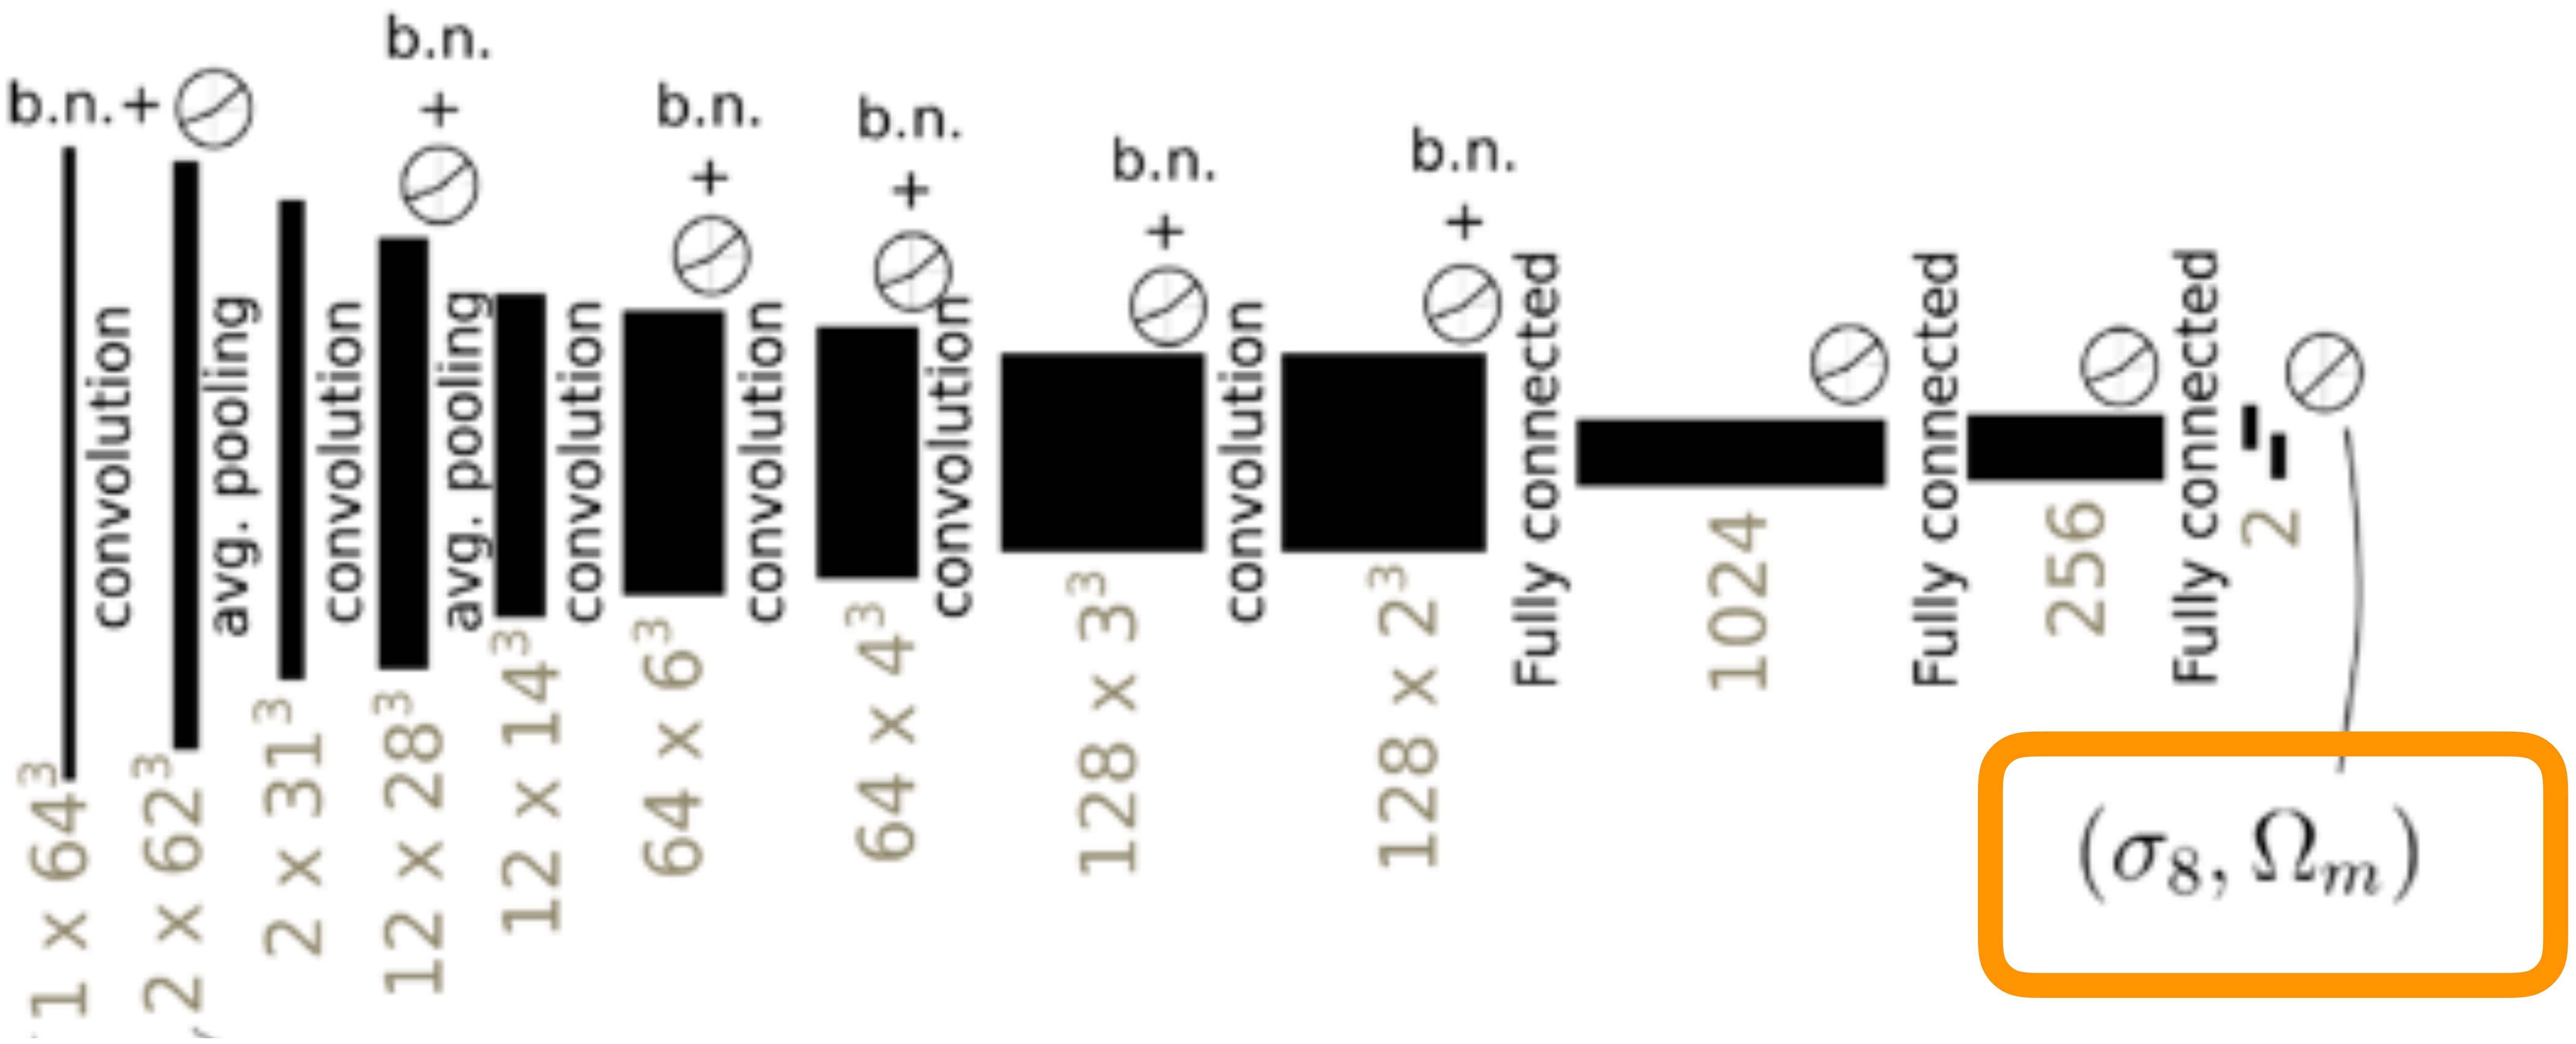

Ravanbakhsh et al. (2016)

Gaussian statistics, PCA | manifold learning | **scattering transform** | convolutional neural nets

