A Texton for Fast and Flexible Gaussian Texture Synthesis

Arthur Leclaire[1,2]
Joint work with B. Galerne[1] and L. Moisan[1]

[1]Université Paris Descartes
MAP5, CNRS UMR8145

[2]ENS Cachan
CMLA, CNRS UMR8536

SMATI
May 7, 2015

Introduction

The class of texture images is not really well defined. However, one can think of a texture as an image whose content consists in repeated patterns with a certain amount of randomness [Wei et al., 2009].

One can distinguish between

- macro-textures,

- micro-textures.

Introduction

In this talk, we will be interested in the example-based synthesis of microtextures.

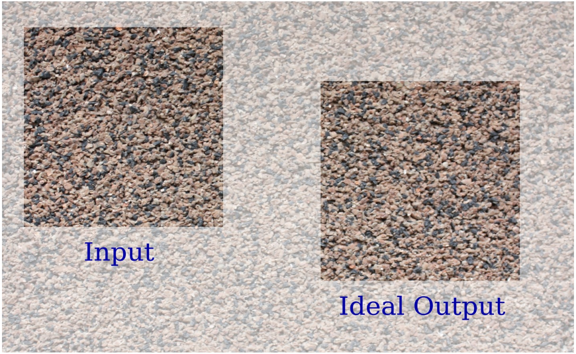

We are looking for a **fast** and **flexible** synthesis scheme.

Introduction

We will first present the Gaussian and Spot Noise texture models.

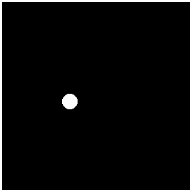
1 spot

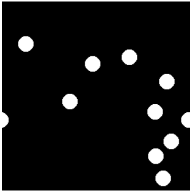
10 spots

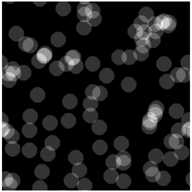
100 spots

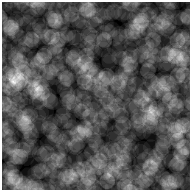
1000 spots

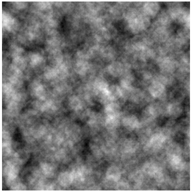
10000 spots

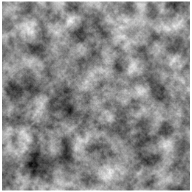
Gaussian limit

[Van Wijk, 1991], [Galerne et al., 2011]

Introduction

Main Question: Given a small support S and an exemplar texture, how to find a kernel h with support $\subset S$ such that the spot noise model associated to h is a good approximation of the Gaussian texture, even with a reasonably low number of spots?

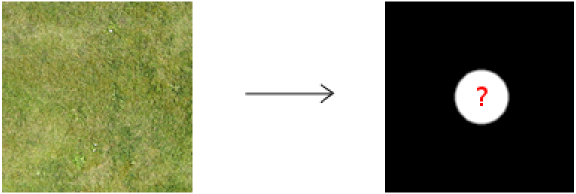

Introduction

At the end of the talk, we will be able to do this:


Exemplar


1 spot

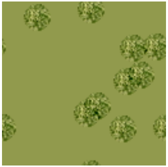
10 spots

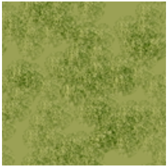
100 spots


Kernel

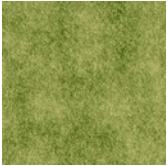
1000 spots

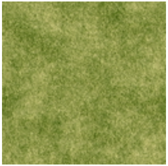
10000 spots


Gaussian limit

Outline

Non-normalized Spot Noise

Let $h : \mathbb{Z}^2 \to \mathbb{R}^d$ be a function with compact support S_h.

The **discrete Spot Noise** on $\mathbb{Z}^2$ with spot h and intensity λ is the stationary random process $F_{h,\lambda} : \mathbb{Z}^2 \to \mathbb{R}^d$ defined by

$$\forall \mathbf{x} \in \mathbb{Z}^2, \quad F_{h,\lambda}(\mathbf{x}) = \sum_{i \geqslant 1} h(\mathbf{x} - \mathbf{X}_i) \ ,$$

where $(\mathbf{X}_i)$ is a Poisson point process on $\mathbb{Z}^2$ with intensity λ.

