Représentation parcimonieuse des textures à phase aléatoire

Lionel Moisan

MAP5, Université Paris Descartes

lionel.moisan@parisdescartes.fr

en collaboration avec Agnès Desolneux (CNRS, ENS Cachan) et Samuel Ronsin (MAP5)

Random Phase Textures

[van Wijk 1991] [Galerne, Gousseau, Morel 2011]

Two simple texture models associated to a given texture sample u (a $M \times N$ image):

- the Random Phase Noise (RPN)

$$\hat{U}_{RPN}(p,q) = e^{i\varphi(p,q)} \cdot \hat{u}(p,q),$$

  where the random variables $\varphi(p,q)$ are uniform on $[0, 2\pi]$ and independent within the symmetry constraint forced by the fact that U_{RPN} is real-valued;

- the Gaussian texture

$$\hat{U}_g(p,q) = \hat{W}(p,q) \cdot \hat{u}(p,q)$$

  (or equivalently, $U_g = W \star u$), where W is a normalized Gaussian white noise.

Remarks:

- These definitions still hold for color (or vector-valued) images

- In practice, to avoid periodization artifacts (due to the use of the Discrete Fourier transform), the initial texture sample u is replaced by its periodic component p [M. 2011]

Example of a RPN texture

initial sample (grass)

associated Random Phase Noise (RPN)

Asymptotics of RPN textures

The only difference between a RPN texture and the corresponding Gaussian model is a multiplication by a Rayleigh white noise $R(p,q)$ in Fourier domain. Indeed, one has

$$\hat{U}_g(p,q) = R(p,q) \cdot \hat{U}_{RPN}(p,q).$$

Theorem : Let $S : \mathbb{R}^2 \to R_+$ be a piecewise-continuous, $2\pi\mathbb{Z}^2$ periodic and bounded function such that $S(0) = 0$. Then, the (M,N)-periodic RPN texture defined by

$$\hat{U}(p,q) = \sqrt{MN} \cdot S\left(\frac{2\pi p}{M}, \frac{2\pi q}{N}\right) \cdot e^{i\varphi(p,q)}$$

for $|p| < M/2$ and $|q| < N/2$ converges (when M and N go to $+\infty$), in the finite dimensional joint distribution sense, to a zero-mean stationnary Gaussian field U defined on $\mathbb{Z}^2$ and such that

$$\mathbb{E}\left(U(0,0)U(k,l)\right) = \frac{1}{4\pi^2} \int_{[-\pi,\pi]^2} S^2(\alpha,\beta)\cos(k\alpha + l\beta)\, d\alpha d\beta.$$

Conclusion : The natural extension to $\mathbb{Z}^2$ of the RPN model is the Gaussian model

Sparse modeling of RPN textures

All images v such that $|\hat{v}| = |\hat{u}|$ define the same RPN model
$\rightarrow$ how to find a sparse representation of the RPN model associated to u ?

Motivations: texture analysis (notion of scale, comparison, etc.), synthesis of large texture images from a limited number of parameters

Natural ideal: find an image v that minimizes an energy criterion E (designed to favor concentration) under the constraint $|\hat{v}| = |\hat{u}|$.

Problem: the space $\{v \in \mathbb{R}^{\Omega}, |\hat{v}| = |\hat{u}|\}$ (a high-dimensional torus) is highly non-convex

Sparse modeling of RPN textures

All images v such that $|\hat{v}| = |\hat{u}|$ define the same RPN model
$\rightarrow$ how to find a sparse representation of the RPN model associated to u ?

Motivations: texture analysis (notion of scale, comparison, etc.), synthesis of large texture images from a limited number of parameters

Natural ideal: find an image v that minimizes an energy criterion E (designed to favor concentration) under the constraint $|\hat{v}| = |\hat{u}|$.

Problem: the space $\{v \in \mathbb{R}^{\Omega}, |\hat{v}| = |\hat{u}|\}$ (a high-dimensional torus) is highly non-convex

BUT....

$$\text{for} \quad E(v) = \sum_{k,l} A(k,l) v(k,l)^2 \quad \text{with} \quad A(k,l) = \sin^2\left(\frac{\pi k}{M}\right) + \sin^2\left(\frac{\pi l}{N}\right),$$

the minimum can be found **explicitly**.

$\rightarrow$ **Definition:** the **texton** of a discrete grey-level image u is the image

$$T(u) = \mathcal{F}^{-1}(|\hat{u}|).$$

A texton example

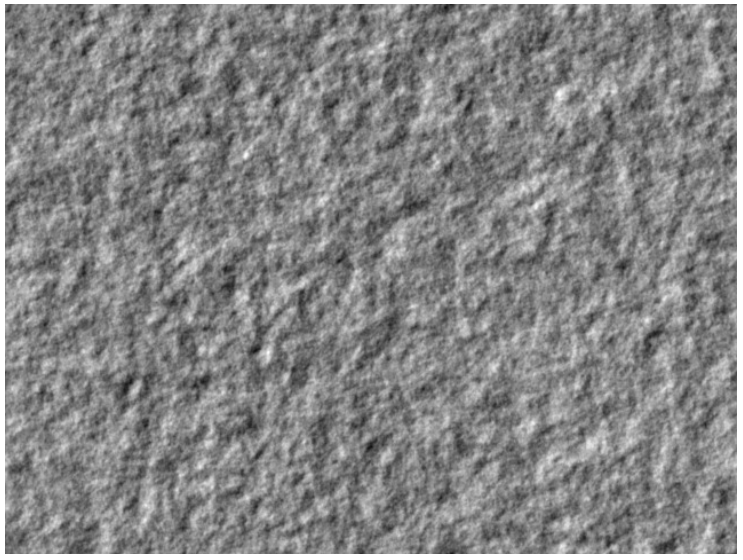

RPN texture sample (wall cover)

