



HAL
open science

AI-based Generation of a Multi-Modal Panorama for the Bayeux Tapestry

Yvain Quéau, Marjorie Redon, Matthieu Pizenberg, Abderrahim Elmoataz

► **To cite this version:**

Yvain Quéau, Marjorie Redon, Matthieu Pizenberg, Abderrahim Elmoataz. AI-based Generation of a Multi-Modal Panorama for the Bayeux Tapestry. *Advances in Applied Mathematics and Learning Approaches for Cultural Heritage and the Arts*, Apr 2024, Cagliari (Sardaigne), Italy. hal-04567887

HAL Id: hal-04567887

<https://hal.science/hal-04567887>

Submitted on 3 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

AI-BASED GENERATION OF A MULTI-MODAL PANORAMA FOR THE BAYEUX TAPESTRY

Yvain QUÉAU

joint work with Marjorie REDON, Matthieu PIZENBERG, Abderrahim ELMOATAZ

CNRS, GREYC, Caen, FRANCE

project funded by ANR IMG

yvain.queau@ensicaen.fr
<https://yqueau.github.io/>



Aim of the talk



Reconstruction of a multimodal panorama of the Bayeux Tapestry, enhancing the accessibility of the artwork. Target modalities, and corresponding imaging tools:

- ▶ daylight → well-tuned heuristics ;
- ▶ multispectral → variational methods ;
- ▶ geometry → deep learning.

The multimodal reconstruction can be explored through a web interface:

<https://redon213.users.greyc.fr/>.

1. The Bayeux Tapestry
2. Keypoint-based daylight panorama generation
3. Low-rank multispectral registration
4. Deep image-to-geometry translation

1. The Bayeux Tapestry
2. Keypoint-based daylight panorama generation
3. Low-rank multispectral registration
4. Deep image-to-geometry translation

The Bayeux Tapestry



70m-long medieval wool and linen *embroidery* telling the conquest of England by William, Duke of Normandy, in 1066.

The Bayeux Tapestry



Story starts with the death of old king Edward.

The Bayeux Tapestry



Edward had chosen his cousin William for his succession on the British throne.

The Bayeux Tapestry



Yet Edward's brother-in-law, Harold, took the throne.

The Bayeux Tapestry



Hence William's armada crossed the Channel.

