M. Daisy, P. Buyssens, O. Lézoray and D. Tschumperlé

19 Avril 2016

Technicolor Workshop on Multimedia Inpainting

Image inpainting with a single-scale approach: From still images to stereoscopic videos.

GREYC
CNRS
ENSICAEN École Nationale Supérieure d'Ingénieurs de Caen
Normandie Université

Why is it a tough problem ?

Why is it a tough problem ?

Why is it a tough problem ?

Why is it a tough problem ?

Doing it manually : Clone tool

Why is it a tough problem ?

Doing it manually with the *Clone* tool

Why is it a tough problem ?

Doing it manually with the *Clone* tool

Why is it a tough problem ?

Doing it manually with the *Clone* tool

Why is it a tough problem ?

Doing it manually with the *Clone* tool

Why is it a tough problem ?

Doing it manually with the *Clone* tool

Why is it a tough problem ?

[Automatic reconstruction by Criminisi, Pérez et Toyama '03]

Why is it a tough problem ?

[Wexler et al. '07 (multi-scale approach)]

A non-exhaustive overview of inpainting approaches

A non-exhaustive overview of inpainting approaches

- Purely geometry-driven methods

[Masnou et Morel '98] - *"Level line based disocclusion"*

Idea : Link together the isophotes external to the inpainting mask

Minimization of an energy functional using dynamic programming

A non-exhaustive overview of inpainting approaches

- Purely geometry-driven methods

[Bertalmio et al. '00] - "Image inpainting"

Idea : Transport equation to propagate isophotes

A non-exhaustive overview of inpainting approaches

- Purely geometry-driven methods

[Chan et al. '01] - "Non-Texture Inpainting by Curvature-Driven Diffusions"

Idea : Diffusion equation guided by the isophotes directions/curvatures

A non-exhaustive overview of inpainting approaches

- Purely geometry-driven methods

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

[Efros and Leung '99] - *"Texture synthesis by non-parametric sampling"*

Idea : Texture synthesis pixel by pixel, exploiting texture auto-similarities

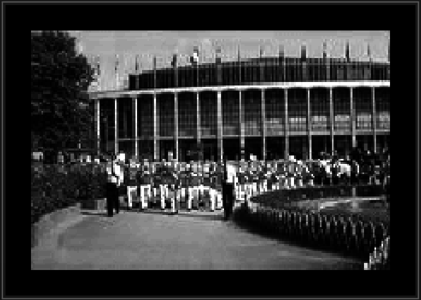

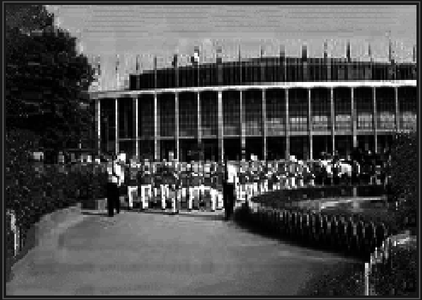

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

[Criminisi et al. '03] - "Region Filling and Object Removal by Exemplar-Based Image Inpainting"

Idea : Mixing texture synthesis and local analysis of the image geometry.

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

[Criminisi et al. 2003], "Region Filling and Object Removal by Exemplar-Based Image Inpainting"

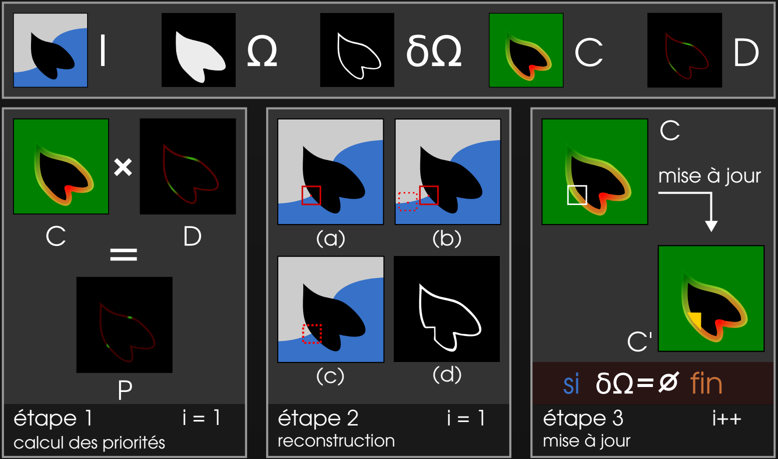

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

[Criminisi et al. 2003], "Region Filling and Object Removal by Exemplar-Based Image Inpainting"

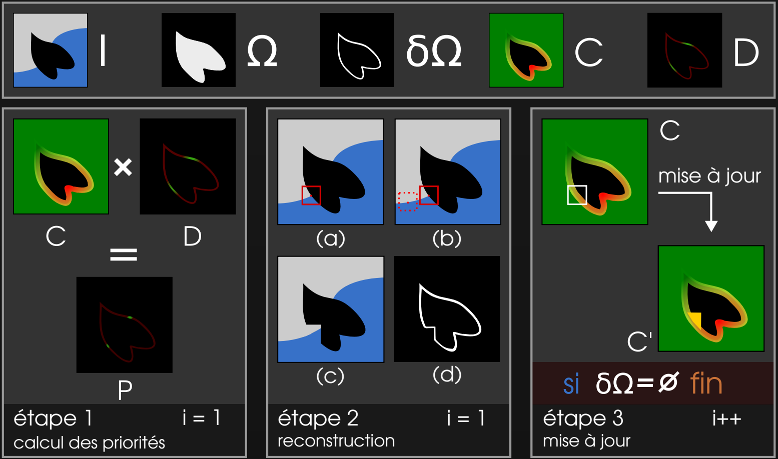

- Patch/texture-based methods

[Criminisi et al. 2003], "Region Filling and Object Removal by Exemplar-Based Image Inpainting"