If we set $P_\lambda(\mathbf{y}) = |\{i \geqslant 1 \ | \ \mathbf{X}_i = \mathbf{y}\}|$ one has $F_{h,\lambda} = h * P_\lambda$.

$h \qquad P_\lambda \qquad F_{h,\lambda}$

Normalized Spot Noise and Gaussian Limit

Denoting $\Sigma h = \sum_{\mathbf{x}} h(\mathbf{x})$ and $\tilde{h}(\mathbf{x}) = h(-\mathbf{x})$, we have

$$\mathbb{E}\Big(F_{h,\lambda}(\mathbf{x})\Big) = \lambda \cdot \Sigma h \, ,$$
$$\mathbb{E}\Big((F_{h,\lambda}(\mathbf{x}) - m)(F_{h,\lambda}(\mathbf{x}+\mathbf{v}) - m)^T\Big) = \lambda \cdot h * \tilde{h}^T(\mathbf{v}) \, .$$

Therefore, the normalized Spot Noise

$$\frac{F_{h,\lambda} - \mathbb{E}(F_{h,\lambda})}{\sqrt{\lambda}} = \frac{1}{\sqrt{\lambda}}\Big(h * P_\lambda - \lambda \cdot \Sigma h\Big) \, ,$$

has zero-mean and covariance function $h * \tilde{h}^T$.

Theorem (Papoulis, 1971)

$$\mathrm{DSN}_\lambda(h) := \frac{1}{\sqrt{\lambda}}\Big(h * P_\lambda - \lambda \cdot \sum h\Big) \xrightarrow[\lambda \to +\infty]{(d)} h * W =: \mathrm{ADSN}(h) \, ,$$

where W is a scalar Gaussian white noise on $\mathbb{Z}^2$ with mean 0 and variance 1.

Circular Discrete Spot Noise

One can define "circular" or periodic spot noise models on

$$\Theta = \mathbb{Z}/M\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z} \; .$$

Using circular convolutions, one can define

$$\mathrm{CDSN}_\lambda(h) = \frac{1}{\lambda}\left(h \odot P_\lambda - \lambda \cdot \sum h\right) \; ,$$
$$\mathrm{CADSN}(h) = h \odot W \; .$$

The discrete Fourier transform gives

$$\widehat{h \odot W} = \hat{h}\hat{W} \; ,$$

leading to a simple algorithm of simulation for the CADSN (with complexity $\mathcal{O}(MN\log(MN))$). Idem with the CDSN.

DSN simulation on a finite domain

Notice that the value at $\mathbf{x}$ is only affected by spots located at $\mathbf{x} - S_h$.

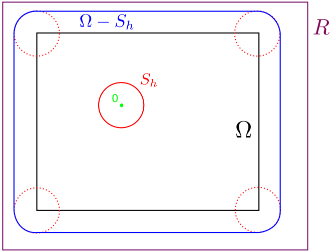

$$\Omega - S_h := \{ \mathbf{x} - \mathbf{z} \; ; \; \mathbf{x} \in \Omega \; , \; \mathbf{z} \in S_h \}$$

Consequences:

- Underlying white noise processes need only be sampled on $\overline{\Omega} := \Omega - S_h$.
- The restrictions on Ω of circular or non-circular models computed on R are the same.

DSN simulation on a finite domain

Algorithm 1: DSN simulation on a finite domain Ω

- Set $\bar{\Omega} = \Omega - S_h = \{\mathbf{x} - \mathbf{y}\ ;\ \mathbf{x} \in \Omega, \mathbf{y} \in S_h\}$.
- Draw n with Poisson distribution of intensity $\lambda|\bar{\Omega}|$.
- Draw $\mathbf{x}_1, \ldots, \mathbf{x}_n$ independently and uniformly in $\bar{\Omega}$.
- $\forall \mathbf{x} \in \Omega, \quad g(\mathbf{x}) := \frac{1}{\sqrt{\lambda}} \left(\sum_{i=1}^{n} h(\mathbf{x} - \mathbf{x}_i) - \lambda \sum h \right).$

Linear mean computational cost:

$$N_{imp} = \lambda |S_h|$$

operations per pixel (number of impacts per pixel).

Parallel Spot Noise Synthesis

DSN synthesis can be parallelized using a grid-based simulation scheme for the Poisson point process [Lagae et al., 2009].

