Microtexture Inpainting through Gaussian Conditional Simulation

Arthur Leclaire[1]
Joint work with Bruno Galerne[2] and Lionel Moisan[2]

[1] ENS Cachan
CMLA, CNRS UMR 8536

[2] Université Paris Descartes
MAP5, CNRS UMR 8145

ICASSP16
Thursday, March 24th, 2016

This work has been partially funded by the French Research Agency under grant ANR-14-CE27-001 (MIRIAM).

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
  [Bertalmio et al., 2003], [Elad et al., 2005], [Cao et al., 2011]...
- variational exemplar-based methods
  [Aujol et al., 2010], [Arias et al., 2011]

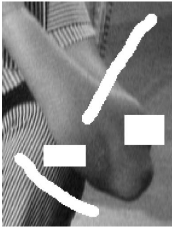

Textural Inpainting by Conditional Simulation

- 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.

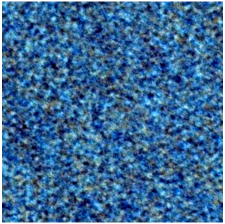

Textural Inpainting by Conditional Simulation

- 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 $x \in \Omega \setminus M$.

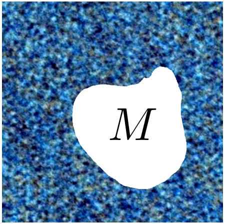

Textural Inpainting by Conditional Simulation

- 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.

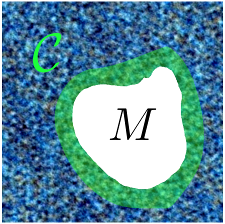

Textural Inpainting by Conditional Simulation

- 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.

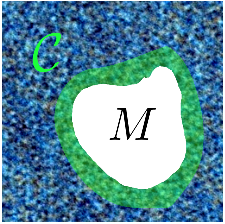

- **Main idea: Sample the conditional distribution of F knowing $F_{|\mathcal{C}} = u_{|\mathcal{C}}$.**

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

Outline

By-example Synthesis of Microtextures

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

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

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

$\to$ 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

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

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

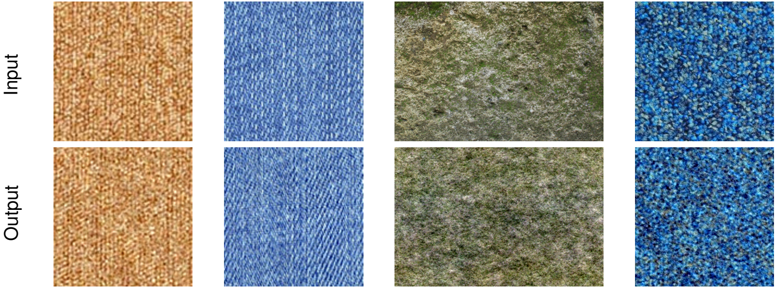

- The convolutions can be computed efficiently with the FFT.
- Technical detail: In order to avoid potential directional artifacts, the border discontinuity of t_u can be attenuated by a smooth window [Galerne et al., 2011].

Outline

Microtexture Synthesis by Example

Gaussian Texture Inpainting

Results

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

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

- 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.

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

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

Algorithm

Trick: With our Gaussian model, $\Gamma_{|\Omega\times\mathcal{C}}\psi$ can be computed efficiently by convolving the zero-padding extension of ψ by $t_v * \tilde{t}_v$.

Color extension: Three values for each pixel $\rightarrow$ Replace Ω by $\Omega \times \{R, G, B\}$.

Algorithm: Microtexture inpainting

Input: Mask $M \subset \Omega$, texture u on $\Omega \setminus M$

Define $\mathcal{C}$ as a border of M with width $w = 3$ pixels

- From a restriction v of u to $\omega \subset \Omega \setminus M$, compute $\bar{v}$, t_v, and $t_v * \tilde{t}_v$
- Store the matrix $\Gamma_{|\mathcal{C}\times\mathcal{C}}(c, d) = t_v * \tilde{t}_v(c - d),\ c, d \in \mathcal{C}$
- Draw a Gaussian sample $F = t_v * W$
- Compute $\psi_1 = \Gamma^{-1}_{|\mathcal{C}\times\mathcal{C}}(u_{|\mathcal{C}} - \bar{v})$ and $\psi_2 = \Gamma^{-1}_{|\mathcal{C}\times\mathcal{C}}F_{|\mathcal{C}}$
- Extend ψ_1 and ψ_2 by zero-padding to get Ψ_1 and Ψ_2
- Compute $(u - \bar{v})^* = t_v * \tilde{t}_v * \Psi_1$, $\quad F^* = t_v * \tilde{t}_v * \Psi_2$

Output: $\bar{v} + (u - \bar{v})^* + F - F^*$

Overall complexity: $\mathcal{O}(|\mathcal{C}|^3 + |\Omega| \log |\Omega|)$.

Outline

Validation

First, we inpaint a texture with an oracle Gaussian model.

Original

Validation

First, we inpaint a texture with an oracle Gaussian model.