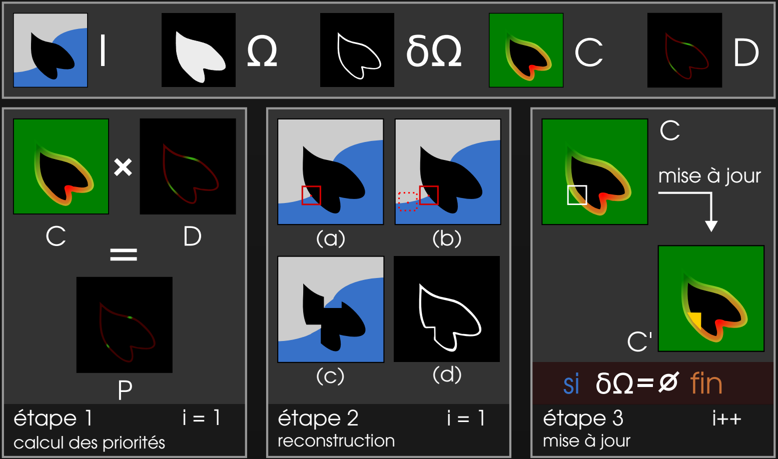

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

[Lemeur et al. '11] - "Examplar-based inpainting based on local geometry"

Idea : Improve *[Criminisi et al '03]* with a better geometry analysis (tensor-based) and a better patch synthesis (means of K best patches).

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

PatchMatch *[Barnes et al. '09]*

- A very fast algorithm for matching similar patchs between two images
- Works for *full* images (non-masked patchs)

Used in several efficient image and video inpainting techniques *(multi-scale approaches)* :

- *[Wexler et al. '07]* + PatchMatch ▶ Photoshop
- *[Newson et al. '14]* ▶ PatchMatch 3D

▶ Solving the problem for *different image scales*, using the solution at one scale as an initialization for the upper scale (*may increase error propagation !*)

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

2 key steps with space for improvements

1. The analysis of the image geometry
2. The synthesis of the best patchs to copy/paste

- Problems with curved structures to reconstruct
- Visible bloc effects

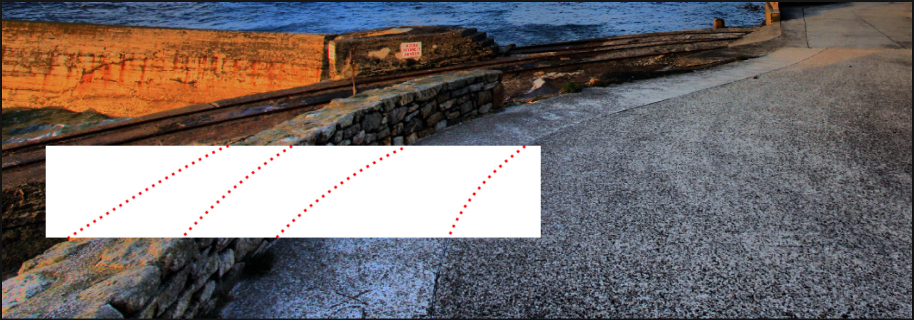

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

2 key steps with space for improvements

1. The analysis of the image geometry
2. The synthesis of the best patchs to copy/paste

- Problems with curved structures to reconstruct
- Visible bloc effects

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

2 key steps with space for improvements

1. The analysis of the image geometry
2. The synthesis of the best patchs to copy/paste

- Problems with curved structures to reconstruct
- Visible bloc effects

A non-exhaustive overview of inpainting approaches

- Patch/texture-based methods

2 key steps with space for improvements

1. The analysis of the image geometry
2. The synthesis of the best patchs to copy/paste

- Problems with curved structures to reconstruct
- Visible bloc effects

Contributions

A few improvements to
[Criminisi, Pérez et Toyama ’03]

Single-scale inpainting : data term

$$D_p = \frac{|\overrightarrow{\nabla I_p}^{\perp} \,.\, \overrightarrow{n_p}|}{\alpha} \qquad (1)$$

"The gradient ∇I_p is computed as the maximum value of the image gradient in $\Psi_p \cap I$."

$$\overrightarrow{\nabla I}_p^{\perp} = \{\overrightarrow{\nabla I}_q^{\perp} \mid \operatorname*{arg\,max}_{q \in ((\mathcal{I}-\Omega)\cap\psi_p)} \|\overrightarrow{\nabla I}_q\|\} \qquad (2)$$

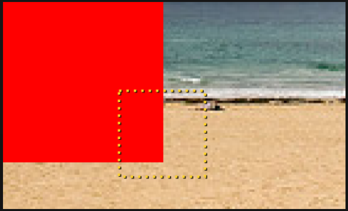

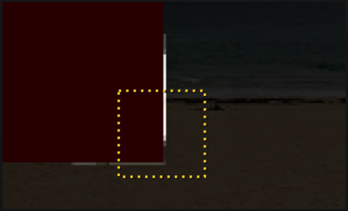

Single-scale inpainting : data term

$$D_p = \frac{|\overrightarrow{\nabla I_p}^{\perp} \,.\, \overrightarrow{n_p}|}{\alpha} \tag{1}$$

"The gradient ∇I_p is computed as the maximum value of the image gradient in $\Psi_p \cap I$."

$$\overrightarrow{\nabla I}_p^{\perp} = \{\overrightarrow{\nabla I}_q^{\perp} \mid \underset{q \in ((\mathcal{I}-\Omega)\cap\psi_p)}{\arg\max} \|\overrightarrow{\nabla I}_q\|\} \tag{2}$$

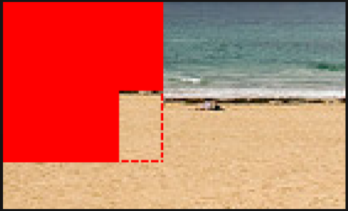

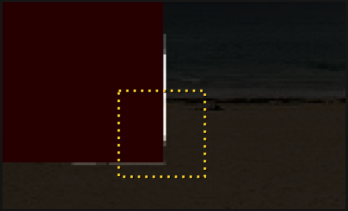

A better data term

Data term based on the smoothed structure tensor field

$$\tilde{D}_p = \|\mathbf{G}_p \ . \ \overrightarrow{n}_p\| \tag{3}$$

with

$$\mathbf{G}_p = \sum_{q \in \mathcal{N}_p \cap (\mathcal{I} - \Omega)} w_p(q) \ \overrightarrow{\nabla I}_q \overrightarrow{\nabla I}_q^T \tag{4}$$

w_p : 2d gaussian function, centered on p.