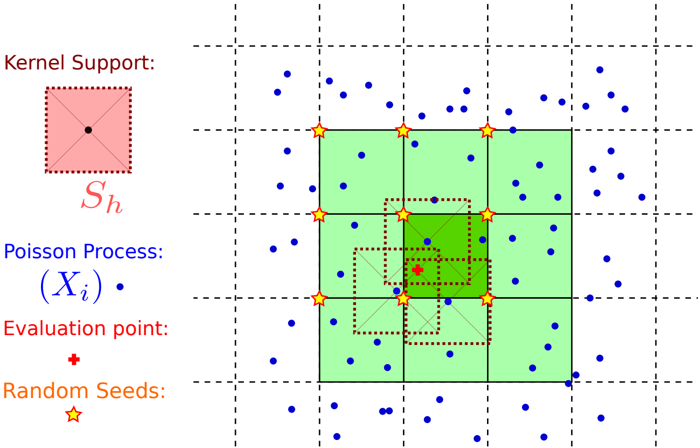

Spot Noise examples

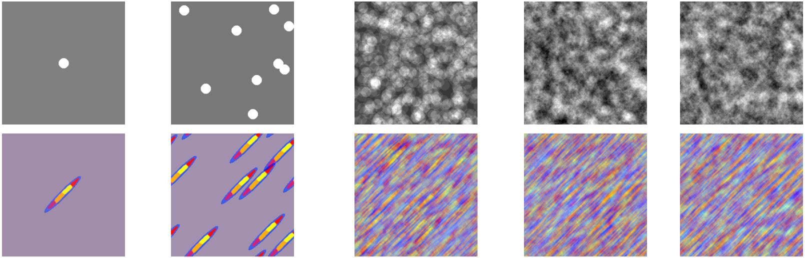

h $\quad \mathrm{CDSN}_{10^{-4}}(h)$ $\quad \mathrm{CDSN}_{10^{-2}}(h)$ $\quad \mathrm{CDSN}_{1}(h)$ $\quad \mathrm{CADSN}(h)$

Estimation of a circular spot noise model

Goal: Synthesize an exemplar $u : \Omega \to \mathbb{R}^d$ on the same domain.

- The mean value can be estimated by

$$\bar{u} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} u(\mathbf{x}) \ .$$

- For the covariance, identifying Ω to a circular domain Θ, we can consider the estimator

$$c_u(\mathbf{v}) = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \Big(u(\mathbf{x}) - \bar{u} \Big) \Big(u\big(\mathbf{x} + \mathbf{v} \bmod \Omega\big) - \bar{u} \Big)^T \ ,$$

which rewrites $c_u = t_u \odot \tilde{t}_u^T$ where $t_u = \frac{1}{\sqrt{|\Omega|}}(u - \bar{u})$.

Examples

Therefore, a circular Gaussian texture synthesis is obtained by drawing a realization of $\bar{u} + \mathrm{CADSN}(t_u) = \bar{u} + t_u \odot W$.

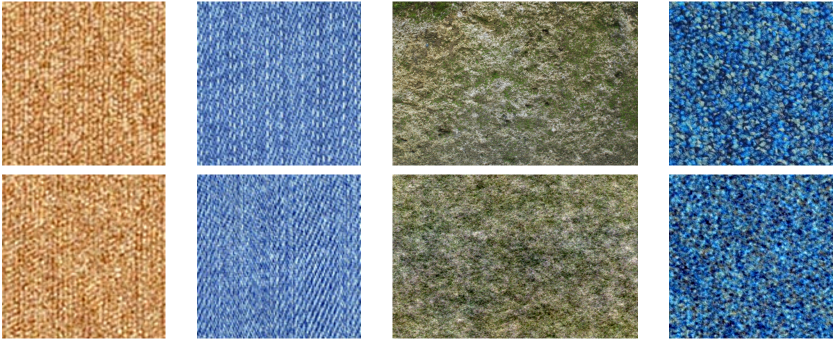

First row: Original textures.
Second row: CADSN synthesis.

Remarks

- A circular model is not legitimate for non-periodic exemplars.
  $\rightarrow$ Periodicity can be forced by taking the periodic component of the exemplar as a pre-computation step [Moisan, 2011].