The Bayeux Tapestry



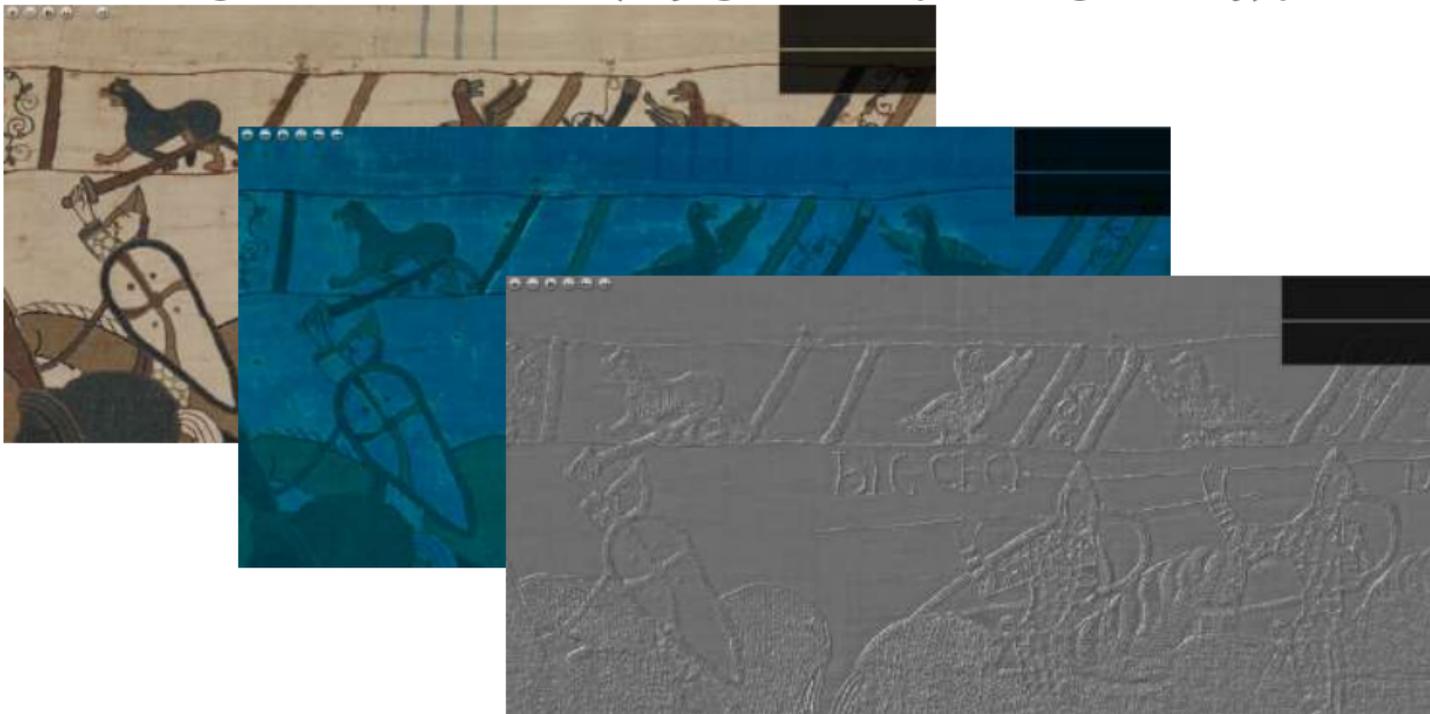
And William eventually defeated Harold during the battle of Hastings, hence becoming the King of England.



- ▶ Research complicated by: the number of visitors (400,000 per year), the exceptional size (70m long and 50cm high), the profusion of scenes (55), and the protective glass.
- Zoomable online interface for remote exploration.
- ▶ Unknown stages of the artistic process of its assembly and realization, inspection of the current state of the artwork.
- Multispectral and geometric modalities.

The multimodal digital twin

Zoomable, registered multimodal (daylight, multispectral, geometry) panorama:



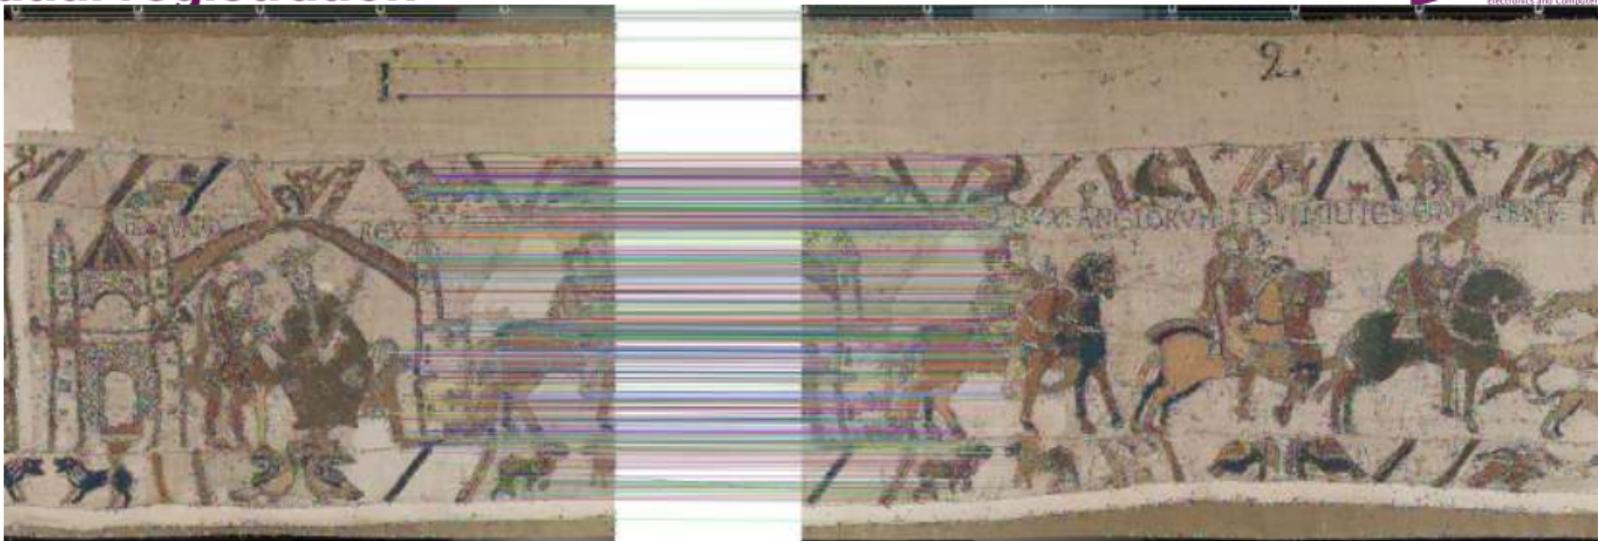
Online viewing: <https://redon213.users.greyc.fr/>.