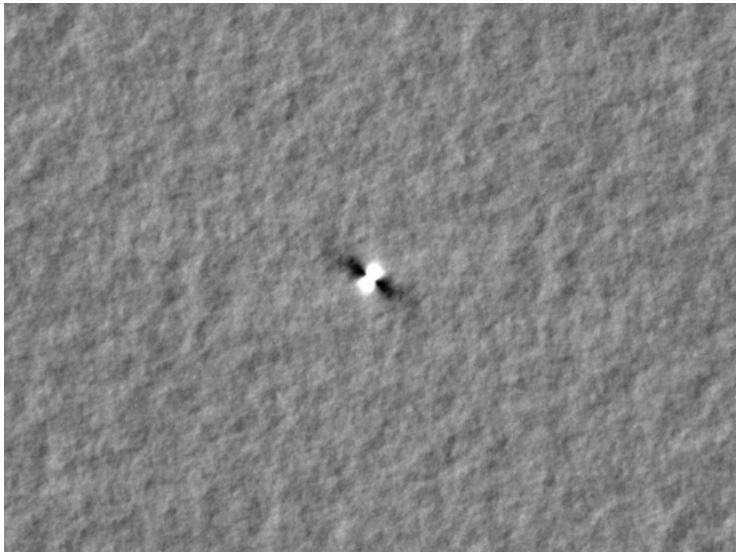

the associated texton

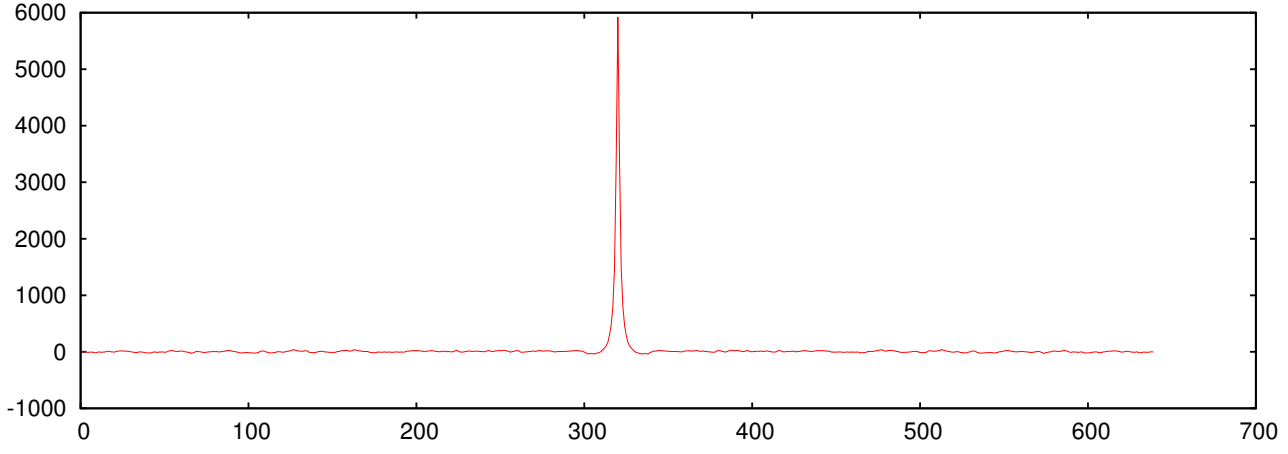

slice of the texton image on the line $y = 0$

Texton properties

- $T(u)$ is a symmetric, real-valued image
- The operator T is translation-invariant, and 1-Lipschitz (for the ℓ^2 norm)
- $T(\alpha u + \beta) = \alpha T(u) + \beta$ for any non-negative α, β, and u
- $T(T(u)) = T(u)$ (that is, $T(u)$ is its own texton)
- $T(k \star u) = T(k) \star T(u)$ (in particular, the "texton" operator commutes with kernels with positive Fourier Transform, e.g. Gaussian smoothing kernels)
- $T(v)$ is the unique image v that maximizes $v(0,0)$ under the constraint $|\hat{v}| = |\hat{u}|$
- The previous property remains true when $v(0,0)$ is replaced by any even-order derivative $\dfrac{\partial^{2m+2n} U}{\partial x^{2m} \partial y^{2n}}(0,0)$
- Let A be a symmetric real-valued image such that $A(0,0) = 0$ and $\widehat{A}(p,q) \leqslant 0$ for any $(p,q) \neq 0$. Then $A \geqslant 0$ and $T(u)$ minimizes

$$E(v) = \sum_{k,l} A(k,l) v(k,l)^2$$

under the constraint $|\hat{v}| = |\hat{u}|$.

What is the texton useful for?

