Complétion d'Images de Textures

Arthur Leclaire
Institut de Mathématiques de Bordeaux

Joint work with
Bruno Galerne (IDP, Orléans), Julien Rabin (GREYC, Caen)

Congrès SMAI
Vendredi 25 Juin 2021

Motivation

Inpainting consists in **filling missing regions of an image**.

This problem has been addressed with

- variational/PDE-based methods
  [Masnou & Morel, 1998], [Bertalmio et al.,2000], [Chan & Shen, 2001], [Tschumperlé et al., 2006] ...
- stochastic/exemplar-based methods
  [Igehy Pereira, 1997], [Efros Leung, 1999], [Criminisi et al., 2004], [Wexler et al., 2007] ...
- hybrid methods (structure+texture e.g.)
  [Bertalmio et al., 2003], [Elad et al., 2005], [Cao et al., 2011]...
- variational exemplar-based methods
  [Aujol et al., 2010], [Arias et al., 2011]

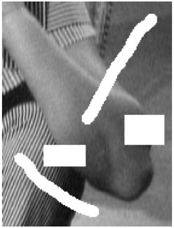

Textural Inpainting by Conditional Simulation

- In the case of random texture models, inpainting can be formulated as **conditional simulation**.
- **Notation:**

$\Omega \subset \mathbb{Z}^2$ is a discrete rectangle.

F is a random texture model on Ω.

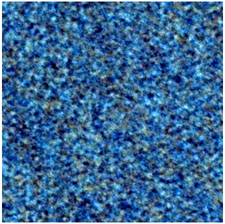

Textural Inpainting by Conditional Simulation

- In the case of random texture models, inpainting can be formulated as **conditional simulation**.
- **Notation:**

$\Omega \subset \mathbb{Z}^2$ is a discrete rectangle.

F is a random texture model on Ω.

$M \subset \Omega$ is a mask.

The values $u(x)$ are known for $x \in \Omega \setminus M$.

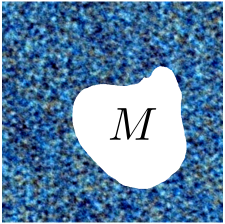

Textural Inpainting by Conditional Simulation

- In the case of random texture models, inpainting can be formulated as **conditional simulation**.
- **Notation:**

$\Omega \subset \mathbb{Z}^2$ is a discrete rectangle.

F is a random texture model on Ω.

$M \subset \Omega$ is a mask.

The values $u(x)$ are known for $x \in \Omega \setminus M$.

$\mathcal{C} \subset \Omega \setminus M$ is a set of conditioning points.

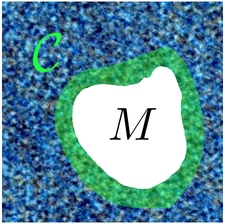

Textural Inpainting by Conditional Simulation

- In the case of random texture models, inpainting can be formulated as **conditional simulation**.
- **Notation:**

$\Omega \subset \mathbb{Z}^2$ is a discrete rectangle.

F is a random texture model on Ω.

$M \subset \Omega$ is a mask.

The values $u(x)$ are known for $x \in \Omega \setminus M$.

$\mathcal{C} \subset \Omega \setminus M$ is a set of conditioning points.

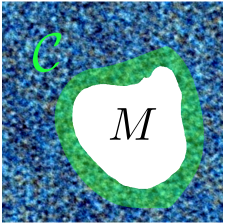

- **Main idea: Sample the conditional distribution of F knowing $F_{|\mathcal{C}} = u_{|\mathcal{C}}$.**

If F is a Gaussian model, this can be done perfectly.

By-example Synthesis of Microtextures

Goal: Synthesize an exemplar microtexture $u : \Omega \to \mathbb{R}^d$.

→ We estimate the mean value by $\bar{u} = \frac{1}{|\Omega|} \sum_{x \in \Omega} u(x)$.

→ We consider the "normalized spot" $t_u(x) = \frac{1}{\sqrt{|\Omega|}}(u(x) - \bar{u})\mathbf{1}_{x\in\Omega}, \ (x \in \mathbb{Z}^2)$.

→ We sample the Gaussian model

$$\bar{u} + t_u * W(x) = \bar{u} + \sum_{y \in \mathbb{Z}^2} W(y) t_u(x-y), \quad (x \in \Omega),$$

where W is a normalized Gaussian white noise on $\mathbb{Z}^2$ ($W(x) \sim \mathcal{N}(0,1)$).

Original u

Synthesis $\bar{u} + t_u * W$

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

A Precise Model for Microtextures

- $F = t_u * W$ has zero mean and covariance function

$$\mathbb{E}(F(x)F(y)^T) = t_u * \tilde{t}_u^T(x-y) = \frac{1}{|\Omega|}\sum_z (u(z)-\bar{u})(u(y-x+z)-\bar{u})^T,$$

where $\tilde{t}_u(x) = t_u(-x)$.