1. The Bayeux Tapestry
2. Keypoint-based daylight panorama generation
3. Low-rank multispectral registration
4. Deep image-to-geometry translation

- ▶ Input: 86 high resolution images, acquired in 2017 by La Fabrique de patrimoines en Normandie under daylight ;
- ▶ Output: $480,000 \times 6,000$ px RGB panorama.



Berthelot et al. (2017), *La reconstruction du panorama de la Tapisserie de Bayeux comme fond de référence d'un système d'information documentaire spatialisée*, In Situ Revue des patrimoines.



Registration procedure

1. SIFT keypoint detection and matching;
2. Moving least-squares deformation estimation;
3. Warp next image and stitch it.

Panorama visualization

The reconstructed $480,000 \times 6,000$ px panorama can obviously not be viewed natively.

→ Custom tiling procedure¹, and visualization using OpenSeaDragon.

consecutive input images

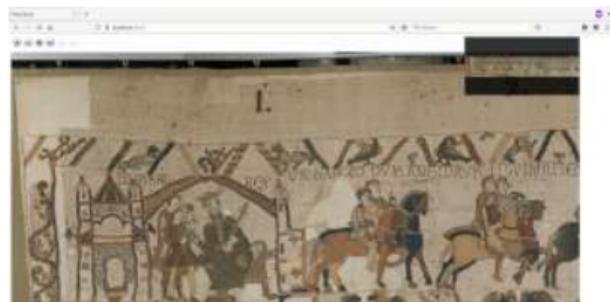


zoomtiler

deepzoom format



zoomable tiled images generation



OpenSeadragon Web image viewer

¹ Pizenberg (2021). *Zoomtiler, a simple CLI program to generate zoomable tiled images.*
<https://github.com/mpizenberg/zoomtiler>.

Online exploration of the artwork in daylight

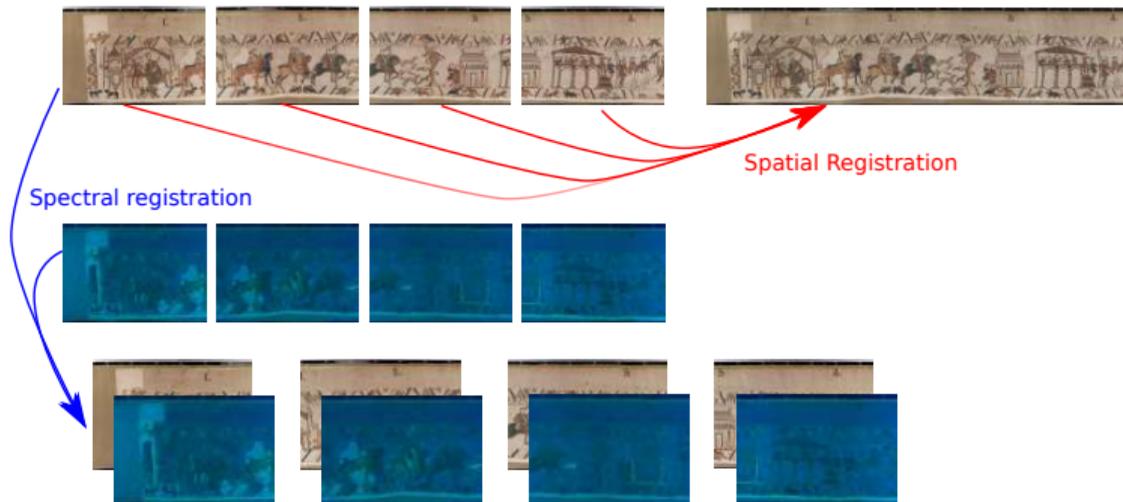
The history told by the artwork can eventually be explored in a zoomable web interface:



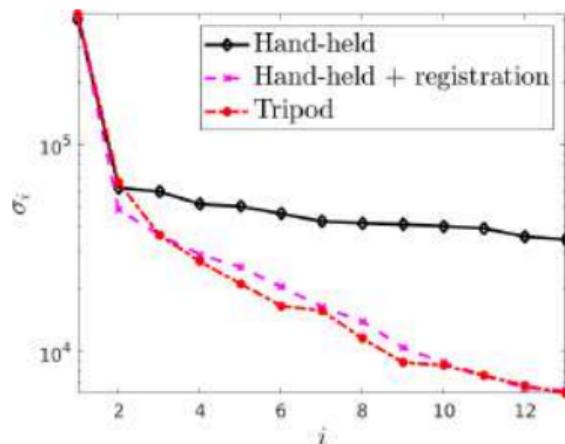
1. The Bayeux Tapestry
2. Keypoint-based daylight panorama generation
- 3. Low-rank multispectral registration**
4. Deep image-to-geometry translation

Spectral registration

- ▶ Input: **86 images**, captured under **6 imaging modalities** (natural lighting, UV, IR, etc.), with $\approx 100px$ sensor displacements between modalities.
- ▶ Output: 86 pixel-accurate registered multispectral images, which can then be stitched exactly as the daylight ones.



Pizenberg et al. (2021), *Low-rank registration of images captured under varying, unknown lighting, SSVM*.



Singular values of the observations matrix

Reorganize observations into a matrix $M \in \mathbb{R}^{N_{\text{observations}} \times N_{\text{modalities}}}$, then find the transformation \mathcal{T} such that

$$\text{rank}(\mathcal{T}(M)) \rightarrow \min .$$

This nonconvex problem can be relaxed into:

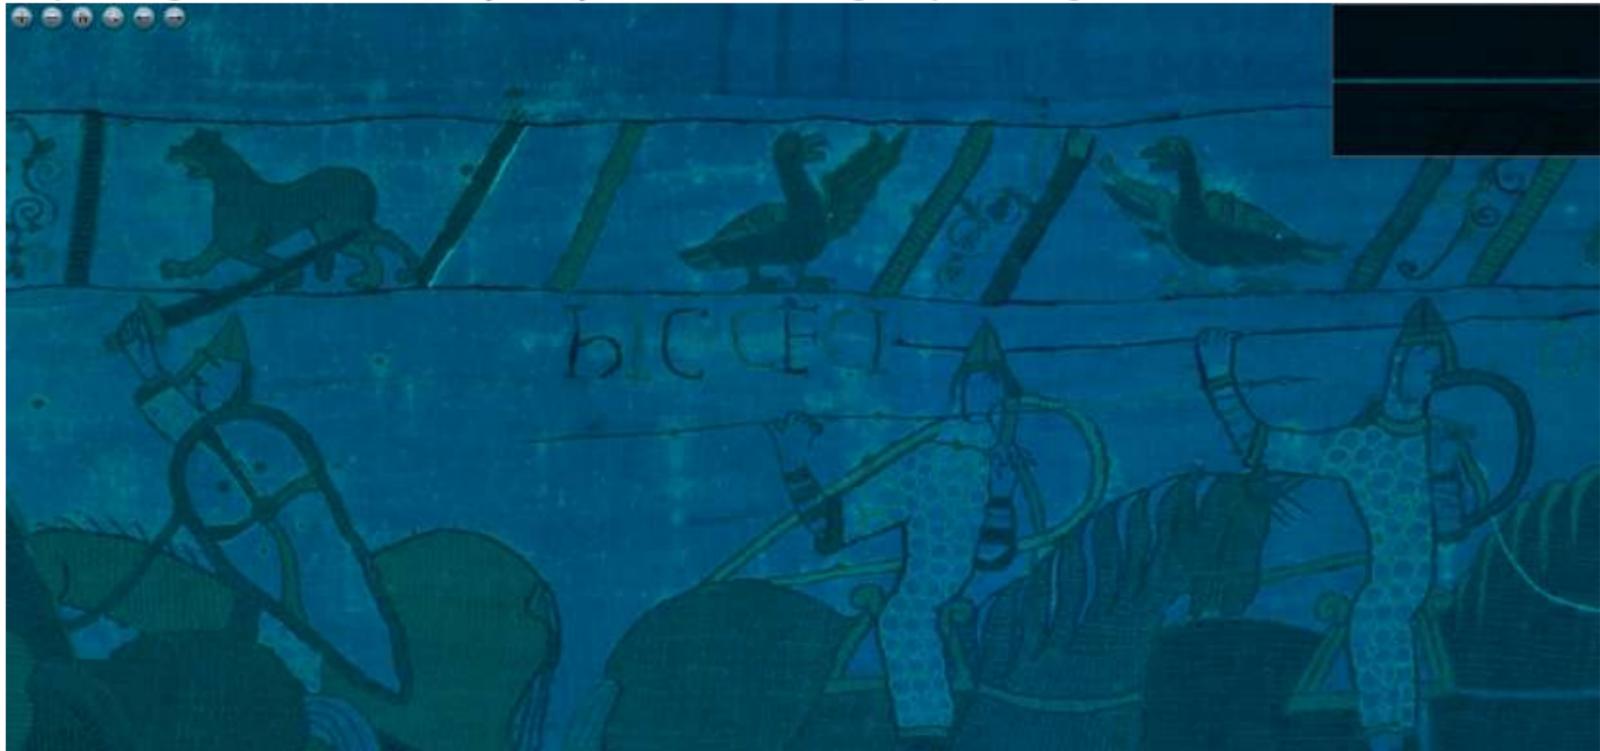
$$\min_{A, \mathcal{T}} \|A\|_{\star} + \lambda \|A - \mathcal{T}(M)\|_1 .$$

and solved efficiently using augmented Lagrangian methods:

$$\begin{aligned} \min_{A, e, \mathcal{T}} \|A\|_{\star} + \lambda \|e\|_1 \\ \text{s.t. } e = A - \mathcal{T}(M). \end{aligned}$$