- Among the kernels h such that

$$h \odot \tilde{h}^T = c_u \ ,$$

  is there one that is more interesting?
  $\rightarrow$ For the graylevel case, the authors of [Desolneux et al., 2012] isolated the one with zero-phase, called canonical texton of u. It is obtained in Fourier domain as

$$\widehat{t}_{can} = \sqrt{\widehat{c}_u} = |\widehat{t}_u| \ .$$

Synthesis on $\mathbb{Z}^2$

Goal: Synthesize an exemplar $u : \Omega \to \mathbb{R}^d$ on a **wider domain.**

$\rightarrow$ One possible choice is to extend t_u by zero-padding

$$\tau_u(\mathbf{x}) = \begin{cases} t_u(\mathbf{x}) = \frac{1}{\sqrt{|\Omega|}}(u(\mathbf{x}) - \bar{u}) & \text{if} \quad \mathbf{x} \in \Omega \\ 0 & \text{if} \quad \mathbf{x} \notin \Omega \end{cases} .$$

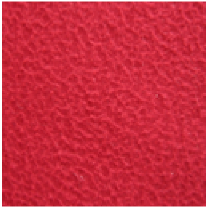

Original u

Synthesis $\bar{u} + \mathrm{ADSN}(\tau_u)$

$\rightarrow$ In order to avoid high-frequency artifacts, the border discontinuity of t_u can be attenuated by a smooth window [Galerne et al., 2011].

Optimal Transport Distance and Model Projection

- [Xia et al., 2013] showed that the L^2 optimal transport distance between $\mu_0 = \mathrm{CADSN}(h_0)$ and $\mu_1 = \mathrm{CADSN}(h_1)$ is given by

$$d_{OT}(\mu_0, \mu_1)^2 = \frac{1}{|\Theta|} \sum_{\boldsymbol{\xi} \in \Theta} \|\hat{h}_0(\boldsymbol{\xi})\|^2 + \|\hat{h}_1(\boldsymbol{\xi})\|^2 - 2|\hat{h}_0(\boldsymbol{\xi})^* \hat{h}_1(\boldsymbol{\xi})| \, .$$

- This allows us to define a projection of h_1 on the set of kernels associated to the model μ_0 as a solution of

$$\operatorname*{Argmin}_{h,\ h \odot \tilde{h} = h_0 \odot \tilde{h}_0} \quad d_{OT}(\mathrm{CADSN}(h), \mathrm{CADSN}(h_1)) \, .$$

- One particular solution $p_{h_0}(h_1)$ can be computed by

$$\widehat{p_{h_0}(h_1)} = \widehat{h_0} \, \frac{\widehat{h_0}^* \widehat{h_1}}{|\widehat{h_0}^* \widehat{h_1}|} \mathbf{1}_{\widehat{h_0}^* \widehat{h_1} \neq 0} + \widehat{h_0} \, \mathbf{1}_{\widehat{h_0}^* \widehat{h_1} = 0} \, .$$

Outline

A texton for Spot Noise synthesis?

Let $u : \Omega \to \mathbb{R}^3$, and a finite $S \subset \mathbb{Z}^2$.

Goal: Compute a kernel h with support $S_h \subset S$ such that

1. a piece of $\mathrm{ADSN}(h)$ looks like u (which means $h * \tilde{h}^T \approx c_u$) ,
2. $\mathrm{DSN}_\lambda(h)$ is a good approximation of $\mathrm{ADSN}(h)$, even for reasonably low λ (which means...?) .

This would allow to synthesize u with $\mathrm{DSN}_\lambda(h)$ for low λ.
Benefits:

- **Very fast** (faster than the DFT-based method for large domains).
- Evaluations can be **parallelized** (using a coherent evaluation scheme for the Poisson point process).
- Allows for **local evaluations**.

SOT computation

Inspired by the phase-retrieval litterature [Hayes, 1982], we propose the following algorithm to compute a **Synthesis-Oriented Texton**.

Algorithm 2: SOT computation

- Compute $t_u = \frac{1}{\sqrt{MN}}(u - \bar{u})$ and its DFT.
- **Initialization**: $t \leftarrow$ Gaussian white noise.
- **Main loop:** Repeat ($n \simeq 50$ times)
  - Spectral projection: $t \leftarrow p_{t_u}(t)$.
  - Support projection: $t \leftarrow t \cdot \mathbf{1}_S$.