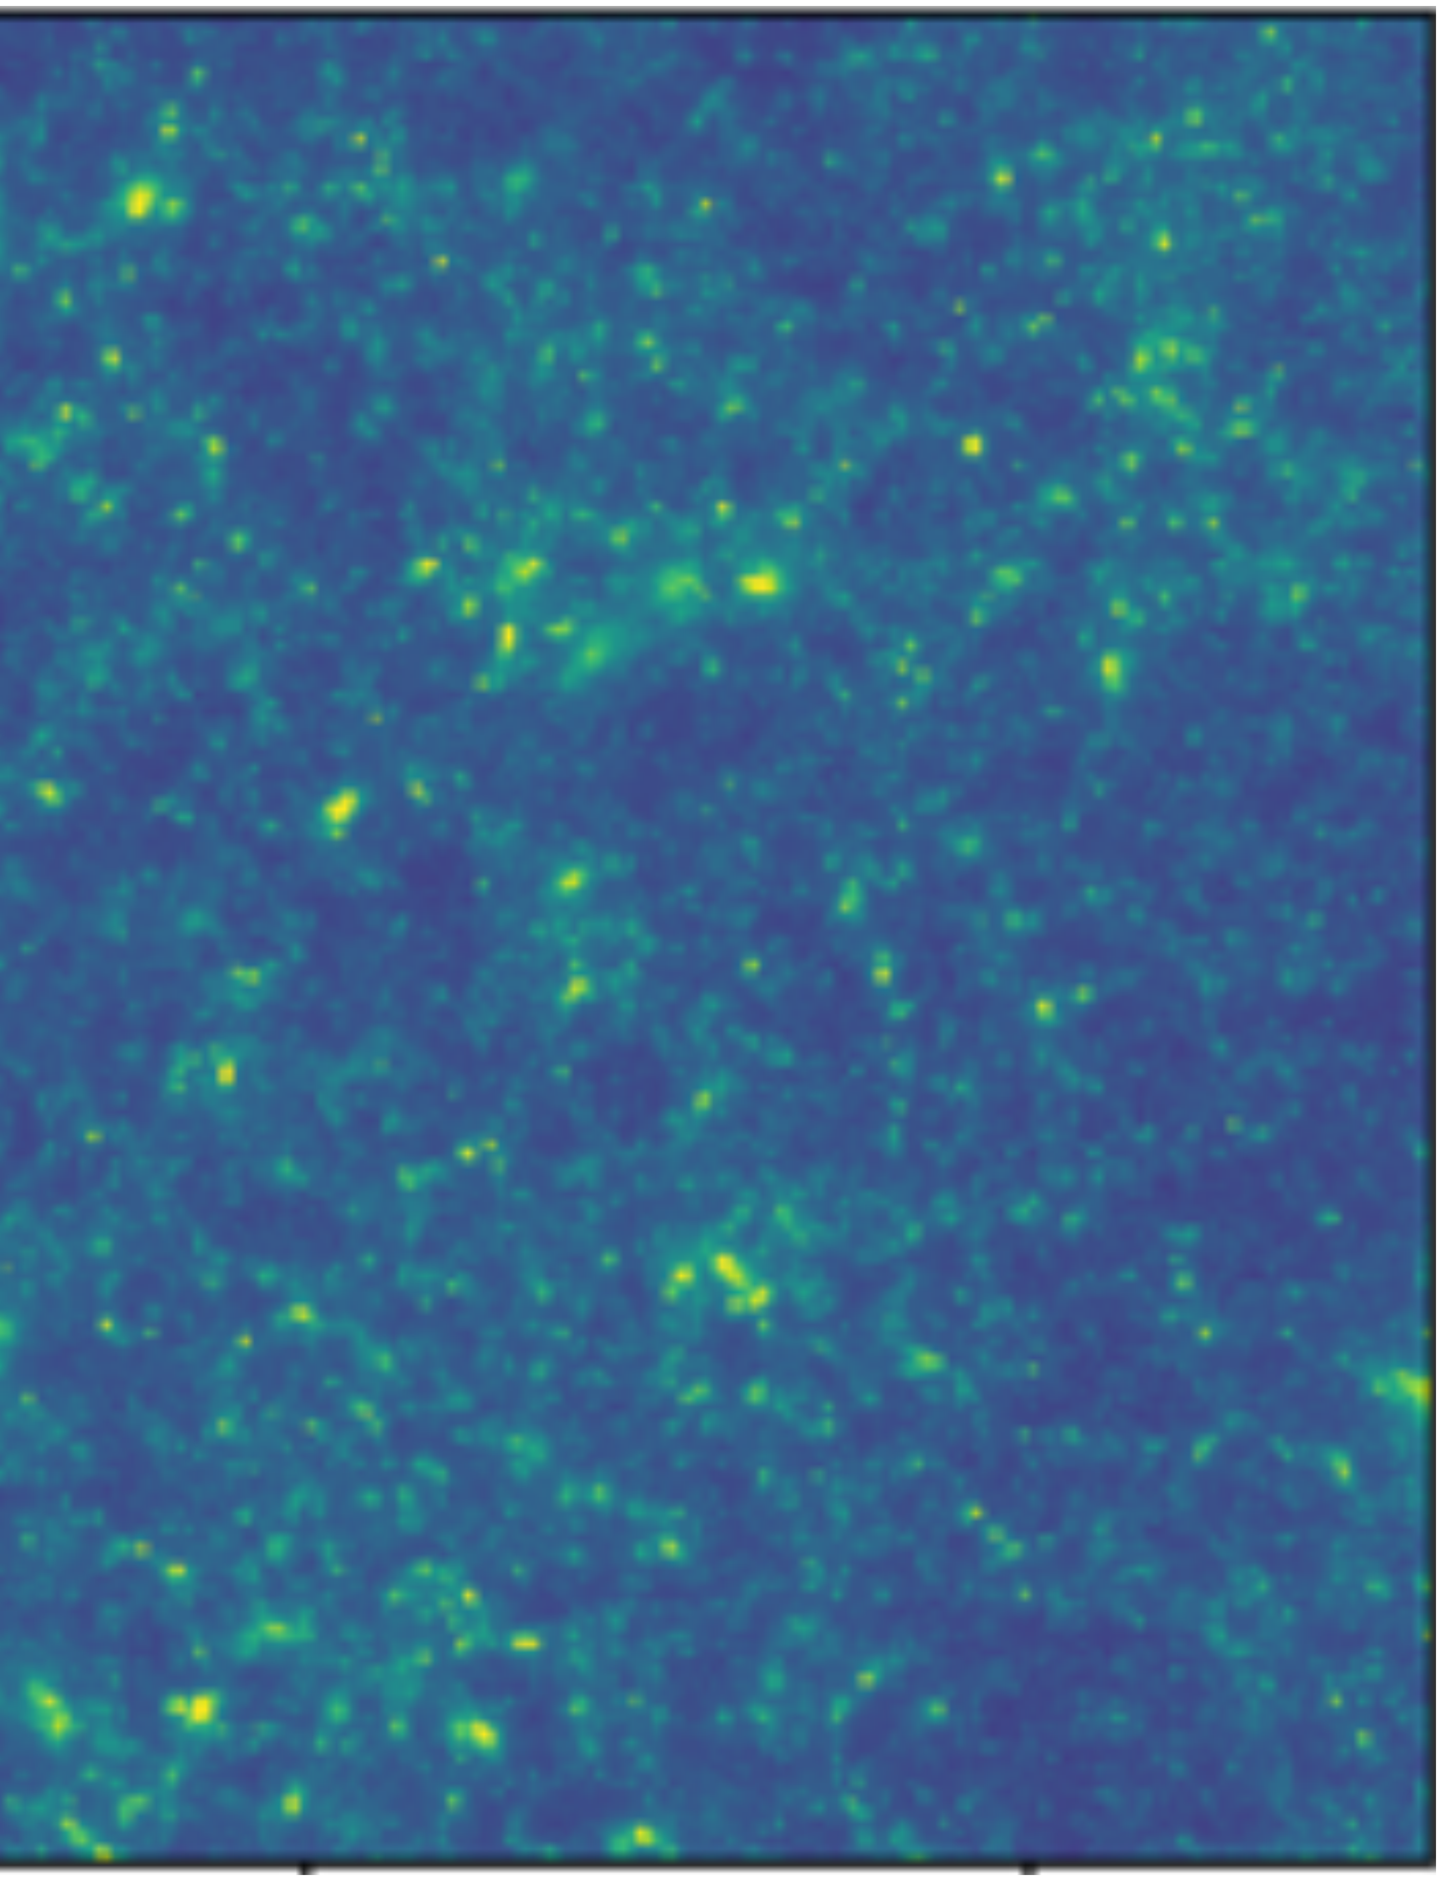

a few hundred scattering coefficients