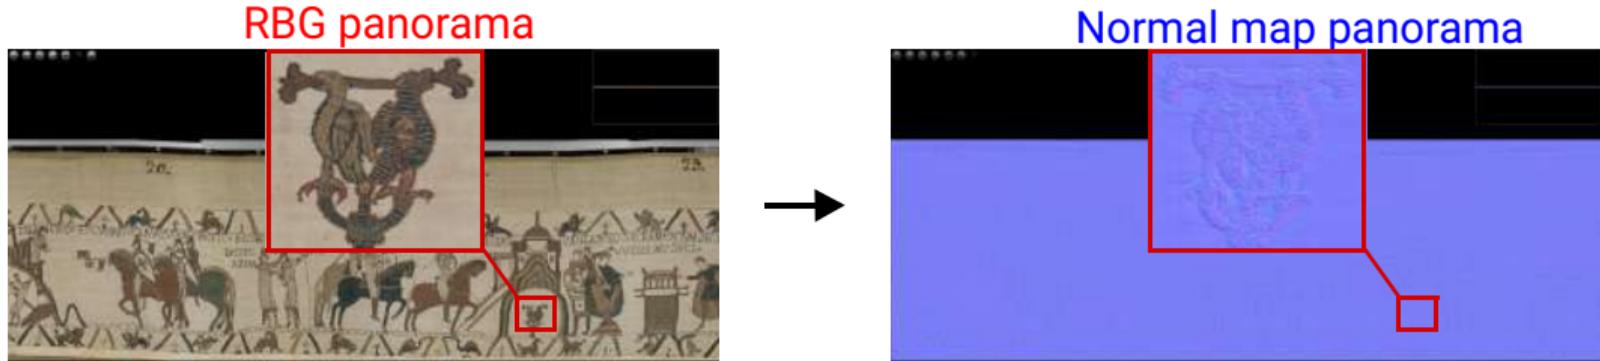
Online exploration of the artwork in UV

Exploring the UV modality may reveal damaged part, e.g. moisture:



1. The Bayeux Tapestry
2. Keypoint-based daylight panorama generation
3. Low-rank multispectral registration
4. Deep image-to-geometry translation

3D-digitization of the Bayeux tapestry



Goal: construct a 2.5D panorama of this 70 m-long medieval wool and linen *embroidery*, in view of geometry inspection and tactile experiments for visually-impaired people.

- ▶ RGB panorama readily available;
- ▶ Can we convert it to geometry?

Redon et al. (2023), *3D surface Approximation of the Entire Bayeux Tapestry for Improved Pedagogical Access*, ICCV workshops.

Proposed strategy



Spatial Registration

Daylight panorama, to be converted to normals

Proposed strategy for 3D-digitization

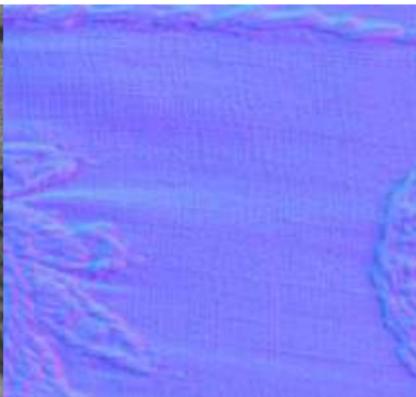
1. Store the estimated daylight spatial registration parameters (cf. Sect. 2);
2. Turn each RGB image to 2.5D using deep learning;
3. Apply the same spatial registration to the 2.5D images → geometric panorama.

Deep image-to-geometry learning

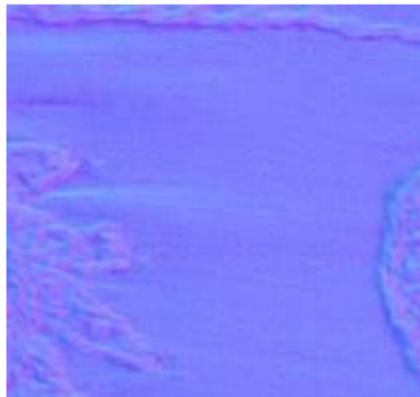
Input
image



Ground truth
normal map



Estimated
normal map



Proposed strategy for 3D-digitization

1. Store the RGB spatial registration parameters;
2. Turn each RGB image to 2.5D using deep learning:
→ a) ground truth acquisition; b) deep network architecture;
3. Apply the same spatial registration to the 2.5D images.

Ground truth geometry acquisition campaign

Photometric stereo-based 3D-reconstruction of a few scenes:



Left: three input images, taken from the same viewing angle but varying lighting.

Right: output high-resolution mesh (5M triangles).