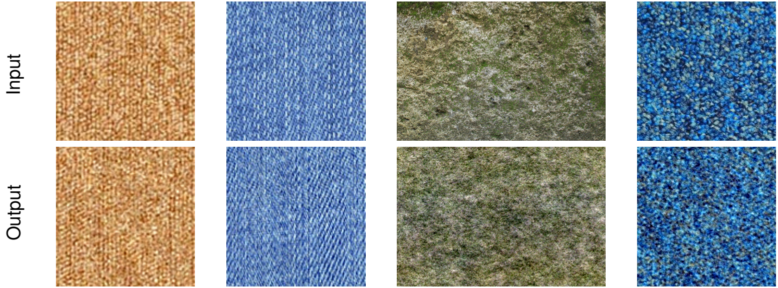

- The convolutions can be computed efficiently with the FFT.
- Technical detail: In order to avoid potential directional artifacts, the border discontinuity of t_u can be attenuated by a smooth window [Galerne et al., 2011].

Gaussian conditional simulation

- Let $(F(x))_{x\in\Omega}$ be a Gaussian vector **with mean zero** and covariance

$$\Gamma(x,y) = \mathrm{Cov}(F(x),F(y)) = \mathbb{E}(F(x)F(y)), \quad x,y\in\Omega.$$

- There exists $(\lambda_c(x))_{c\in\mathcal{C}}$ such that

$$\mathbb{E}(\ F(x) \mid F(c)\ ,\ c\in\mathcal{C}\ ) = \sum_{c\in\mathcal{C}} \lambda_c(x)F(c).$$

- The **simple kriging estimation** is defined by $F^*(x) = \sum_{c\in\mathcal{C}} \lambda_c(x)F(c)$.

Theorem: F^* and $F - F^*$ are independent. [Lantuéjoul, 2002]

Consequence: A conditional sample of F given $F_{|\mathcal{C}} = \varphi$ can be obtained as

$$F \mid F_{|\mathcal{C}} = \varphi \quad \sim \quad \underbrace{\varphi^*}_{\textbf{Kriging component}} + \underbrace{F - F^*}_{\textbf{Innovation component}}.$$

- The **kriging coefficients** $\Lambda = (\lambda_c(x))_{\substack{x\in\Omega\\ c\in\mathcal{C}}}$ satisfy $\Gamma_{|\Omega\times\mathcal{C}} = \Lambda\Gamma_{|\mathcal{C}\times\mathcal{C}}$.

- When $\Gamma_{|\mathcal{C}\times\mathcal{C}}$ is non-singular, $\boxed{\Lambda = \Gamma_{|\Omega\times\mathcal{C}}\Gamma^{-1}_{|\mathcal{C}\times\mathcal{C}}.}$

Application with a Gaussian texture model

We observe a texture $u : \Omega \to \mathbb{R}$ outside a mask $M \subset \Omega$.

We take $\mathcal{C}$ as the outside border of M.

$\to$ **Estimate a Gaussian texture model** on a subdomain $\omega \subset \Omega \setminus M$ by

$$v = u_{|\omega} \ , \quad \bar{v} = \frac{1}{|\omega|} \sum_{x \in \omega} v(x) \ , \quad t_v = \frac{1}{\sqrt{|\omega|}} (v(x) - \bar{v}) \mathbf{1}_{x \in \omega}$$

The Gaussian model is $\bar{v} + F$ where $F = t_v * W$; $\Gamma(x, y) = t_v * \tilde{t}_v(x - y)$

$\to$ **Draw a conditional sample** of $\bar{v} + F$ given $F_{|\mathcal{C}} = u_{|\mathcal{C}} - \bar{v}$, i.e.

$$\bar{v} + (u - \bar{v})^* + F - F^* \quad \text{where} \quad \varphi^* = \Lambda(\varphi_{|\mathcal{C}}) = \Gamma_{|\Omega \times \mathcal{C}} \Gamma^{-1}_{|\mathcal{C} \times \mathcal{C}} \varphi_{|\mathcal{C}}$$

u $\quad F \mid F_{|\mathcal{C}} = u_{|\mathcal{C}} - \bar{v} \quad \sim \quad (u_{|\mathcal{C}} - \bar{v})^* \quad + \quad F - F^*$

Illustration

The texture model is estimated on the masked exemplar.

Original

Illustration

The texture model is estimated on the masked exemplar.


Original

Masked input

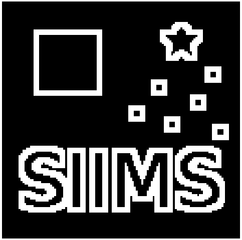

Conditioning set

Illustration

The texture model is estimated on the masked exemplar.