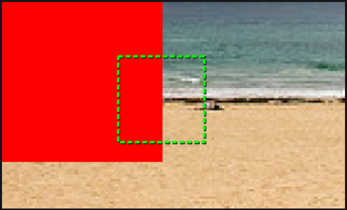

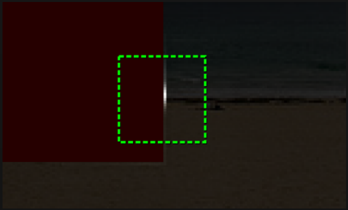

A better data term

Data term based on the smoothed structure tensor field

$$\tilde{D}_p = \|\mathbf{G}_p \ . \ \overrightarrow{n}_p\| \tag{3}$$

with

$$\mathbf{G}_p = \sum_{q \in \mathcal{N}_p \cap (\mathcal{I} - \Omega)} w_p(q) \ \overrightarrow{\nabla I}_q \overrightarrow{\nabla I}_q^T \tag{4}$$

w_p : 2d gaussian function, centered on p.

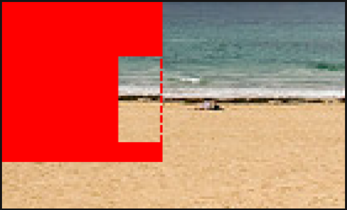

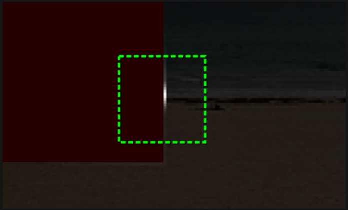

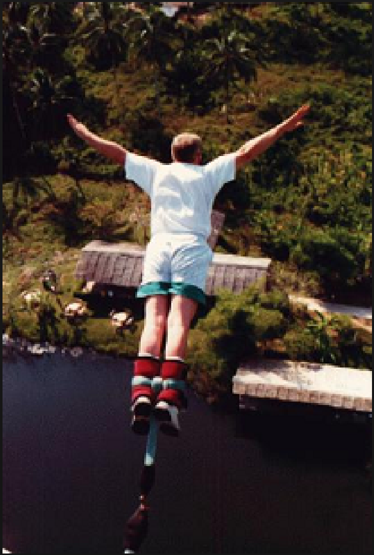

Original image

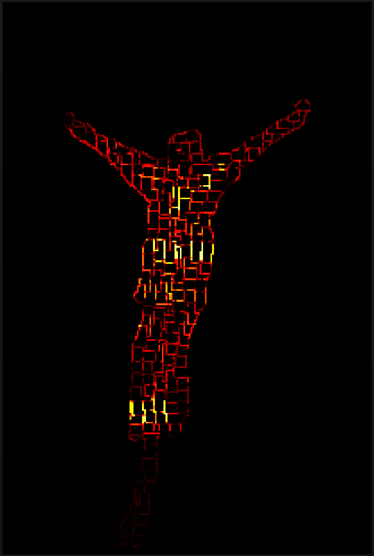

Original data term from *[Criminisi et al. '03]*

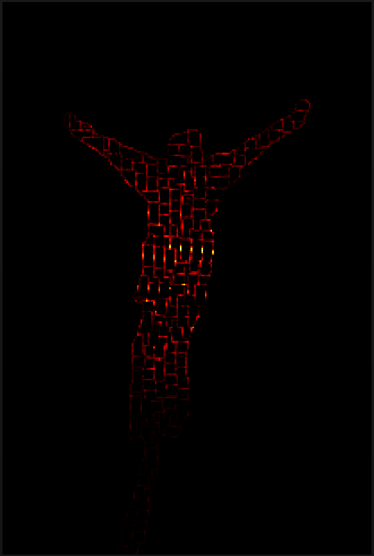

Our data term

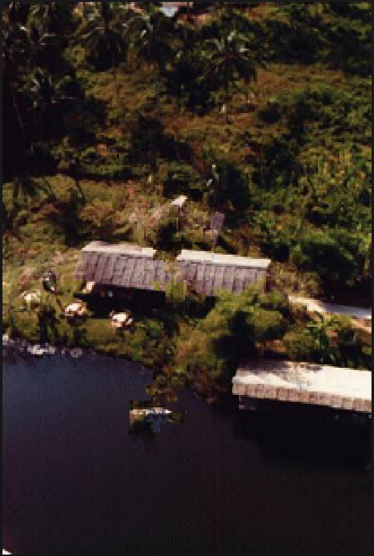

Inpainting result with original data term from *[Criminisi et al. '03]*

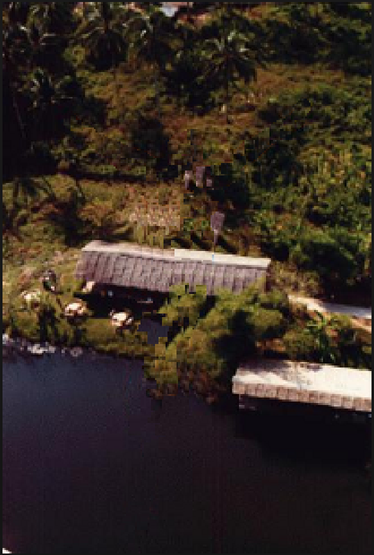

Inpainting result with our data term

A better patch lookup strategy

- Inspired from [Ashikhmin '01, PatchMatch '09]
- Local/Global scheme
- Using location of already-reconstructed patches

Idea : Enhance the geometric cohesion of patches pasted side by side.