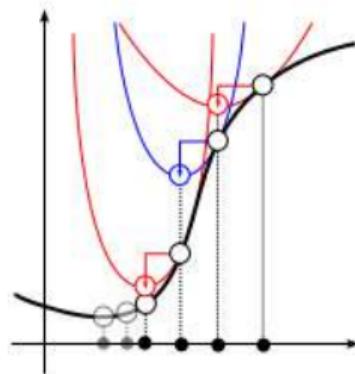
Data acquisition for photometric stereo

Image capture from fixed viewpoint, but varying illumination:

Given images $I^i : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, estimate shape $z : \Omega \rightarrow \mathbb{R}$ and reflectance $\rho : \Omega \rightarrow \mathbb{R}$ as solution of a variational problem:

$$\min_{z, \rho} \sum_{i=1}^m \int_{x \in \Omega} \Phi \left(\left| \rho(x) \ell^i \cdot \left[\nabla z(x)^\top, -1 \right]^\top - I^i(x) \right| \right) dx.$$

- ▶ ϕ : robust (non-convex) redescending M-estimator;
- ▶ Optimization by alternating MM.



QUÉAU, MECCA and DUROU (2016), *Unbiased photometric stereo for colored surfaces*, CVPR.

QUÉAU et al. (2017), *A non-convex variational approach to photometric stereo under inaccurate lighting*, CVPR.

Shape and reflectance reconstruction



Left: input images. Right: image, reconstructed albedo and reconstructed normals.

The reconstructed albedo encodes the intrinsic properties of the object's material ; the reconstructed normals encode the geometric variations of its surface.

Ground truth (image, geometry) pairs



We have ≈ 30 couples (RGB, normals) of size 3000 px^2 .

→ Thousands of 128 px^2 patches for learning the mapping $\text{RGB} \mapsto \text{geometry}$.

Ground truth (image, geometry) pairs



We have ≈ 30 couples (RGB, normals) of size 3000 px^2 .

→ Thousands of 128 px^2 patches for learning the mapping $\text{RGB} \mapsto \text{geometry}$.

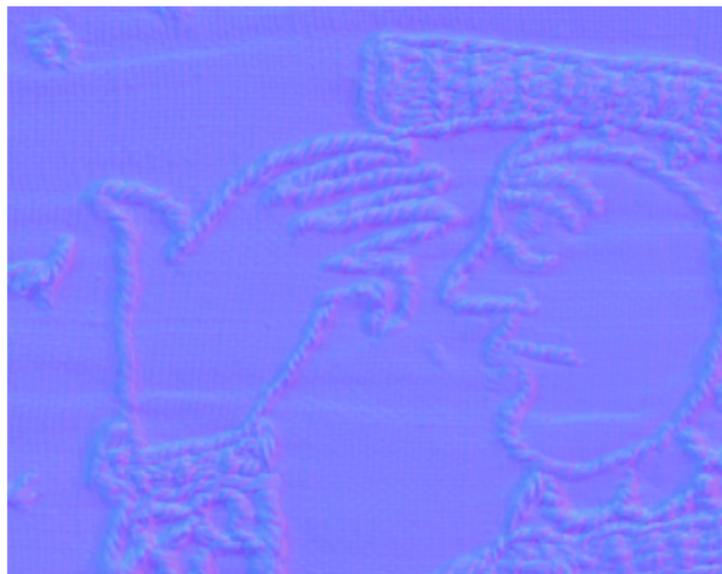
Ground truth (image, geometry) pairs



We have ≈ 30 couples (RGB, normals) of size 3000 px^2 .

→ Thousands of 128 px^2 patches for learning the mapping $\text{RGB} \mapsto \text{geometry}$.