Original

Masked input

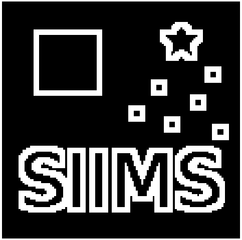

Conditioning set

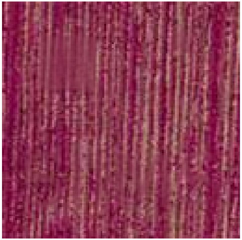

Kriging component
$\bar{v} + (u - \bar{v})^*$

Illustration

The texture model is estimated on the masked exemplar.

Original

Masked input

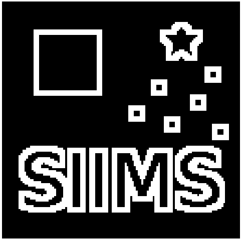

Conditioning set

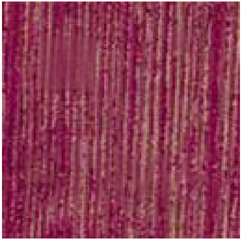

Kriging component
$\bar{v} + (u - \bar{v})^*$

Innovation component
$\bar{v} + F - F^*$

Illustration

The texture model is estimated on the masked exemplar.

Original

Masked input

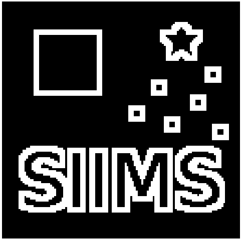

Conditioning set

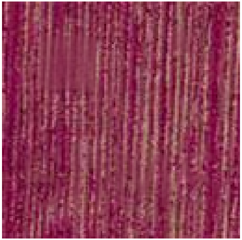

Kriging component
$\bar{v} + (u - \bar{v})^*$

Innovation component
$\bar{v} + F - F^*$

Inpainted result
$\bar{v} + (u - \bar{v})^* + F - F^*$

Inpainting Results

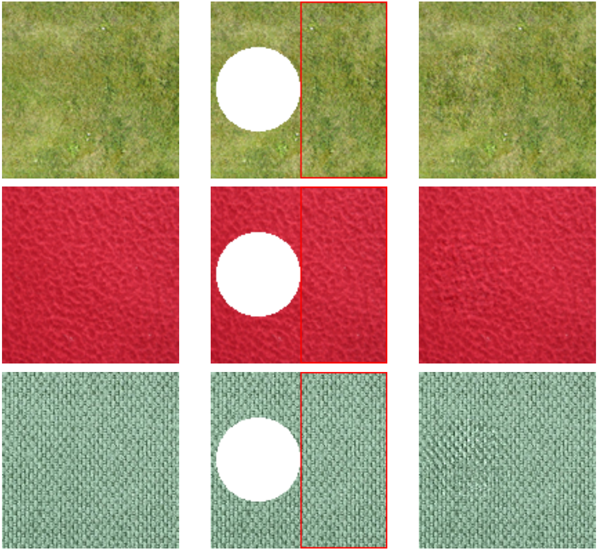

The Gaussian model is estimated on the right part ω of the image.

Comparison (I)

Original

TV inpainting
[Chan & Shen, 2002]

Our result

[Criminisi et al, 2004]

Comparison (II)

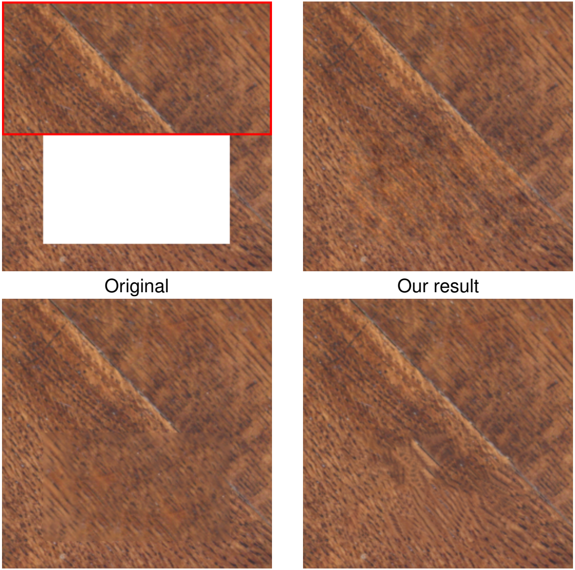

Original | Our result | [Arias et al, 2011] | [Daisy et al, 2015]

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

A Model for Structured Textures

One can enrich the model by applying **optimal transport maps** in the **patch space**.
$\longrightarrow$ Model **TexTO** (Textures with Optimal Transport)