A better patch lookup strategy

$$\psi_{\hat{p}} = \{\psi_q \mid q = \underset{r \in \mathcal{S}(p)}{\arg\min}\, d(\psi_p, \psi_r)\} \quad \text{avec} \quad \mathcal{S}(p) = \bigcup_{i \in \Phi(p)} \mathcal{W}_{\hat{i}+p-i} \qquad (5)$$

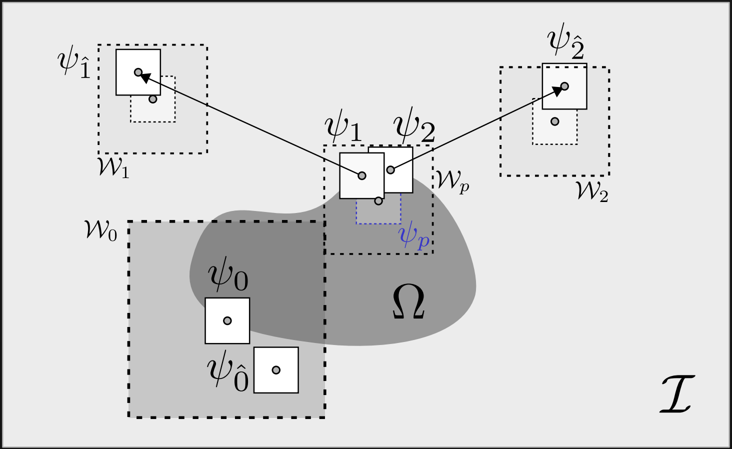

A better patch lookup strategy

$$\psi_{\hat{p}} = \{\psi_q \mid q = \operatorname*{arg\,min}_{r \in \mathcal{S}(p)} d(\psi_p, \psi_r)\} \quad \text{avec} \quad \mathcal{S}(p) = \bigcup_{i \in \Phi(p)} \mathcal{W}_{\hat{i}+p-i} \quad (5)$$

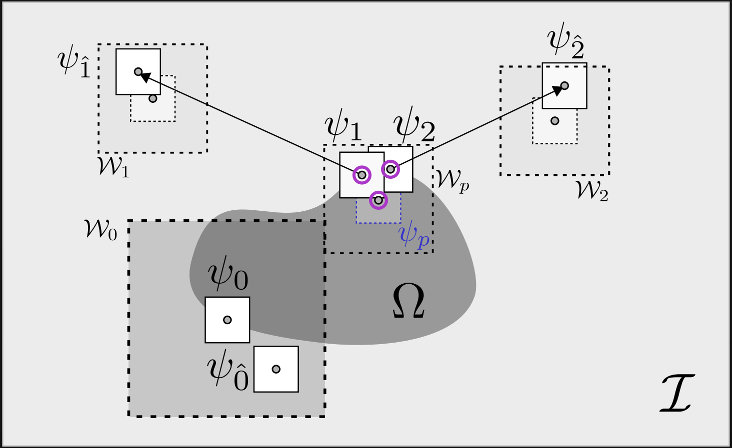

A better patch lookup strategy

$$\psi_{\hat{p}} = \{\psi_q \mid q = \underset{r \in \mathcal{S}(p)}{\arg\min}\, d(\psi_p, \psi_r)\} \quad \text{avec} \quad \mathcal{S}(p) = \bigcup_{i \in \Phi(p)} \mathcal{W}_{\hat{i}+p-i} \quad (5)$$

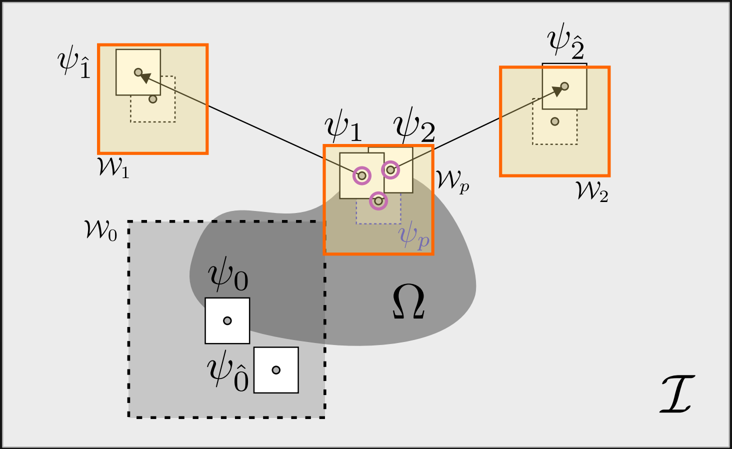

Results

512×384

Original image

512×384

Masked image

Results

512×384

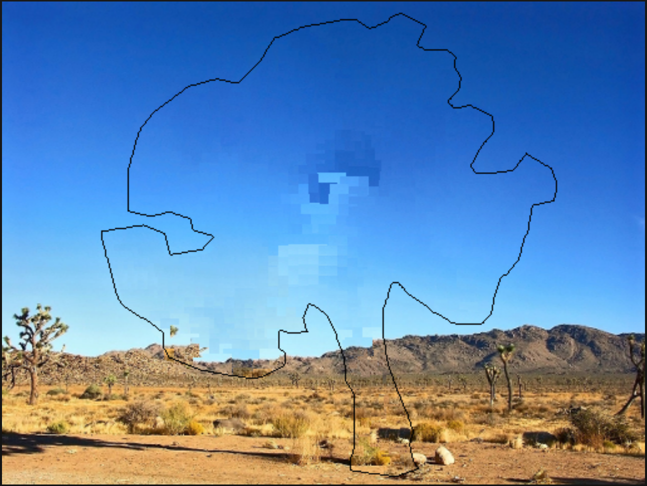

Inpainting result with *[Criminisi et al. '03]* (4.5s)

Results

512×384

Inpainting result with our improvements (3.2s)

Results

800×450

Original image

Results

800×450

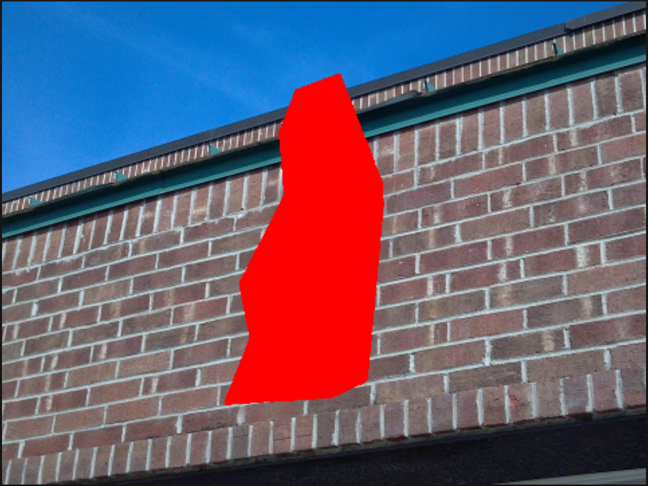

Masked image

Results

800×450

Inpainting result with *[Criminisi et al. '03]* (1.7s)

Results

800×450

Inpainting result with our improvements (0.8s)

Contributions

A patch blending technique for inpainting

Patch blending : Idea

Patch blending : Idea

Patch blending : Idea

Patch blending : Idea

Patch blending : Method

- **Idea :** Visually minimize block artifacts due to side-by-side pasted patches.