Definition: The autocorrelation Γ of a discrete image $u : \Omega \to \mathbb{R}$ is defined by

$$\forall \mathbf{z} \in \Omega, \quad \Gamma(u)(\mathbf{z}) = \sum_{\mathbf{x} \in \Omega} u(\mathbf{x}) u(\mathbf{x} + \mathbf{z}).$$

- Since $\Gamma(RPN(u)) = T(u) \star T(u)$, the texton of u is less blurry (and more concentrated) than its autocorrelation $\Gamma(RPN(u))$
  $\to$ useful property for the analysis of RPN or Gaussian textures
- natural extension of a RPN texture: threshold (or approximate) the texton $T(u)$ to obtain a nearly equivalent texture sample that can be used to produce arbitrary large texture images.
- link between RPN and Gaussian textures: one has

  $$U_g \sim U_{RPN} \star T(W),$$

  where the texton of a white noise W is a kind of noisy Dirac mass:

$$\mathbb{E}\, T(W)(\mathbf{x}) = \sqrt{\frac{\pi}{2MN}} \cdot \begin{cases} MN - 1 & \text{if} \quad \mathbf{x} = 0, \\ 1 & \text{if} \quad \mathbf{x} \neq 0. \end{cases}$$

Texton versus Autocorrelation

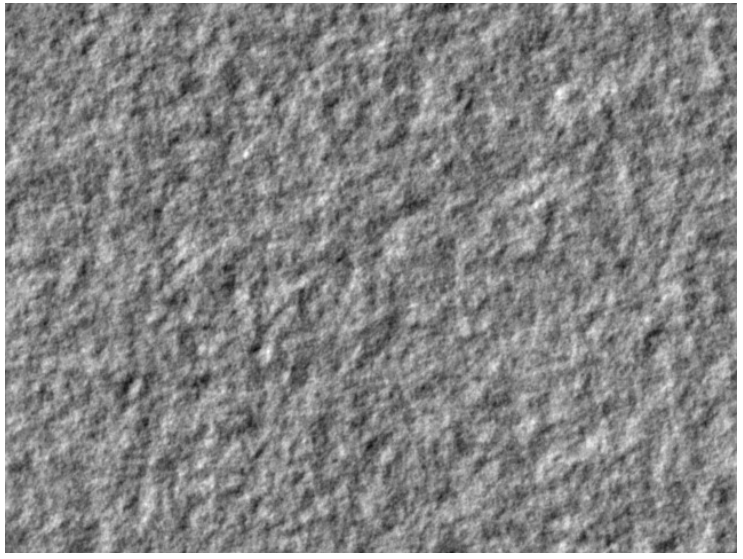

a RPN texture

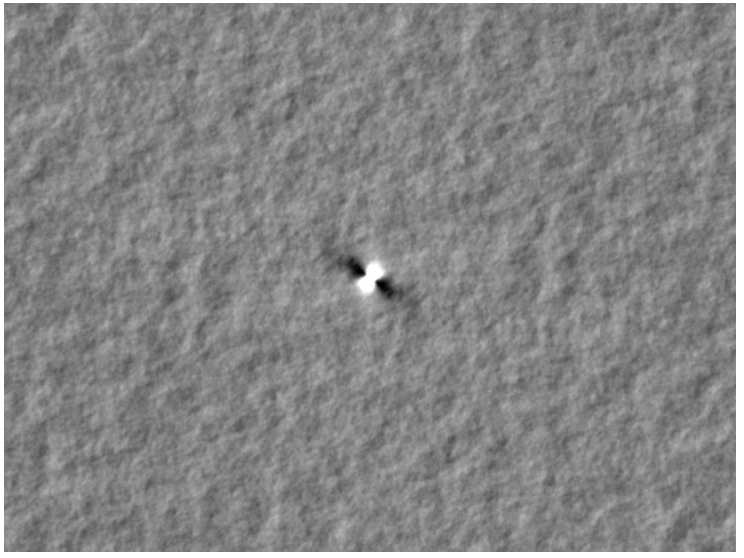

its texton

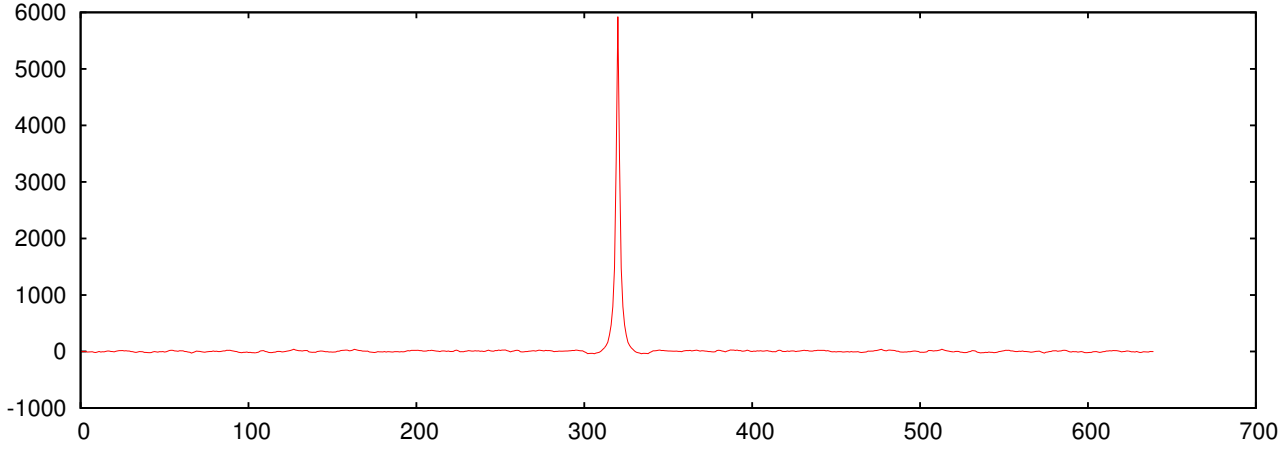

slice of the texton on the line $y = 0$

Texton versus Autocorrelation

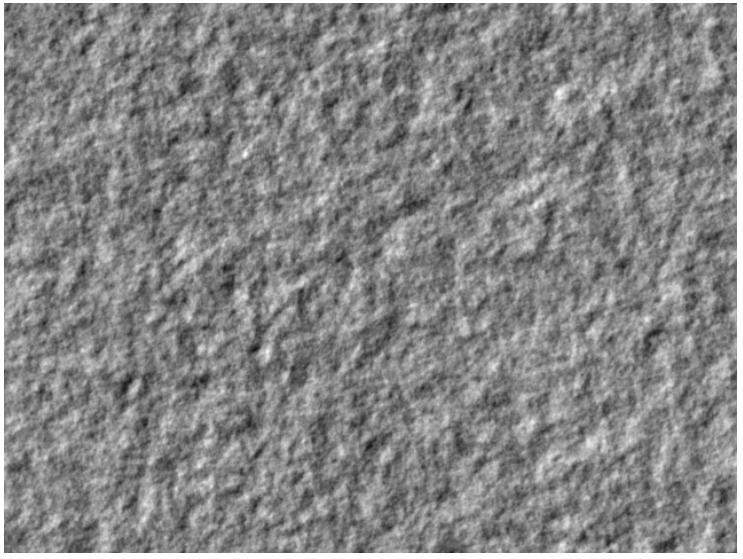

a RPN texture

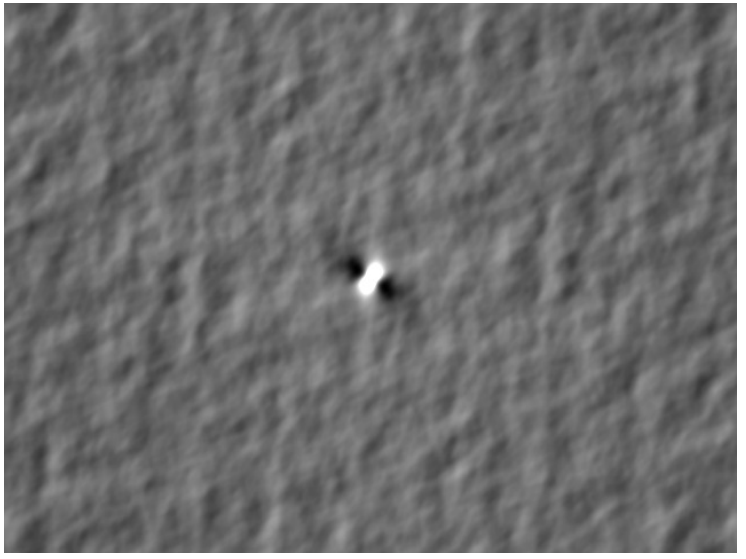

its autocorrelation

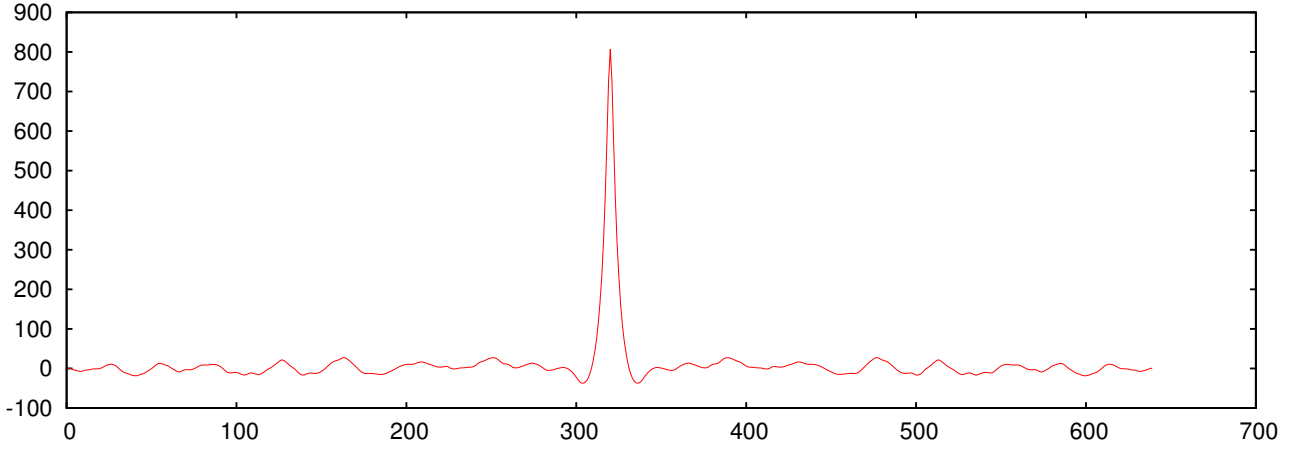

slice of the autocorrelation on the line $y = 0$

Extension to color

The autocorrelation matrix function of $\mathbf{u} : \Omega \to \mathbb{R}^d$ $(d = 3)$ is defined by

$$\forall \mathbf{z} \in \Omega, \quad \Gamma(\mathbf{u})(\mathbf{z}) = \sum_{\mathbf{x} \in \Omega} \mathbf{u}(\mathbf{x}) \mathbf{u}(\mathbf{x} + \mathbf{z})^T$$

or equivalently by $\quad \forall \xi \in \Omega, \quad \widehat{\Gamma(\mathbf{u})}(\xi) = \widehat{\mathbf{u}}(\xi)^* \, \widehat{\mathbf{u}}(\xi)^T.$