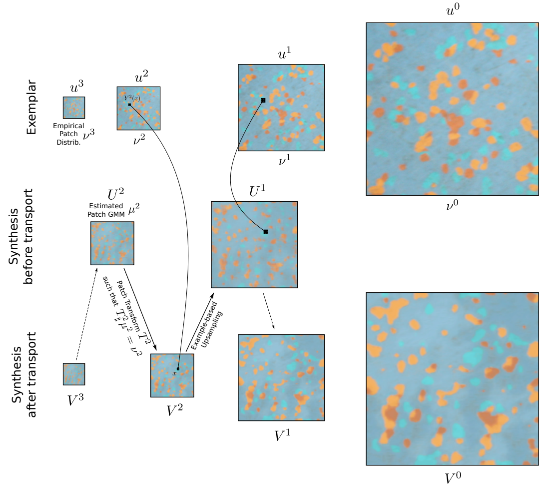

Texto Model

- Compute exemplar at L resolutions: u^0, u^1, ..., u^{L-1}.
- Start with Gaussian model U^{L-1} at scale $L-1$.
- For $\ell = L-1, \dots, 0$,
  - On all patches (p_i) of U^ℓ, apply a well-chosen transform

    $$T_\psi(p_i) = \operatorname*{Argmin}_{q_j} \|p_i - q_j\|^2 - \psi_j$$

    where (q_j) are the patches of u^ℓ, and (ψ_j) mildly adjusted so that

    $$T_\psi \sharp \mu^\ell = \nu^\ell \ .$$

  - Recompose an image V^ℓ by average of the transformed patches $(T_\psi(p_i))$.
  - Perform an exemplar-based upsampling $U^{\ell-1}$ to initialize the next scale.