Recall that

$$\widehat{p_h(t)} = \widehat{h}\,\frac{\widehat{h}^*\widehat{t}}{|\widehat{h}^*\widehat{t}|}\mathbf{1}_{\widehat{h}^*\widehat{t}\neq 0} + \widehat{h}\,\mathbf{1}_{\widehat{h}^*\widehat{t}=0}\ .$$

Convergence to a random point

- We observed that Algorithm 2 converges to a random point.
- Considering the relative model error

$$\mathrm{RME}(t, t_u)^2 = \frac{\sum_\xi \left(|\widehat{t_u}|^2 + |\widehat{t}\,|^2 - |\widehat{t_u^*}\widehat{t}\,| \right)(\xi)}{\sum_\xi |\widehat{t_u}|^2(\xi)},$$

the mean RME stabilizes after a few iterations.

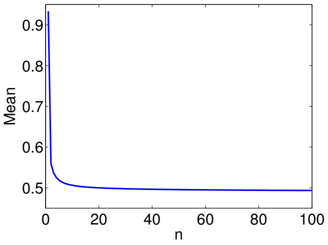

DSN synthesis of a natural texture, comparison

u

$\mathrm{DSN}(t)$, 50 imp./px

$\mathrm{DSN}(t_{lum}^{crop})$

$\mathrm{DSN}(t_{rpn})$

t
RME = 0.50

t_{lum}^{crop}
RME = 0.52

t_{rpn}
RME = 0.70

DSN results obtained with kernels of size 31×31:

- a synthesis-oriented texton t,
- the cropped luminance texton t_{lum}^{crop} [Desolneux, et al. 2012],
- a cropped realization of the random phase noise associated to u.

Synthesis of color textures

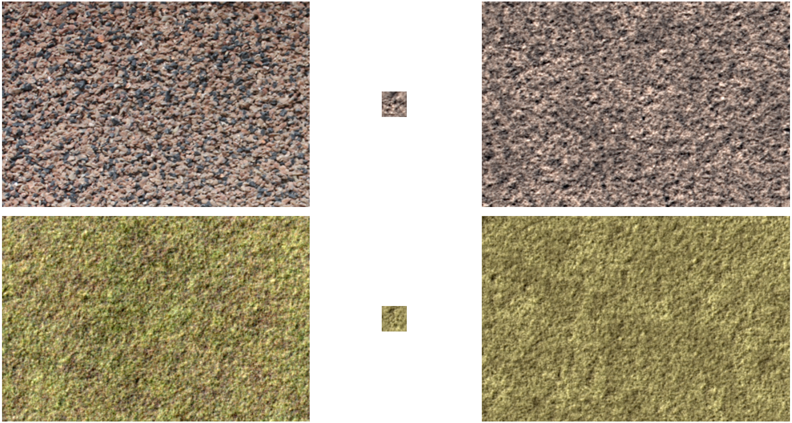

Original
(384 $\times$ 256)

SOT
(31 $\times$ 31)

DSN synthesis
(50 impacts per pixel)

Possible color correction

Let $t : \mathbb{Z}^2 \to \mathbb{R}^3$ be a color SOT computed with Algorithm 2.

- The covariance of one pixel value of $\mathrm{ADSN}(t)$ is the matrix

$$C(0) = \sum_{\mathbf{x} \in S} t(\mathbf{x}) t(\mathbf{x})^T .$$

Nothing ensures that $C(0) = c_u(0) \implies$ Possible color loss.

- As in [Desolneux et al., 2012], we suggest to apply a color transformation matrix

$$M = \sqrt{c_u(0)} \sqrt{C^{-1}}$$

Results with color correction

Original (384 × 256)

t
RME = 0.51

t_{cc}
RME = 0.54

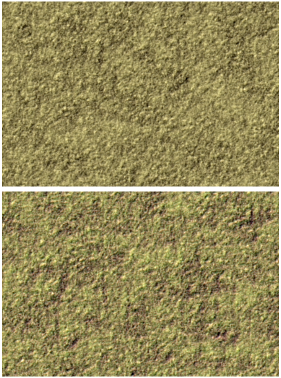

DSN (50 imp./px)

DSN synthesis with SOT of size 31 × 31.
Up: without color correction.
Bottom: with color correction.