1. Detect the locations of the artifacts
   - Remember the location of the pasted patches
2. Spatial blending of the patches
   - Reduce the bloc artifacts

Detect artifact locations

We build an estimated "map of artifacts"

- 2 hypotheses on their location
  - Local discontinuity in Color/Luminosity
  - Discontinuity in the source patch localization
- Model by

$$P_{\mathcal{A}} = \| \nabla I \| \,.\, div(\phi)$$

with ϕ the shiftmap of the inpainting step done previously

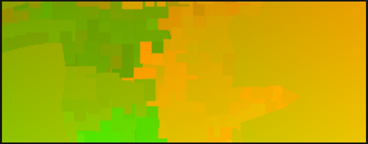

Example of ϕ

Detect artifact locations

We build an estimated "map of artifacts"

Detect artifact locations

We build an estimated “map of artifacts”

Detect artifact locations

We build an estimated “map of artifacts”

Detect artifact locations

Keep only the more visible

- Apply a simple threshold by a value $\tau \in [0,1]$

$$L_{\mathcal{A}} = \{p \mid P_{\mathcal{A}}(p) > \tau\}$$

Detect artifact locations

Keep only the more visible

- Apply a simple threshold by a value $\tau \in [0, 1]$

$$L_{\mathcal{A}} = \{p \mid P_{\mathcal{A}}(p) > \tau\}$$

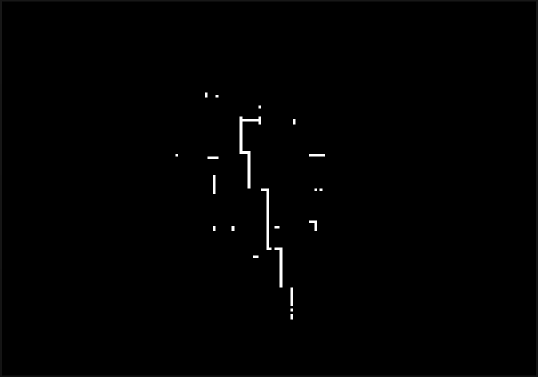

Detect artifact locations

Build a "blending amplitude map".

$$\sigma(p) = \rho \, . \, \frac{\sum_{q \in L_{\mathcal{A}}} w(p,q)}{\max_{r \in \mathcal{I}} \sum_{q \in L_{\mathcal{A}}} w(p,q)} \tag{6}$$

with

$$w(p,q) = \exp\left(- \frac{\|p-q\|^2}{2.P_{\mathcal{A}}(q)^2} \right) \tag{7}$$

where ρ is the maximum blending bandwidth

Detect artifact locations

Build a “blending amplitude map”.

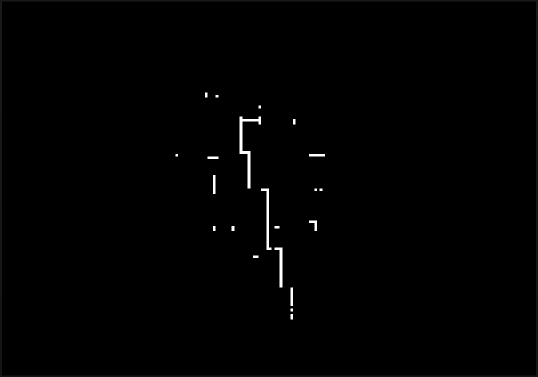

Detect artifact locations

Build a “blending amplitude map”.

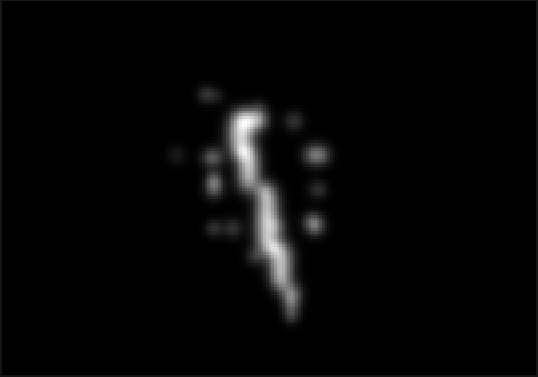

Spatial blending of patches

Compute a weighted average of pixels from patchs that overlap :

$$J(p) = \frac{\sum_{\psi_q \in \Psi_p} w(p,q)\ .\ \psi_q(p-q)}{\sum_{\psi_q \in \Psi_p} w(q,p)} \tag{6}$$

- $w(p,q) = \exp\left(-\frac{d(p,q)^2}{2.\sigma(p)^2}\right)$
- $d(p,q) = \min_{r \in \mathcal{N}_q} \|p-q\|$
- $\Psi_p = \{\psi_q \mid \psi_q \cap \psi_p \neq \varnothing\}$

- Computational complexity equivalent to a 2d convolution (with a spatially varying kernel)
  - ▶ But, fast algorithm when σ is quantized.