REMARK: This model can be used for inpainting with

1. Gaussian conditional simulation at coarse scale,
2. using patches (q_j) only outside the mask at each scale.

Texto Inpainting

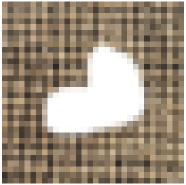

$s = 4$, Masked

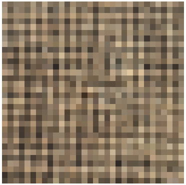

$s = 4$, Gaussian

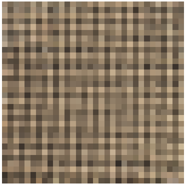

$s = 4$

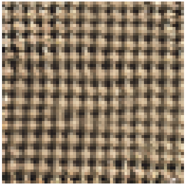

$s = 3$

$\downarrow$

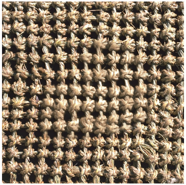

Output

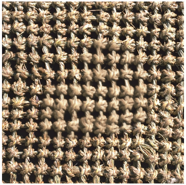

$s = 0$

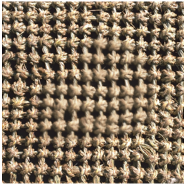

$s = 1$

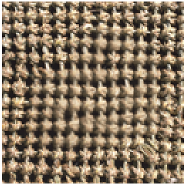

$s = 2$

Texto Inpainting

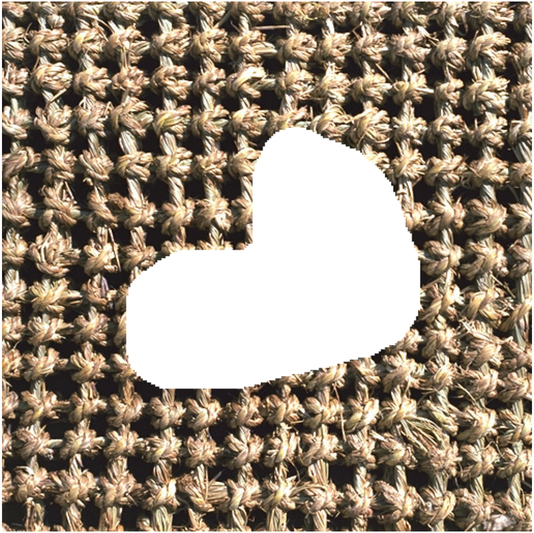

Masked

Texto Inpainting

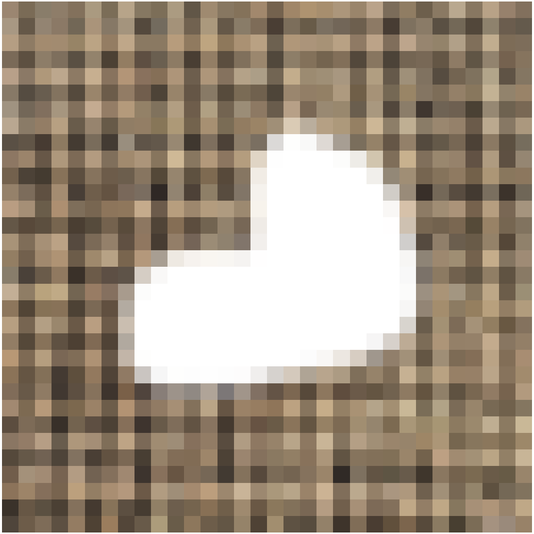

$s = 4$, Masked

Texto Inpainting

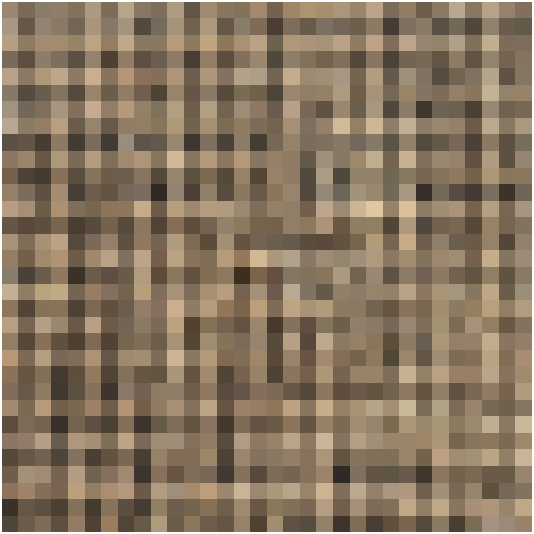

$s = 4$, Gaussian

Texto Inpainting