Results with color correction

Original (384 × 256)

t
RME = 0.49

t_{cc}
RME = 0.52

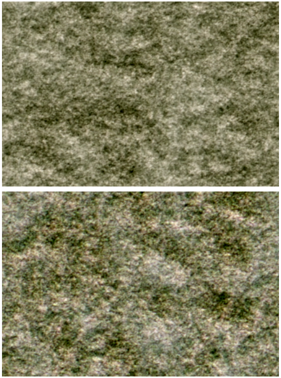

DSN (50 imp./px)

DSN synthesis with SOT of size 51 × 51.
Up: without color correction.
Bottom: with color correction.

Influence of the support size

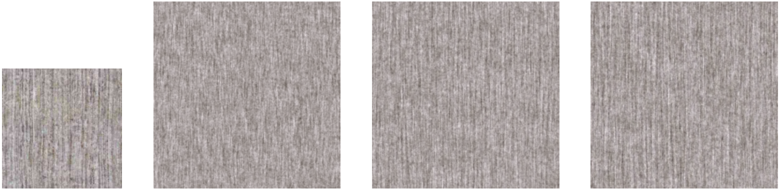

u (51×51) | $\mathrm{ADSN}(t^{11})$ RME = 0.56 | $\mathrm{ADSN}(t^{31})$ RME = 0.49 | $\mathrm{ADSN}(t^{51})$ RME = 0.43

From left to right, a Gaussian texture, and samples of the models obtained with SOTs t^r with circular supports of radius $r \in \{5, 15, 25\}$.

Conclusion

- The SOT is a very compact summary of a Gaussian texture (solution of inverse texture synthesis [Wei et al., 2008] in the Gaussian case).
- With the SOT, the "visual convergence" of the DSN is very fast. We get a satisfying synthesis with **only 50 impacts per pixel!**
- The DSN is naturally defined on $\mathbb{Z}^2$ and allows for parallel evaluation.
- The DSN associated to the SOT is thus a **fast** and **flexible** method for Gaussian texture synthesis, with a **small memory footprint**.

Our paper, codes and several examples are available on my webpage:

`http://www.math-info.univ-paris5.fr/~aleclair/sot/`

Questions and perspectives:

- How to measure the DSN visual convergence towards the ADSN?
- Continuous version of the SOT... for procedural texture synthesis.
- Dual approach: compact approximation of the covariance inverse.

A. Desolneux, L. Moisan, and S. Ronsin.
A compact representation of random phase and Gaussian textures.
In *ICASSP'12*, pages 1381–1384.

B. Galerne, Y. Gousseau, and J.M. Morel.
Random Phase Textures: Theory and Synthesis.
IEEE TIP, 20:1:257–267, 2011.

M.H. Hayes.
The Reconstruction of a Multidimensional Sequence from the Phase or Magnitude of its Fourier Transform.
IEEE Trans. on Acoustics, Speech, Sign. Proc., 30:2:140–154, 1982.

A. Lagae, S. Lefebvre, G. Drettakis, and P. Dutré.
Procedural Noise Using Sparse Gabor Convolution.
In *SIGGRAPH'09*, volume 28:3, pages 54:1–54:10.

L. Moisan.
Periodic plus smooth image decomposition.
Journal of Mathematical Imaging and Vision, 39(2):161–179, 2011.

J. J. van Wijk.
Spot noise texture synthesis for data visualization.
In *proc. SIGGRAPH'91*, volume 25, pages 309–318, 1991.

L.-Y. Wei, J. Han, K. Zhou, H. Bao, B. Guo, and H.-Y. Shum.
Inverse texture synthesis.
In *ACM ToG*, volume 27, 2008.

L.-Y. Wei, S. Lefebvre, V. Kwatra, and G. Turk.
State of the art in example-based texture synthesis.
In *Eurographics 2009, State of the Art Report, EG-STAR*, pages 93–117, 2009.

G.-S. Xia, S. Ferradans, G. Peyré, and J.-F. Aujol.
Synthesizing and Mixing Stationary Gaussian Texture Models.
to appear in SIAM J. on Imaging Science, 2013.

G. Xia, S. Ferradans, G. Peyré, and J. Aujol.
Synthesizing and Mixing Stationary Gaussian Texture Models.
SIAM Journal on Imaging Sciences, 7(1):476–508, 2014.

Thanks!

THANK YOU FOR YOUR ATTENTION