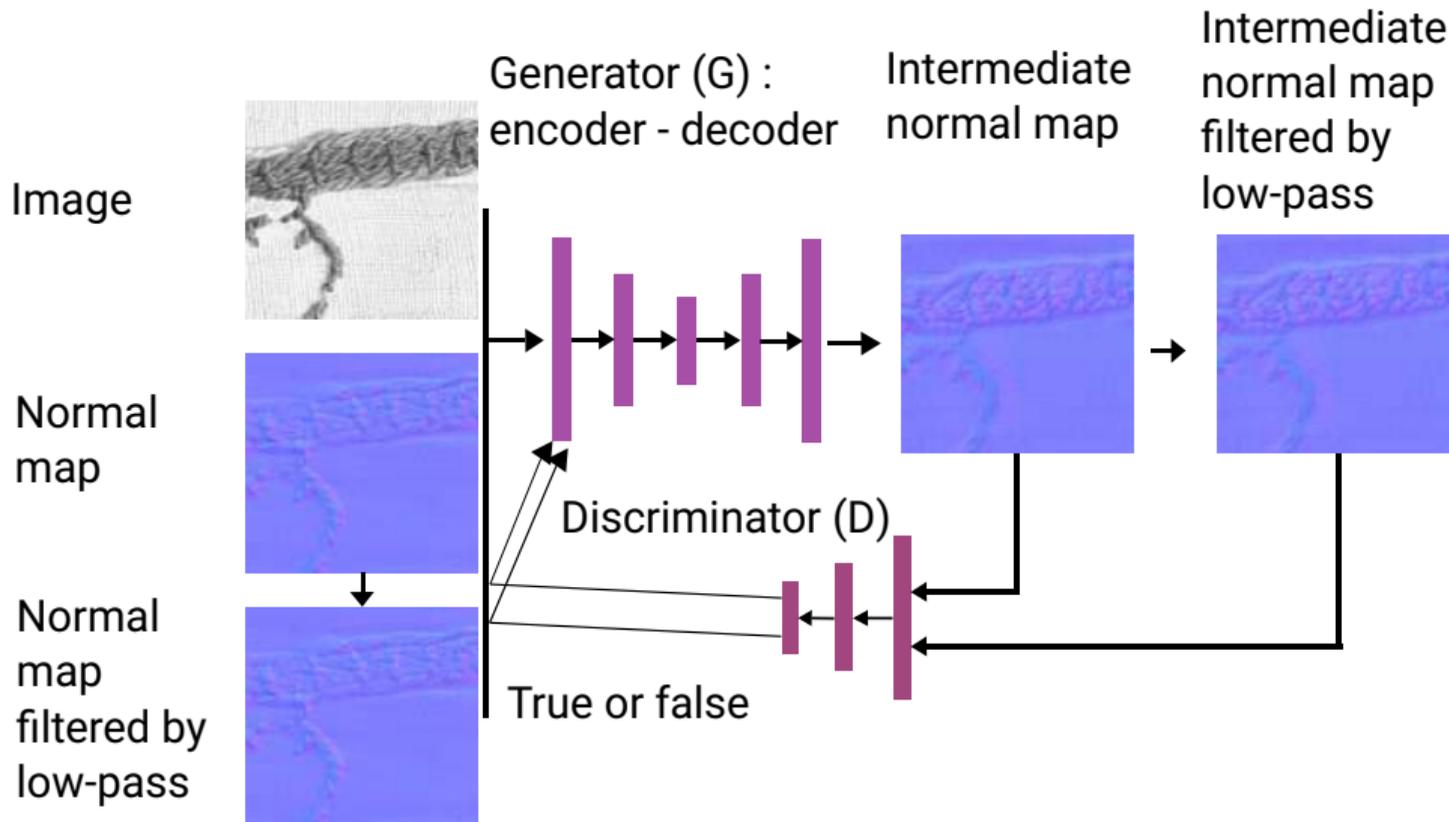
Ground truth (image, geometry) pairs



We have ≈ 30 couples (RGB, normals) of size 3000 px^2 .

→ Thousands of 128 px^2 patches for learning the mapping $\text{RGB} \mapsto \text{geometry}$.

Regularized GAN architecture



Proposed strategy for 3D-digitization

1. Store the RGB spatial registration parameters
2. Turn each RGB image to 2.5D using deep learning
→ a) ground truth acquisition ; b) deep network architecture
3. **Apply the same spatial registration to the 2.5D images**



Online exploration of the artwork's geometry

From the normal map images, we can eventually render shaded views of the artwork, highlighting the surface details:



Thank you

<https://redon213.users.greyc.fr>

yvain.queau@ensicaen.fr