Spatial blending of patches

Masked image

Spatial blending of patches

Inpainting result with improved version of *[Criminisi et al. '03]*

Spatial blending of patches

Result after patch blending

Spatial blending of patches

Original image

Spatial blending of patches

Masked image

Spatial blending of patches

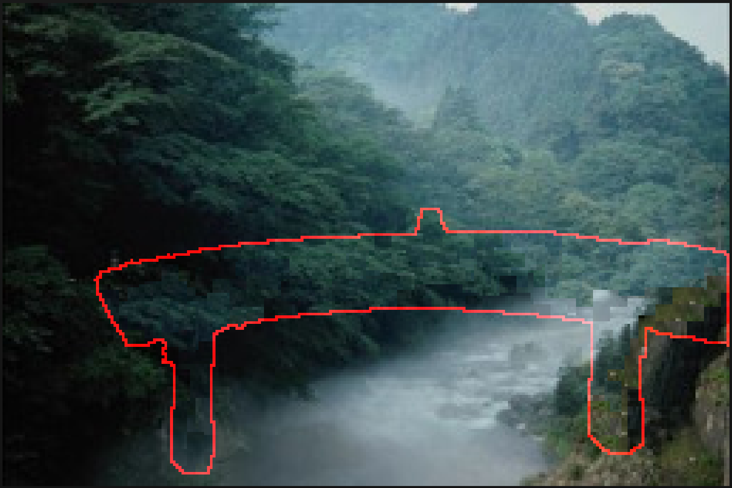

Inpainting result with improved version of *[Criminisi et al. '03]*

Spatial blending of patches

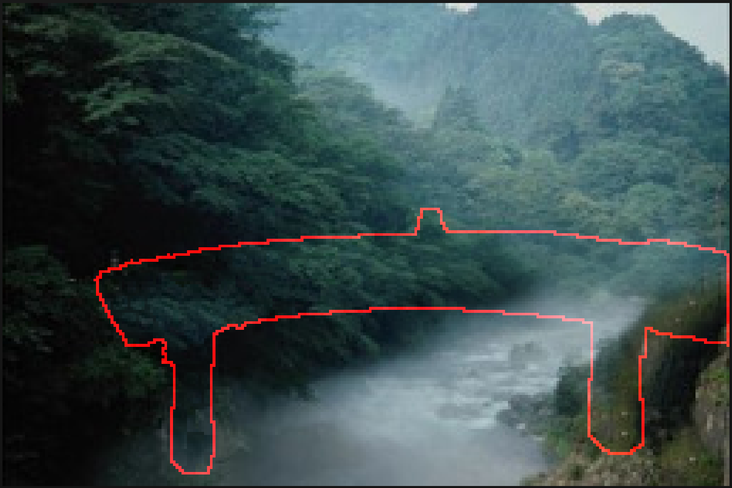

Result after patch blending (1.2 s)

Inpainting result with improved version of *[Criminisi et al. '03]*

Result after patch blending (**isotropic blending**)

Is it possible to take the local geometry of the image contours into account ?

Analogy with image regularization

Anisotropic image regularization

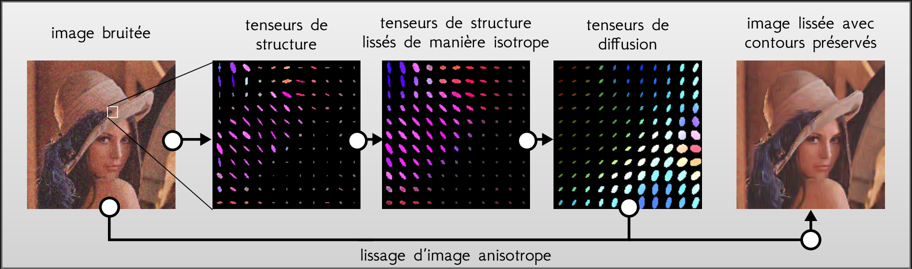

Analogy with image regularization

Anisotropic blending of patches

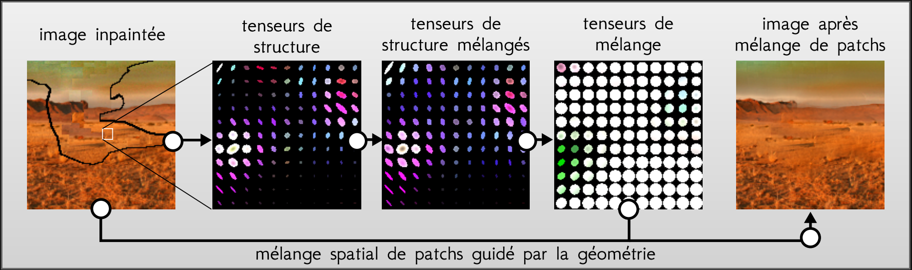

A tensor model for patch blending

From structure tensors :

- Eigenvalues of the blending tensors

$$\lambda_{\mathbf{B}i} = \frac{1}{\left(1 + \hat{\lambda}_{\mathbf{S}1} + \hat{\lambda}_{\mathbf{S}2}\right)^{\gamma_i}}$$

with γ_i a parameter related to the overall anisotropy of the tensors

- Blending tensors are computed as :

$$\mathbf{B} = \lambda_{\sigma\mathbf{B}1} \mathbf{e}^{\perp}_{\mathbf{S}1} . ^T \mathbf{e}^{\perp}_{\mathbf{S}1} + \lambda_{\sigma\mathbf{B}2} \mathbf{e}^{\perp}_{\mathbf{S}2} . ^T \mathbf{e}^{\perp}_{\mathbf{S}2} \quad (7)$$

  - $\mathbf{e}_{\mathbf{S}i}$: eigenvectors of the (inpainted) structure tensor
  - $\lambda_{\sigma\mathbf{B}i} = \sigma_{\mathbf{B}} \lambda_{\mathbf{B}i}$
  - $\sigma_{\mathbf{B}}$: maximal bandwidth of the blending

A tensor model for patch blending

Example of blending tensor field

A tensor model for patch blending

$$J(p) = \frac{\sum_{\psi_q \in \Psi_p} w(p,q) \,.\, \psi_q(p-q)}{\sum_{\psi_q \in \Psi_p} w(p,q)}$$

$$w(p,q) = \begin{cases} \exp\left(-\frac{\|p-q\|^2}{2\sigma(p)^2}\right) & \text{isotropic} \\ \exp\left(X^T \mathbf{B}(p)^{-1} X\right) & \text{anisotropic (edge-oriented)} \end{cases}$$

where $X = q - p$

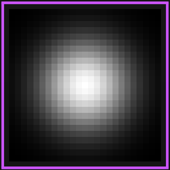

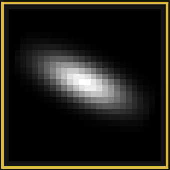

Inpainting result with improved version of *[Criminisi et al. '03]*
without patch blending