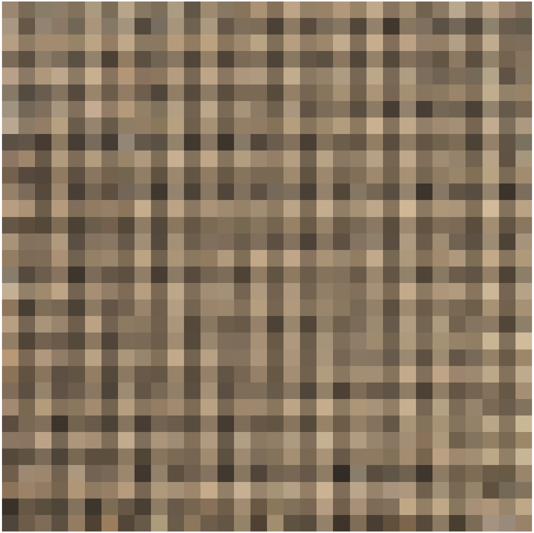

$s = 4$

Texto Inpainting

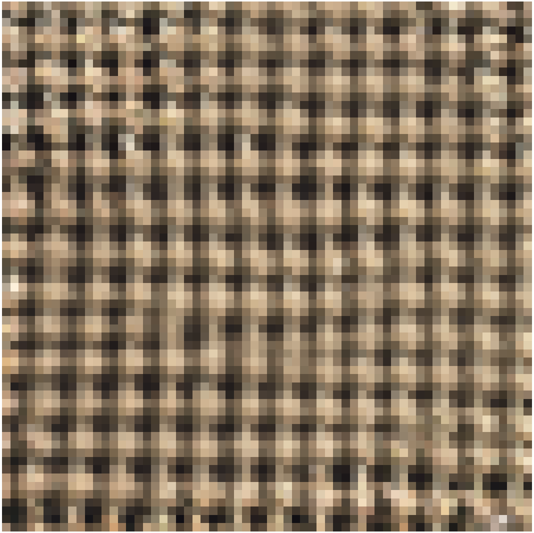

$s = 3$

Texto Inpainting

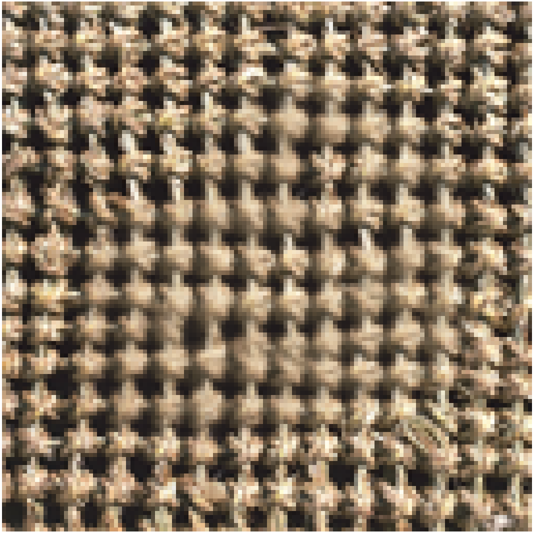

$s = 2$

Texto Inpainting

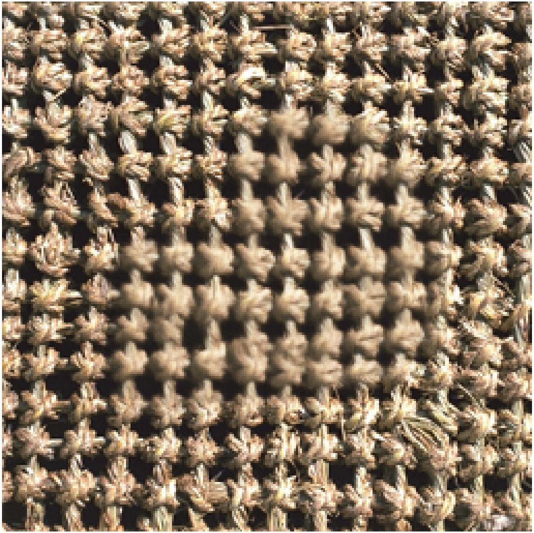

$s = 1$

Texto Inpainting

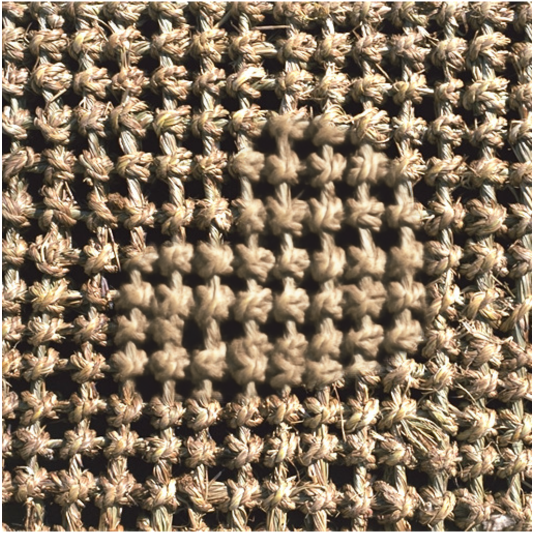

$s = 0$

Texto Inpainting

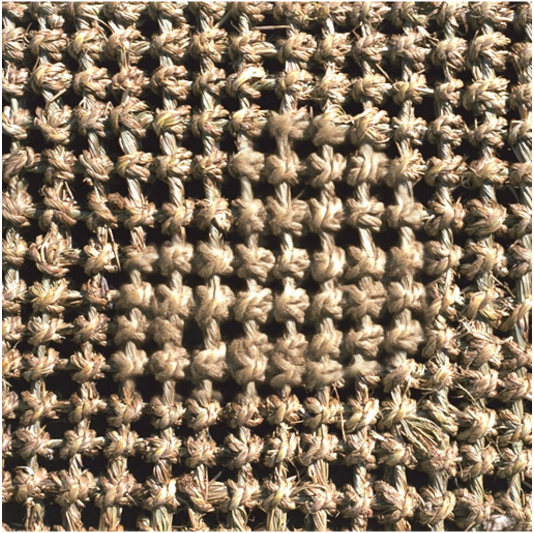

Inpainted

TexTo Inpainting - Results

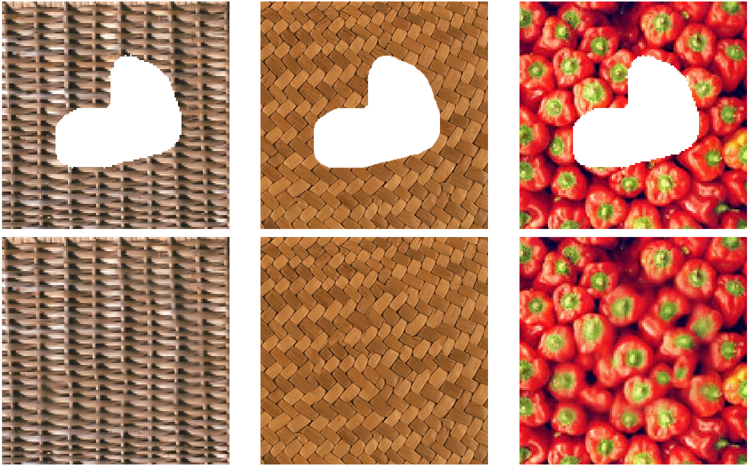

TexTo Inpainting - Results

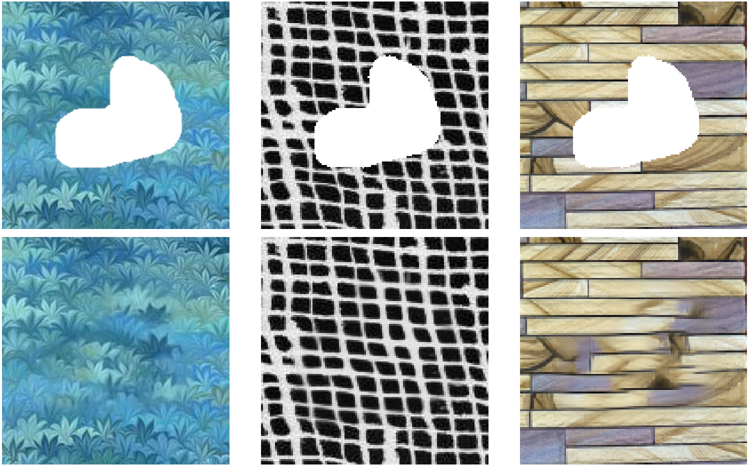

Texto Inpainting - Comparison

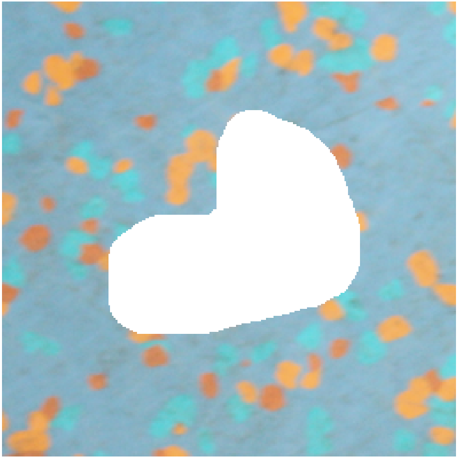

Masked Texture (256 $\times$ 256)

Texto Inpainting - Comparison

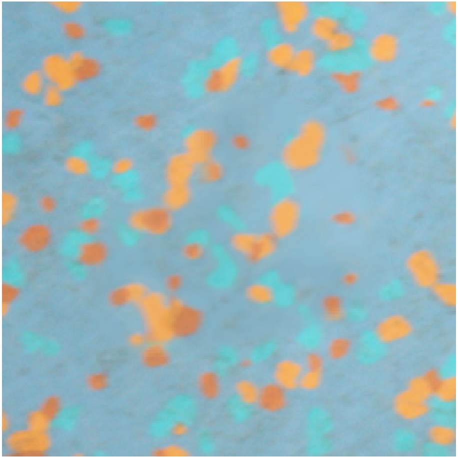

Texto

Texto Inpainting - Comparison

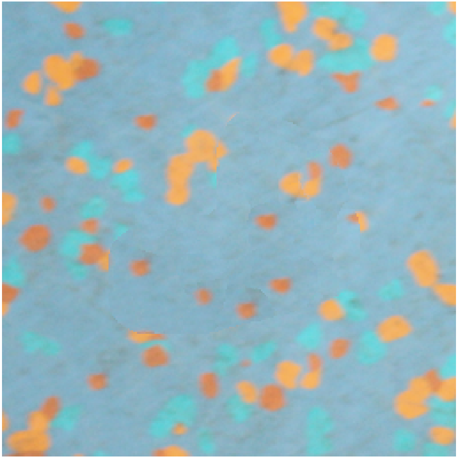

[Newson et al., 2014]

Texto Inpainting - Comparison