$\to \Gamma(\mathbf{u})$ is preserved as soon as a simultaneous phase perturbation is applied to all components of $\hat{\mathbf{u}}(\xi)$, that is

$$\widehat{\mathbf{U}_{RPN}}(\xi) = e^{i\varphi(\xi)} \widehat{\mathbf{u}}(\xi) \qquad \text{and} \qquad \widehat{\mathbf{U}_g}(\xi) = W(\xi) \widehat{\mathbf{u}}(\xi).$$

Definition of a color texton: given a reference normalized color $\alpha \in \mathbb{R}^3$ $(\alpha_1 + \alpha_2 + \alpha_3 = 1)$, one can define the monochrome image $\alpha \cdot \mathbf{u}$, and the color texton associated to $\mathbf{u}$ (and α) is then defined by

$$\widehat{T_\alpha(\mathbf{u})}(\xi) = e^{-i\varphi(\xi)} \cdot \hat{\mathbf{u}}(\xi),$$

where φ is the phase of $\mathcal{F}(\alpha \cdot \mathbf{u})$. One then has the property

$$T(\alpha \cdot \mathbf{u}) = \alpha \cdot T_\alpha(\mathbf{u}).$$

Examples of color texton

color RPN texture

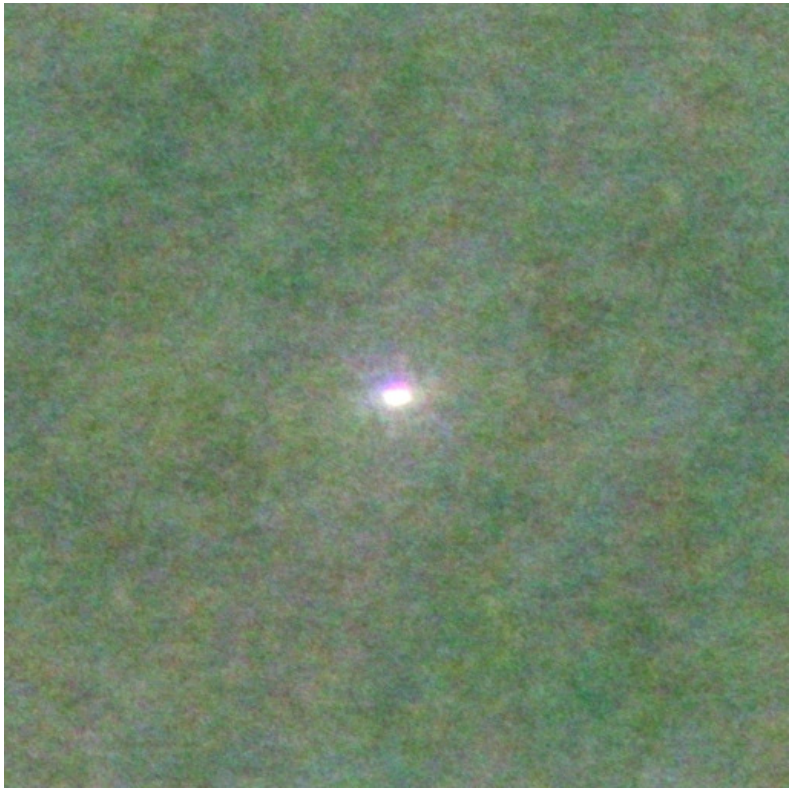

color texton

(color weights: 1/3,1/2,1/6)

Examples of color texton

original texture (wood)

color RPN texture

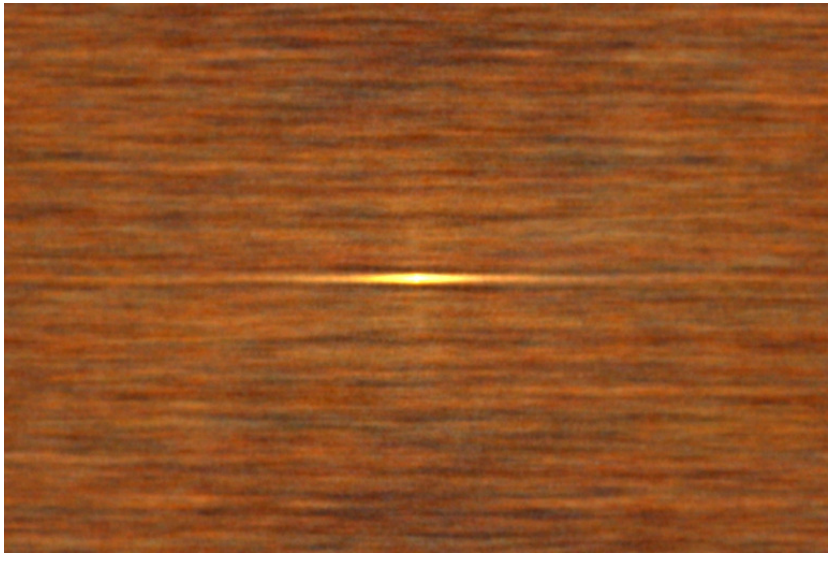

color texton (color weights: 1/3,1/2,1/6)