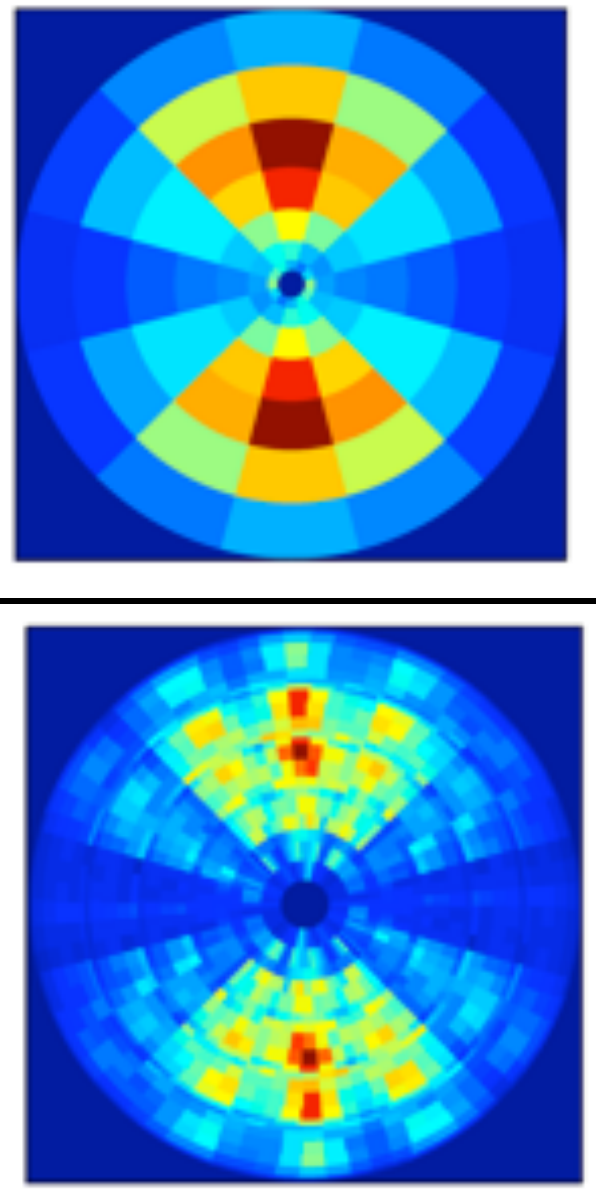

50 parameters (isotropy)