Masked Texture (512×512)

Texto Inpainting - Comparison

Texto

Texto Inpainting - Comparison

[Newson et al., 2014]

Conclusion

- Microtexture inpaiting can be addressed with perfect Gaussian conditional simulation.
- – Gaussian simulation is limited to stationary Gaussian textures.
- \+ It is guaranteed to respect the Gaussian texture model.
- \+ It can inpaint holes of any shape and size in a reasonable time.
- The texture model can be enriched with patch-based optimal transport.

OUR PAPERS, SOURCE CODES AND A TUTORIAL are available on my website
`https://www.math.u-bordeaux.fr/~aleclaire/gaussian_inpainting/`

See also the online demo of Gaussian inpainting at
`https://www.ipol.im/pub/art/2017/198/`

THANK YOU FOR YOUR ATTENTION!

Comparison (III)

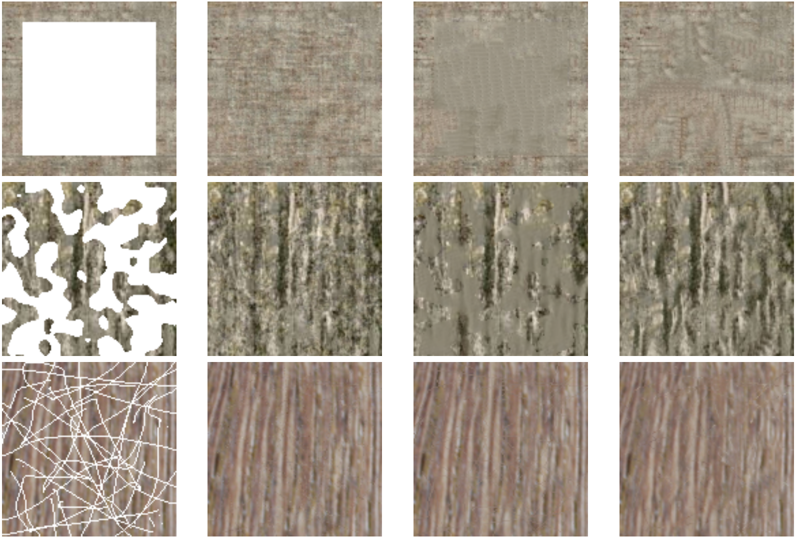

Input | Our method | [Arias et al., 2011] | [Daisy et al., 2015]

