

# Fine-scale 3D-copying for culturage heritage

Yvain QUÉAU

CNRS researcher  
GREYC, University of Caen Normandy, France

Joint work with Matthieu PIZENBERG

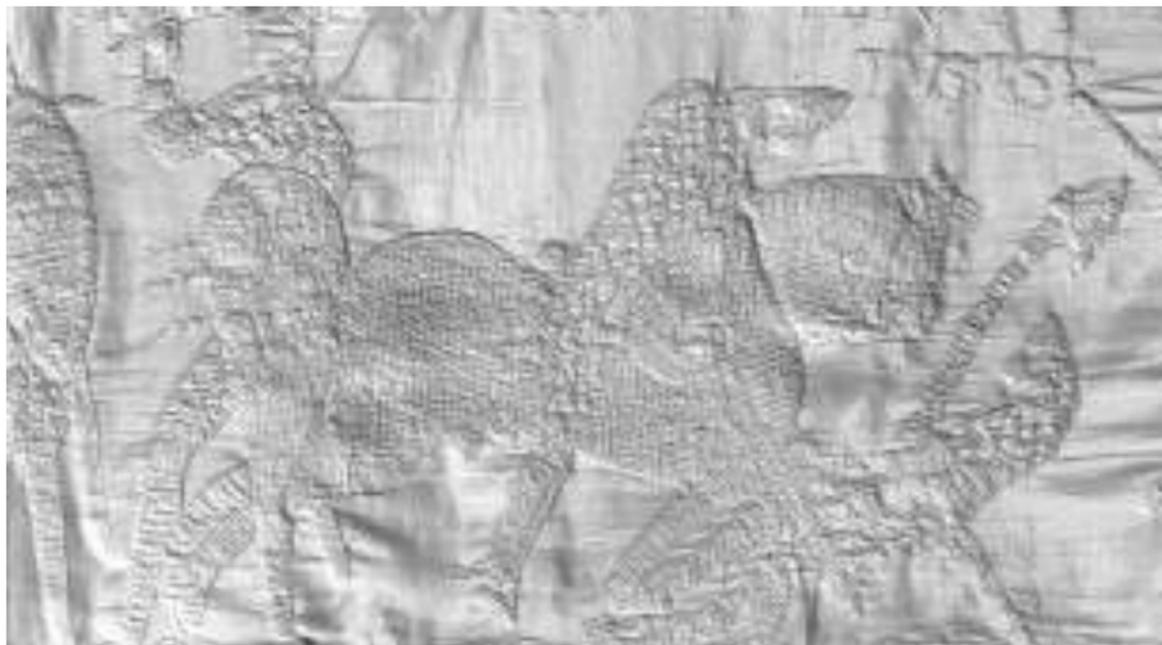
INDAM Workshop: from 3D Vision to 3D Printing  
Feb 12th, 2021

## Aim of the talk



Fine-scale 3D-reconstruction of Bayeux tapestry,  
to create 3D-printed mockups to be touched  
by visually impaired people

# Aim of the talk



Fine-scale 3D-reconstruction of Bayeux tapestry,  
to create 3D-printed mockups to be touched  
by visually impaired people

# Outline

- 1 The Inclusive Museum Guide project
- 2 Shape-from-shading and photometric stereo
- 3 Acquisition campaign

# Outline

- 1 The Inclusive Museum Guide project
- 2 Shape-from-shading and photometric stereo
- 3 Acquisition campaign

# Context of the project

- Funding: Normandy region and ANR projects “Inclusive Museum Guide” (2019-2024)
- Partners: Caen univ., Rouen univ., CNRS, Bayeux museum, Royal Holloway univ., Westminster univ.



AGENCE NATIONALE DE LA RECHERCHE

ANR

# Bayeux tapestry



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

## Bayeux tapestry: Edward



Story starts with the death of old king Edward...

## Bayeux tapestry: William



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

# Bayeux tapestry: Harold



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

# Bayeux tapestry: the Northmen armada



William's armada crossed the Channel

## Bayeux tapestry: the battle of Hastings



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

# The IMG project



Question: how can we make visually-impaired people “feel” such artifacts, knowing that

- they cannot see it
- they cannot touch it (the tapestry is protected by glass)

→ Re-think the way we access visual artworks:  
audio-descriptions, navigation tools, **tactile representations**,  
etc.