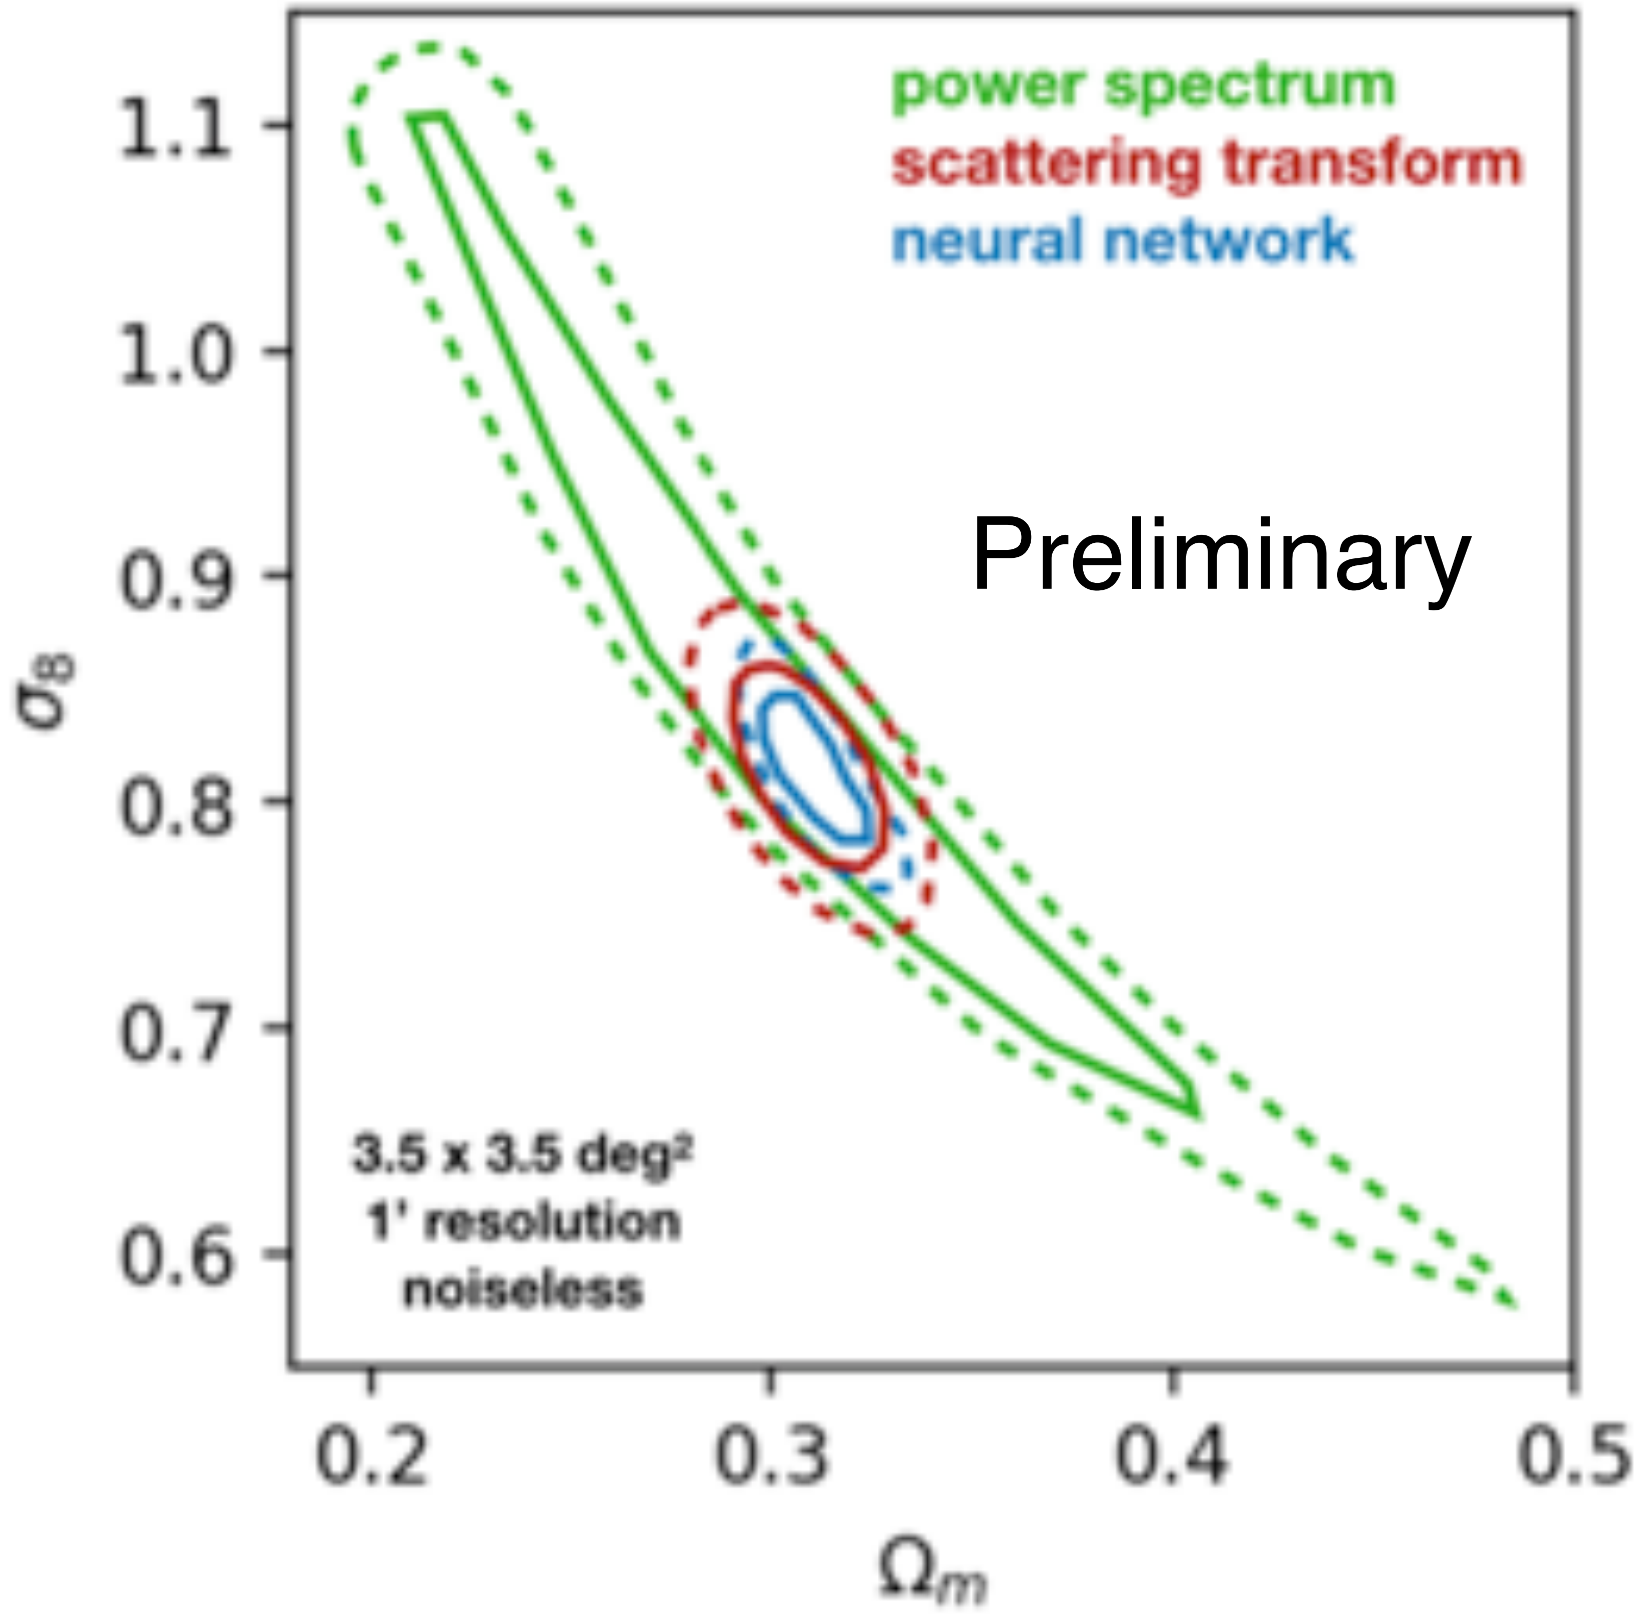

with grad student Sihao Cheng

& Yuan-Sen Ting

Gaussian statistics

manifold learning

scattering transform

convolutional neural nets

complexity

CNN

dictionary