Color matrix texton

Another possibility to enforce the concentration property is to define a color matrix texton $M(\mathbf{u})$ by

$$\forall \xi \in \Omega, \qquad \widehat{M(\mathbf{u})}(\xi) = \frac{1}{\|\widehat{\mathbf{u}}(\xi)\|} \widehat{\mathbf{u}}(\xi) \widehat{\mathbf{u}}(\xi)^{*T},$$

$$\text{that is,} \qquad \widehat{M_{ij}(\mathbf{u})}(\xi) = \frac{\hat{u}_i(\xi)\hat{u}_j(\xi)^*}{\sqrt{\sum_k |\hat{u}_k(\xi)|^2}}.$$

This matrix texton (and approximations of it) permits to synthetize Gaussian fields by

$$\mathbf{U}_g = \mathcal{F}^{-1}\left(\widehat{M(\mathbf{u})\mathbf{W}}\right),$$

where $\mathbf{W}$ is a normalized white noise defined on $(\mathbb{C}^3)^\Omega$.

Note that when an approximation M' of $M(\mathbf{u})$ is used (by restriction, or thresholding), any realization of $\widehat{\Gamma(\mathbf{U}_g)}(\xi)$ will be a rank-one matrix, though

$$\widehat{\mathbb{E}\Gamma(\mathbf{U}_g)}(\xi) = M'(\xi)^* M'(\xi)^T$$

will be a full rank matrix in general.

Color texton versus color matrix texton

color RPN texture

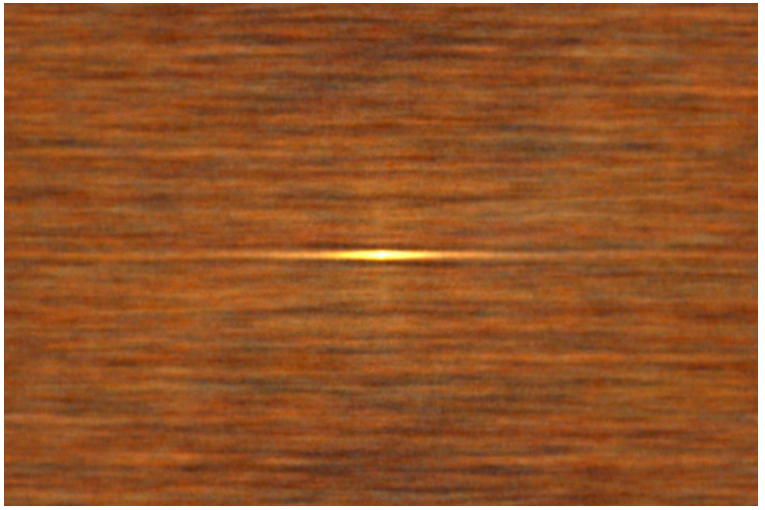

color texton

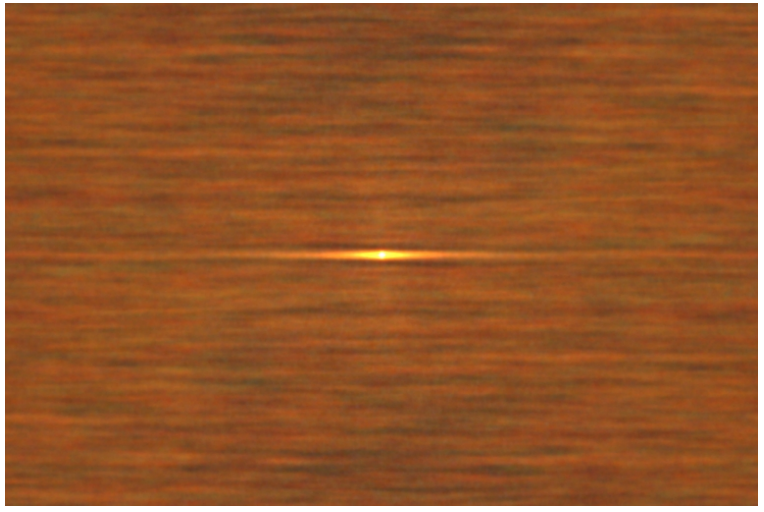

color matrix texton (diagonal)

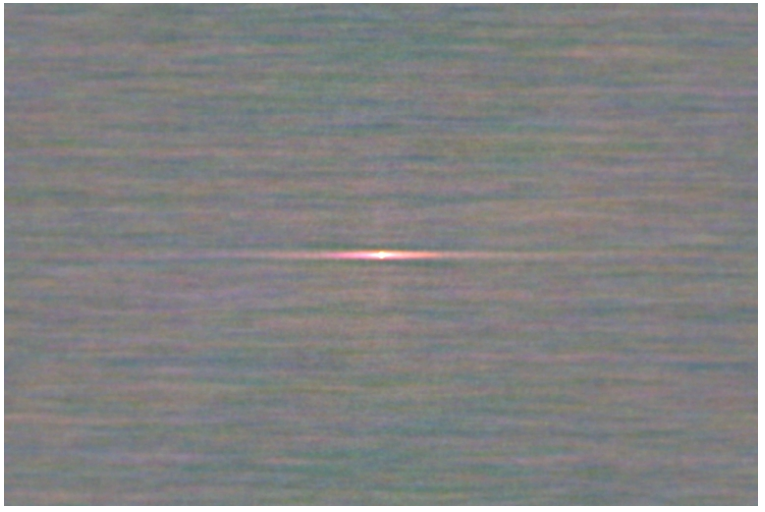

color matrix texton (off diagonal)