After *isotropic* patch blending


After *anisotropic* patch blending

Masked image

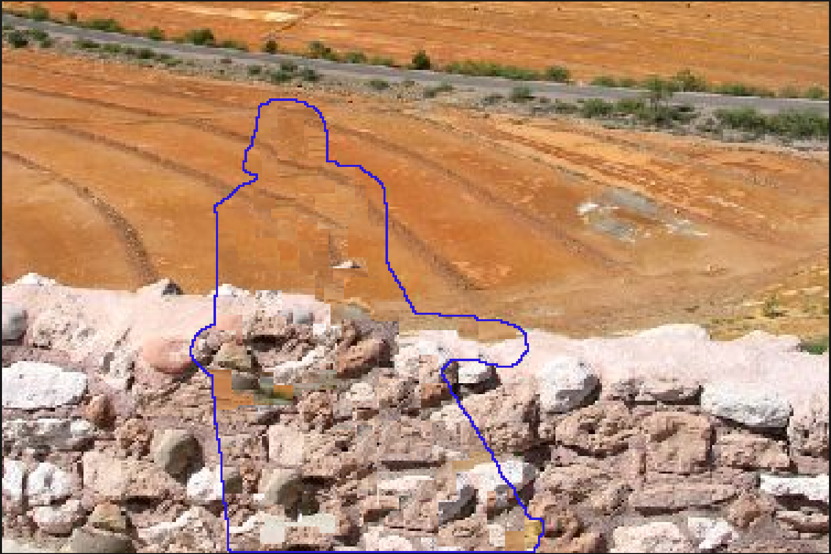

Inpainted without blending

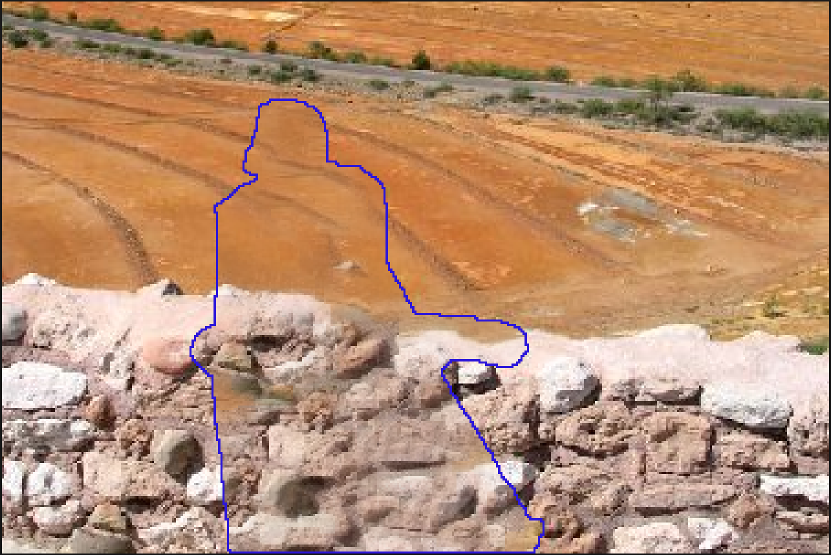

After isotropic blending

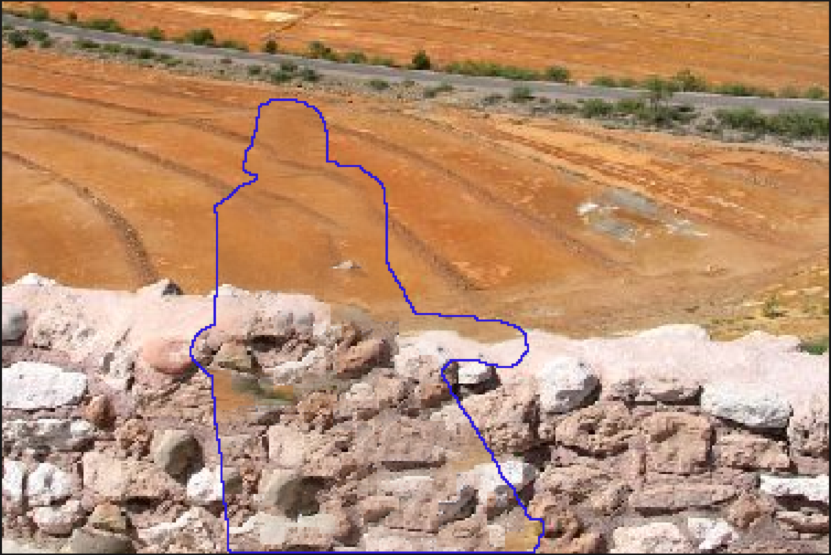

After anisotropic blending

Original image

Masked image

Inpainted with *[Criminisi et al. 2003]*

Inpainting with anisotropic patch blending

Contributions

Extension to video sequences

Video inpainting

- Inpainting priorities are still computed frame by frame
- But patches and lookup regions are now in 3d
  - e.g. with size $5 \times 5 \times 3$

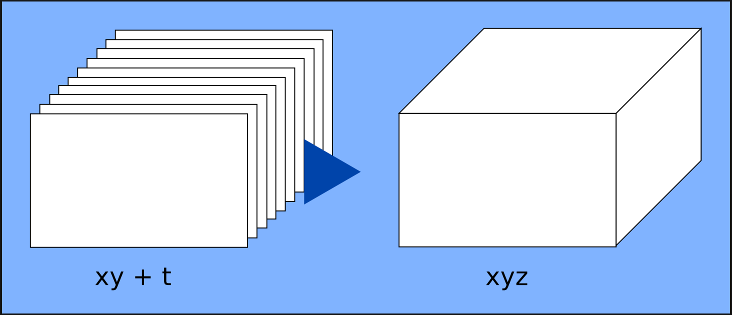

Video inpainting

Patch lookup strategy is easily extended in $3D$

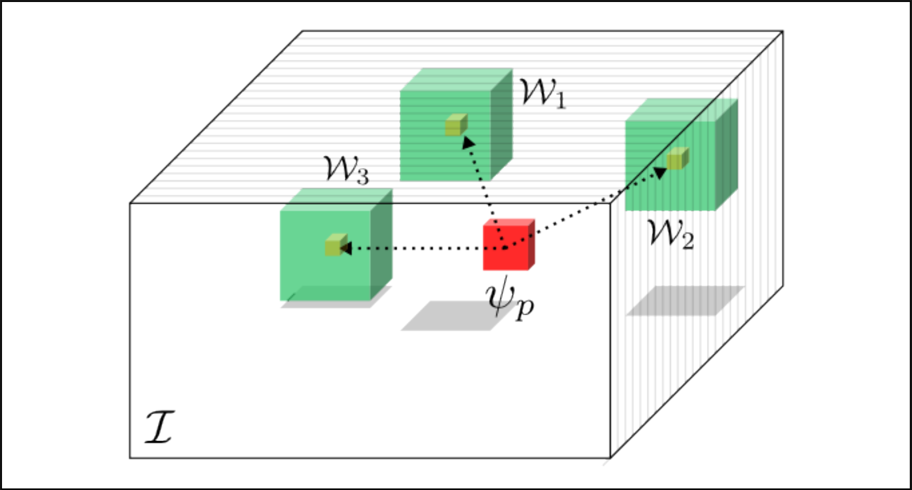

- Adding a temporal term to the blending tensors (that are now 3x3)