Gaussian statistics, PCA

manifold learning

scattering transform

convolutional neural nets

complexity

collaboration with OpenAI

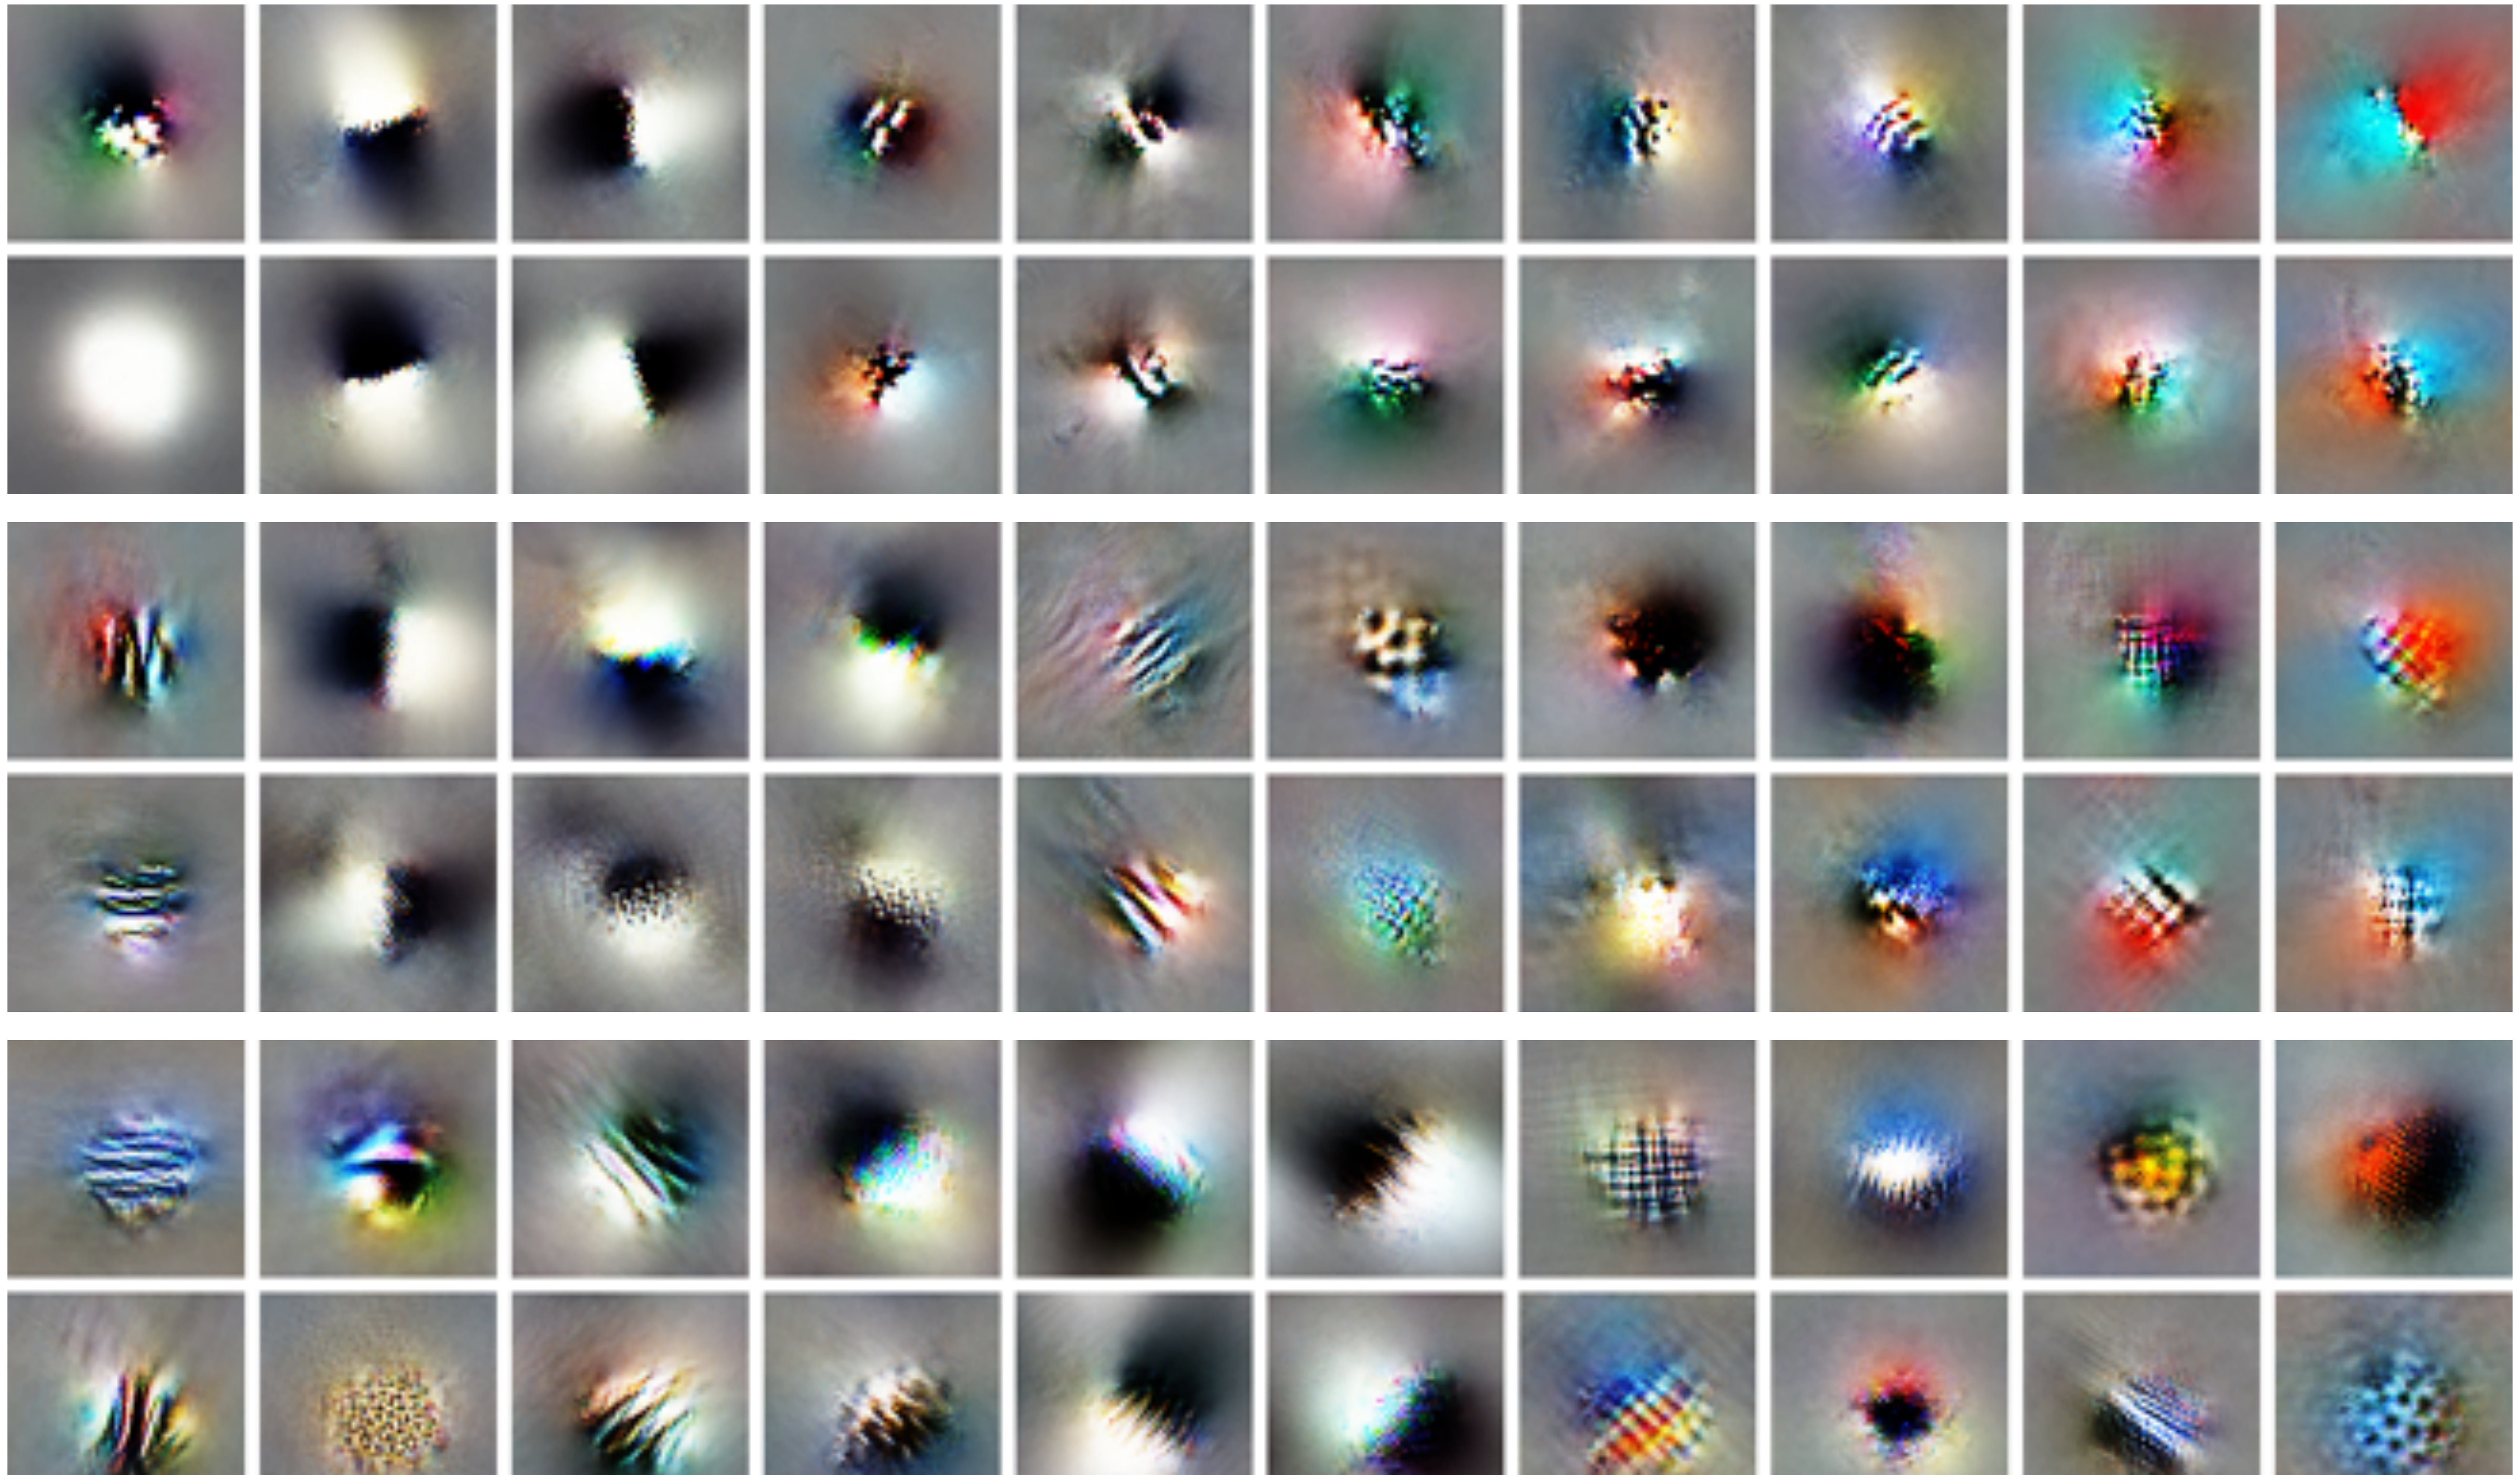

complexity