Compact texton

Idea: approximate the full texton image by an image with small support

Interest:

- reduced storage
- synthesis of a texture with arbitrary large dimensions

Strategies:

- (bi-)thresholding
- clipping with a fixed shape (disc, ellipse, etc.)
- clipping with an adapted shape

How to rate an approximate texton?
$\rightarrow$ use the transport distance between the associated textures

$$d^2(T_1, T_2) = \sum_{\xi} \left(|\hat{T}_1(\xi)| - |\hat{T}_2(\xi)| \right)^2$$

Compact texton

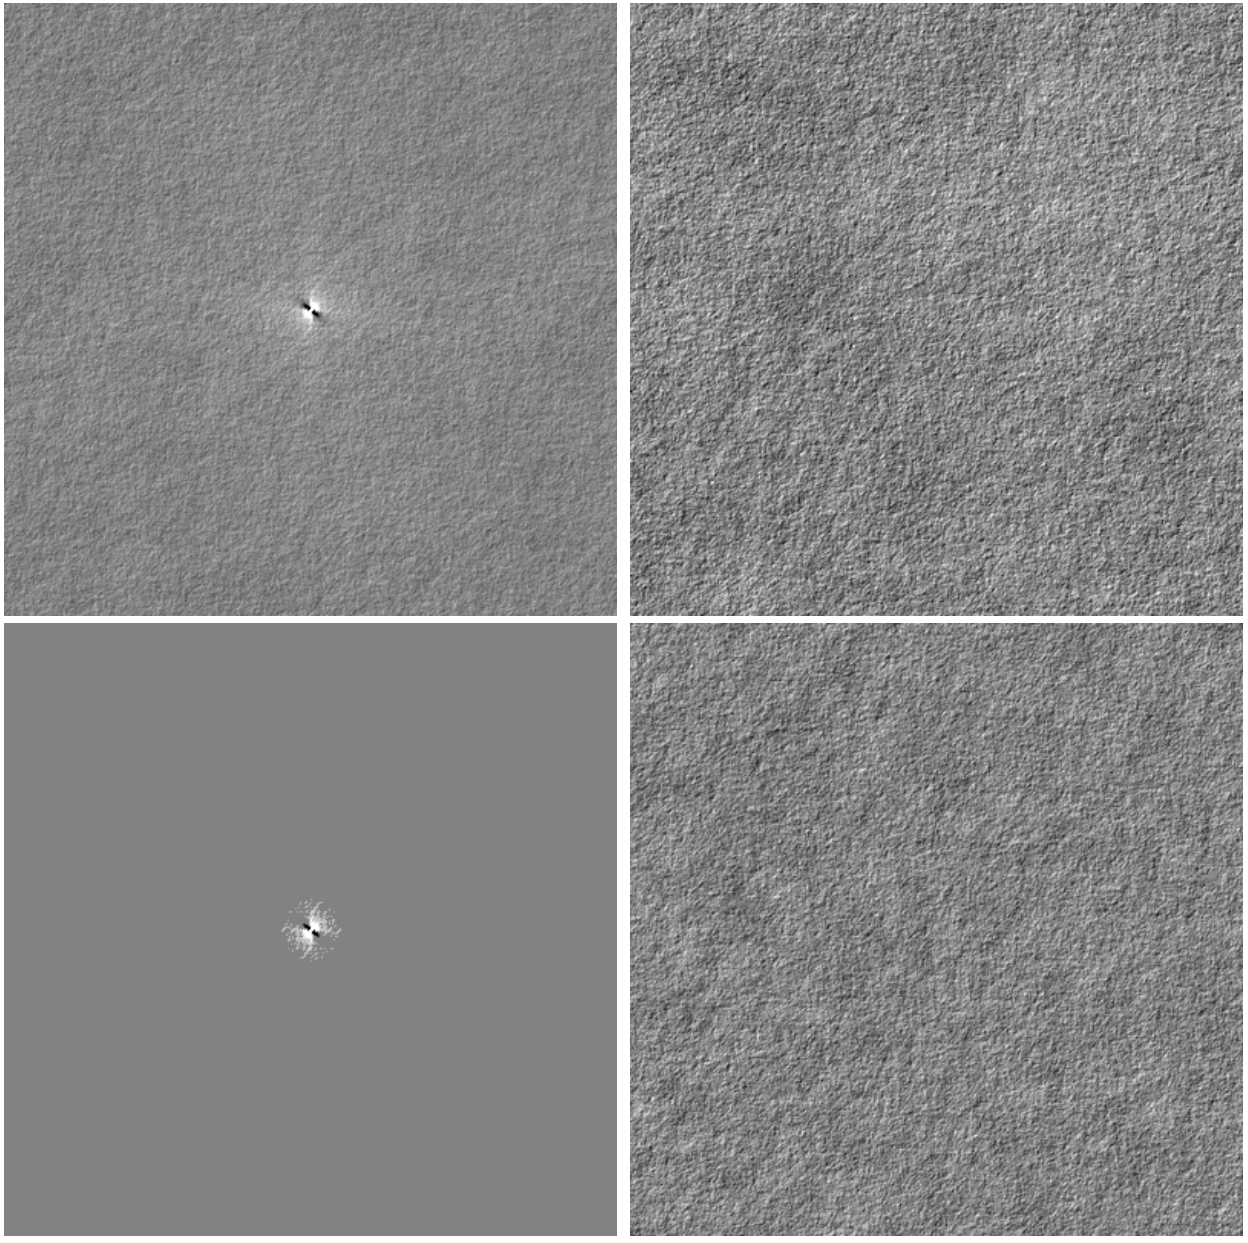

Compact texton

Compact texton

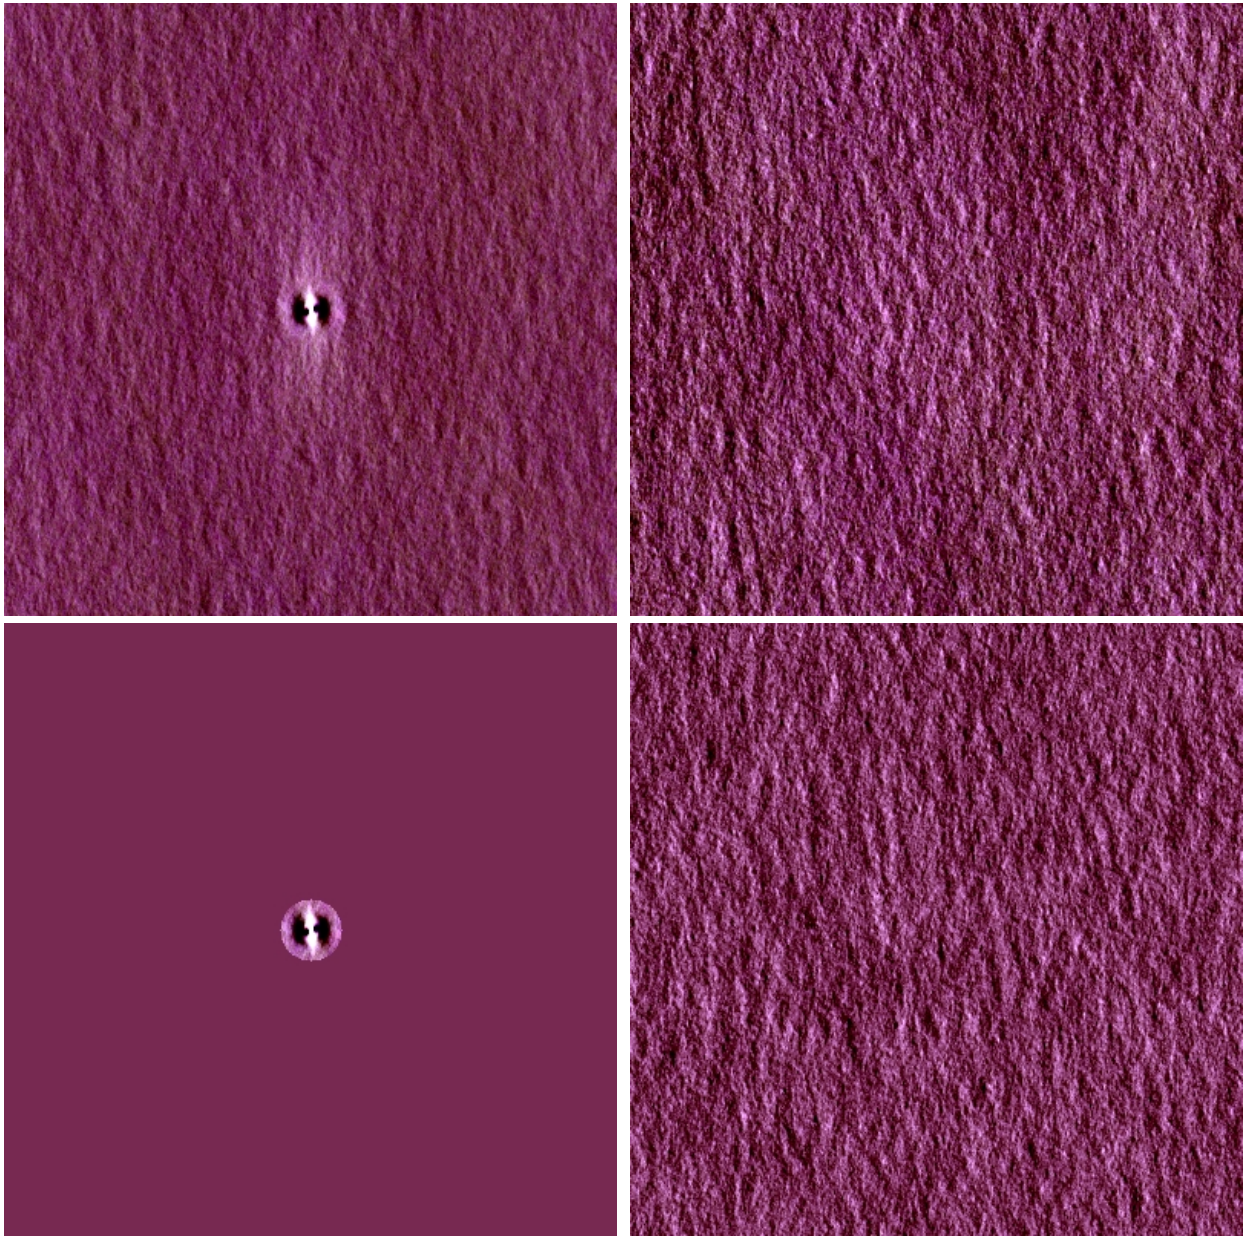

Compact texton