Original

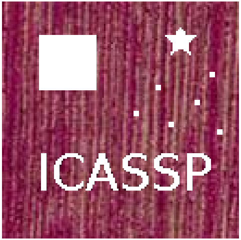

Masked input

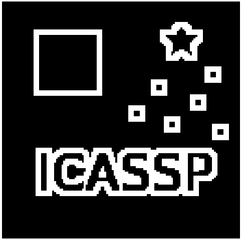

Conditioning set

Validation

First, we inpaint a texture with an oracle Gaussian model.

Original

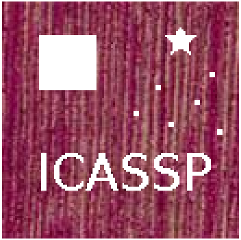

Masked input

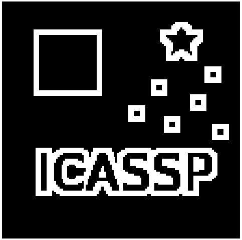

Conditioning set

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

Validation

First, we inpaint a texture with an oracle Gaussian model.

Original

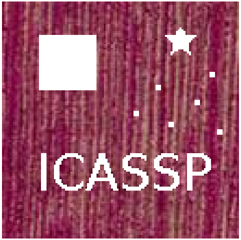

Masked input

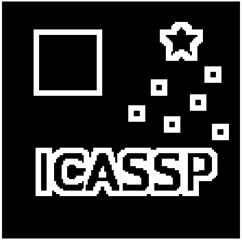

Conditioning set

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

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

Validation

First, we inpaint a texture with an oracle Gaussian model.

Original

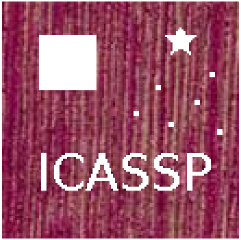

Masked input

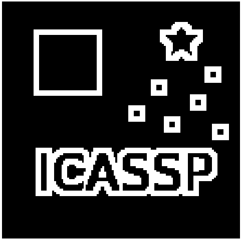

Conditioning set

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

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

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

Inpainting Results

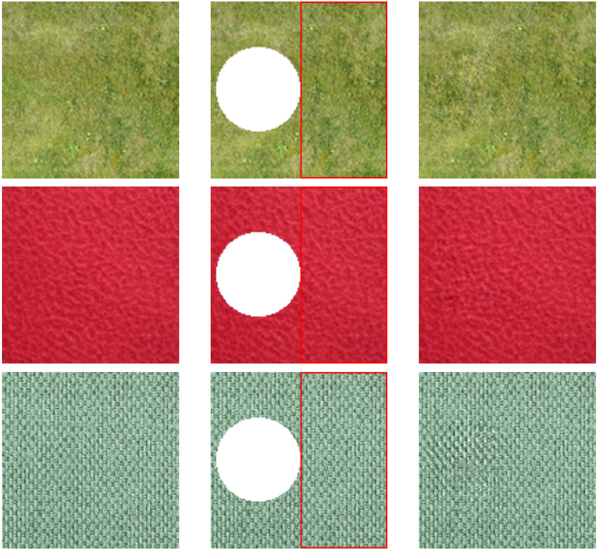

The Gaussian model is estimated on the right part ω of the image.

Comparison (I)

Original

Kriging component

[Efros & Leung, 1999]

Our result

Comparison (II)

Original

TV inpainting
[Chan & Shen, 2002]

Our result

[Criminisi et al, 2004]

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

Inpaint microtexture parts of an image

Inpainting Composite Textures

Original

Inpainted

- Limitation: the estimated Gaussian model is stationary.

Inpainting Composite Textures

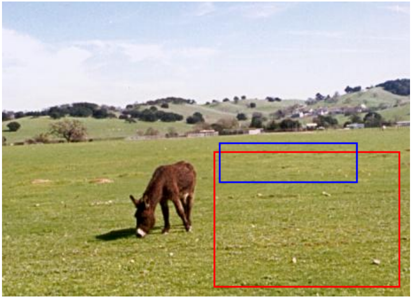

Original

Inpainted with two ADSN models

- Limitation: the estimated Gaussian model is stationary.

Conclusion

CONCLUSIONS:

- Microtexture inpaiting can be addressed with perfect Gaussian conditional simulation.

– Limited to stationary Gaussian textures.

+ Guaranteed to respect the Gaussian texture model.

+ Can inpaint holes of any shape and size in a reasonable time.

PERSPECTIVES:

- Solve more efficiently the linear systems involving $\Gamma_{|\mathcal{C}\times\mathcal{C}}$.
- Use a more involved procedure for the estimation of a Gaussian model on a masked exemplar.
- Extend the conditional simulation to a non-stationary case
  $\rightarrow$ inpainting of non-texture images.

SOURCE CODES AND A TUTORIAL are available on my website
`www.math-info.univ-paris5.fr/~aleclair/gaussian_inpainting/`

THANK YOU FOR YOUR ATTENTION!