As scientists, we want to describe/explain the complex world

We want to go from complexity to simplicity.

This requires an efficient language.



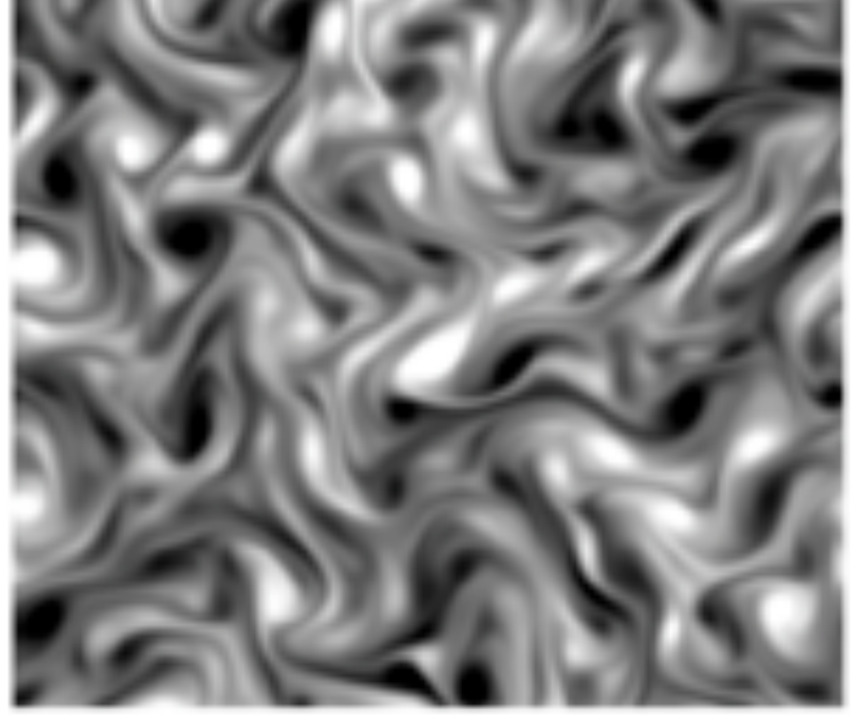
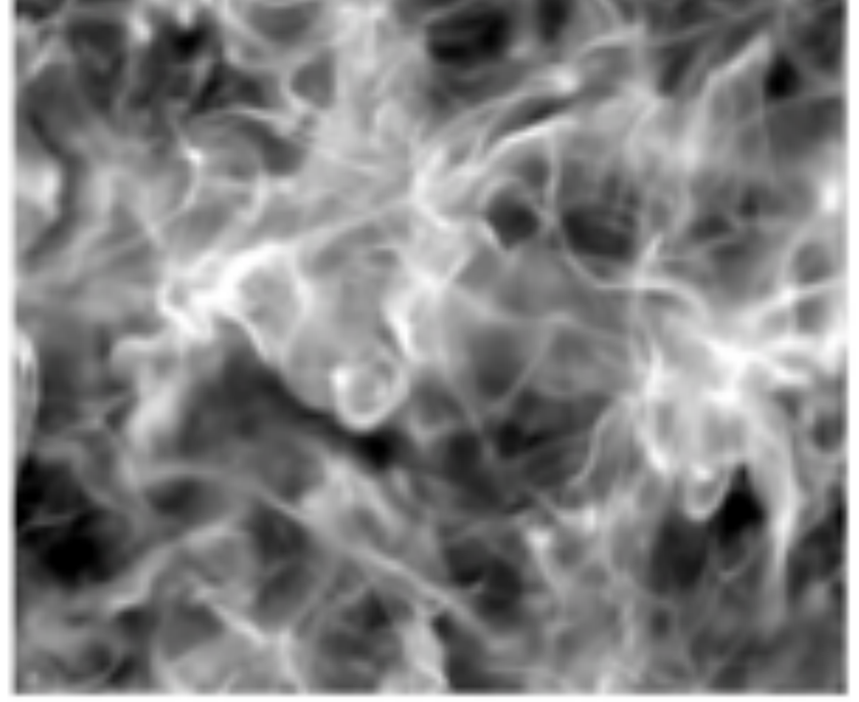


complexity

Gaussian statistics	manifold learning	dictionary learning	convolutional neural nets
parameters	geometry	vocabulary	classes

Simplicity found in:

Challenges: interpretability, control of systematics