$$\mathbf{B} = \underbrace{\lambda_{\sigma\mathbf{B}1}\mathbf{e}_{\mathbf{S}1}^{\perp}.^{T}\mathbf{e}_{\mathbf{S}1}^{\perp} + \lambda_{\sigma\mathbf{B}2}\mathbf{e}_{\mathbf{S}2}^{\perp}.^{T}\mathbf{e}_{\mathbf{S}2}^{\perp}}_{\text{Spatial term}} + \underbrace{\lambda_{\sigma\mathbf{B}3}\mathbf{e}_{\mathbf{S}3}^{\perp}.^{T}\mathbf{e}_{\mathbf{S}3}^{\perp}}_{\text{Temporal term}} \tag{8}$$

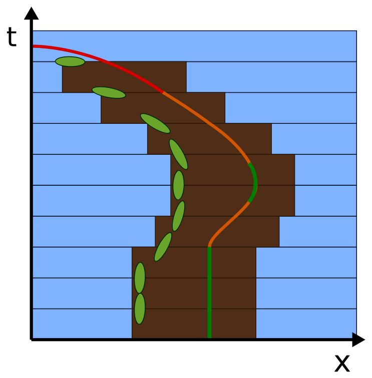

Beach Umbrella

Video inpainting

Fontaine Chatelet

Reconstruction of Resynthetized Stereoscopic Views

Initial image (from stereo pair)

Resynthetized viewpoint (from stereo + depth-map)

Inpainted with standard inpainting algorithm.

Inpainting depth-map

We have designed a specified inpainting algorithm to reconstruct missing regions in depth-map :

1. Segment the original depth-map (the one without holes) J_o
2. Estimate the motion of those super-pixels between J_o and J_s (*registration*)
3. Reconstruct missing pixels in J_s from interpolating several correspondent in J_o.

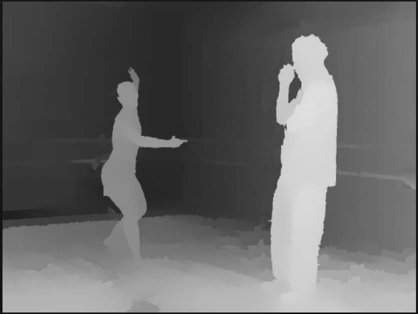

Inpainting with depth-map

Algorithm *[Criminisi, Pérez et Toyama '03]* with our improvements + taking the depth into account :

- Modification of the confidence term C_p
- Modification of the lookup strategy
- Modification of the patch synthesis step
- A single additional parameter
  - λ : Threshold to gather/differenciate pixels considered to have same/different depths

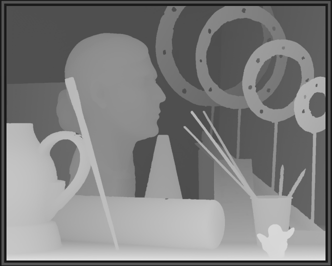

J_s (already inpainted)

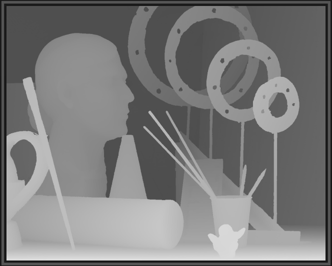

J_o

I_s

I_o

Inpainting with depth-map

- Modification of the confidence term C_p

$$C_p = \frac{1}{\mathcal{N}_p} \sum_{\substack{q \in \mathcal{N}_p \cap (\mathcal{I}_s - \Omega) \\ |J_s(p) - J_s(q)| < \lambda}} C_q \tag{9}$$

Inpainting with depth-map

- Modification of the confidence term C_p

$$C_p = \frac{1}{\mathcal{N}_p} \sum_{\substack{q \in \mathcal{N}_p \cap (\mathcal{I}_s - \Omega) \\ |J_s(p) - J_s(q)| < \lambda}} C_q \tag{9}$$

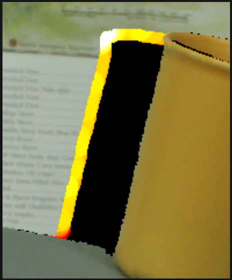

Inpainting with depth-map

- Patch lookup strategy

$$\psi_{\hat{p}} = \left\{ \psi_q \in I_o \mid \underset{|J_o(q)-J_s(p)|<\lambda}{\arg\min} d(\psi_p, \psi_q) \right\} \tag{10}$$

I_s

I_o

- Patch synthesis, only for pixels with same depth

$$\psi_p(q) = \psi_{\hat{p}}(r) \left| \begin{array}{l} q \in \mathcal{N}_p \cap \Omega \\ |J_o(r) - J_s(q)| < \lambda \end{array} \right. \tag{11}$$

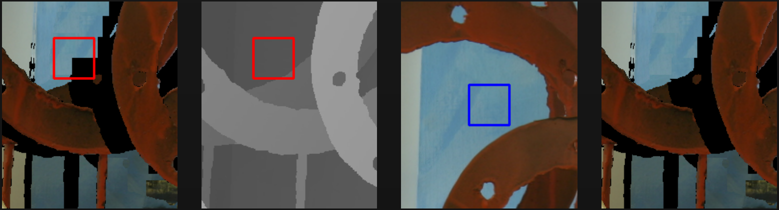

Inpainting with depth-map

Image

Inpainting with depth-map

Resynthesized viewpoint

Inpainting with depth-map

Inpainted viewpoint, with depth constraints

Inpainting with depth-map

Resynthesized viewpoint (without depth constraints)

Inpainting algorithm with spatial patch blending in :

[http ://gmic.eu]