Comparison (IV)

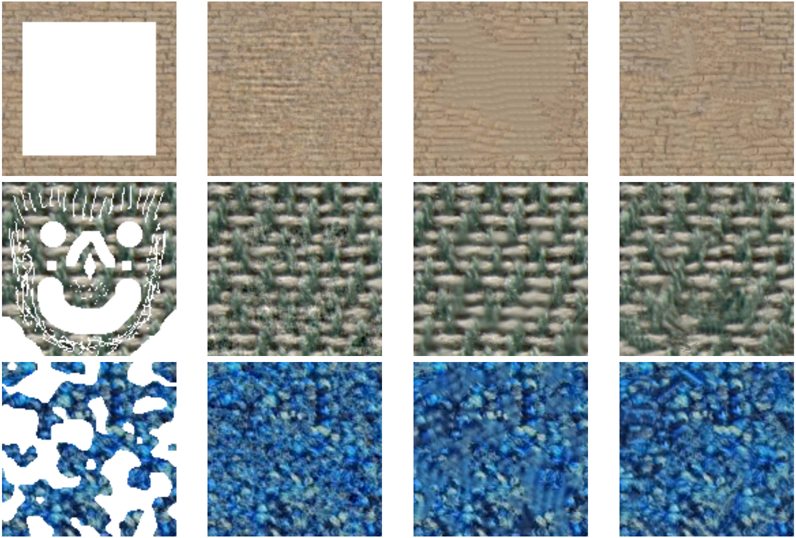

Input | Our method | [Arias et al., 2011] | [Daisy et al., 2015]

Inpainting Composite Textures

Original

Inpainted

- Limitation: the estimated Gaussian model is stationary.

Inpainting Composite Textures

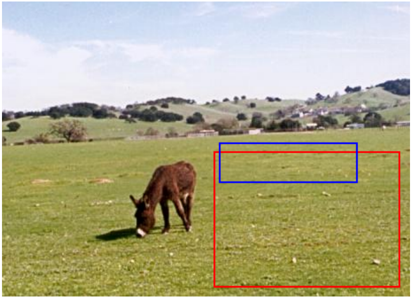
Original


Inpainted with two ADSN models

- Limitation: the estimated Gaussian model is stationary.

Texto Inpainting - Comparison

Masked texture

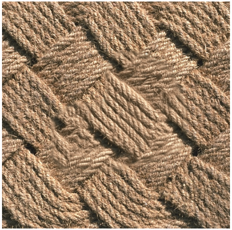

Texto

[Ulyanov et al. 2018]

[Daisy et al., 2015]

Texto Inpainting - Comparison

Masked texture

Texto Inpainting - Comparison

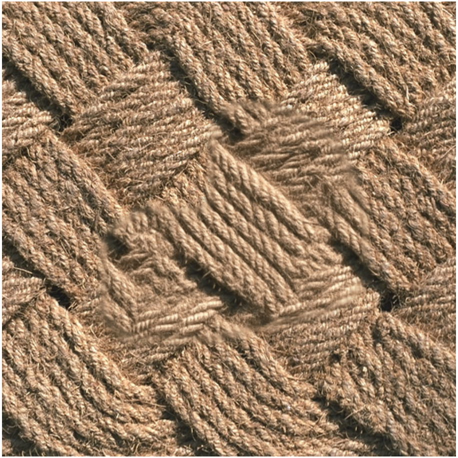

Texto

Texto Inpainting - Comparison

[Ulyanov et al. 2018]

Texto Inpainting - Comparison

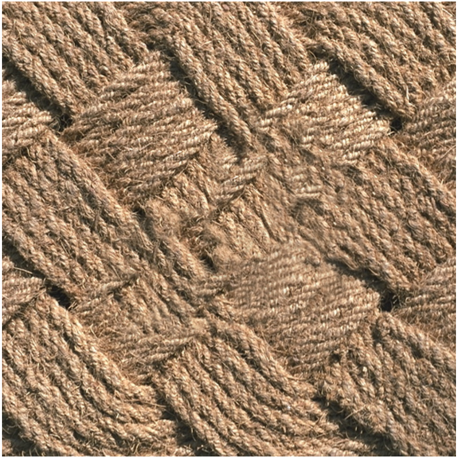

[Daisy et al., 2015]