# Towards a tactile representation of visual artworks



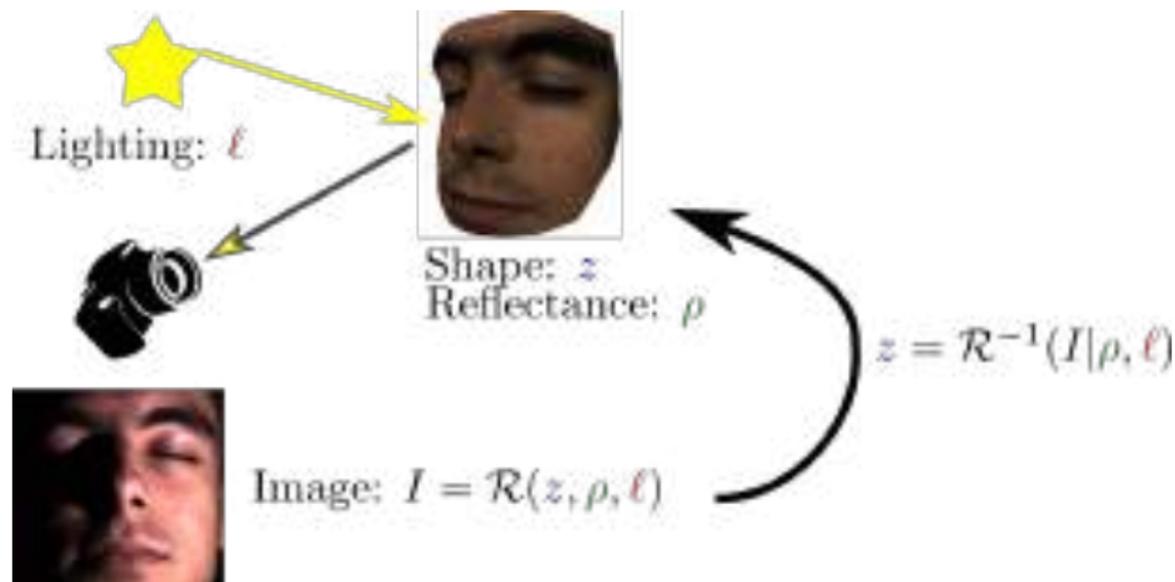
Goal: extract micro-geometry of the scene (“texture”),  
in view of 3D-printing

Challenges: no direct access (behind glass), fragility to light  
exposure, thinness of the wool strings, color changes...

# Outline

- 1 The Inclusive Museum Guide project
- 2 Shape-from-shading and photometric stereo
- 3 Acquisition campaign

# 3D-reconstruction as an inverse problem



Goal: invert the radiance function  $\mathcal{R}$  to recover the shape  $z$  (and, possibly, lighting  $\ell$  and reflectance  $\rho$ ) of the pictured surface

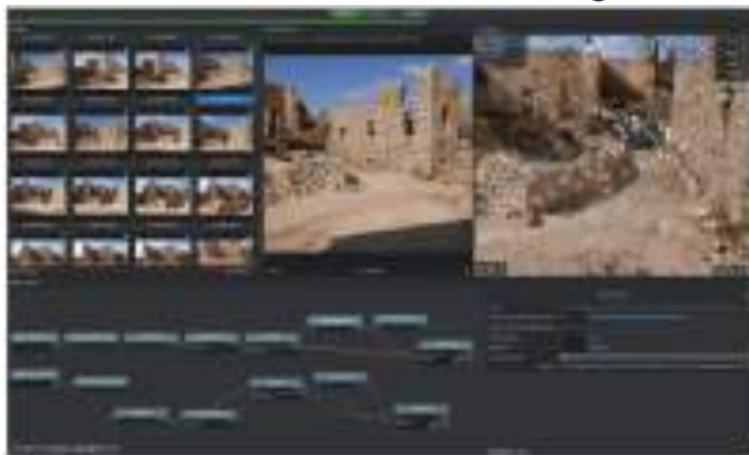
# Context: shape-from-X

	Geometric techniques	Photometric techniques
Single image	Structured light Shape-from-shadows Shape-from-contours Shape-from-texture Shape-from-template	<b>Shape-from-shading</b>
Multi-images	Structure-from-motion Stereopsis Shape-from-silhouettes Shape-from-focus	<b>Photometric stereo</b> Shape-from-polarisation

*Geometric techniques aim at identifying and analysing features. This presentation rather focuses on two photometric techniques (shape-from-shading and photometric stereo), which aim at inverting a physics-based image formation model.*

# Geometric techniques

Mature workflows based on structure-from-motion and multi-view stereo are available, e.g. Alice Vision's Meshroom:

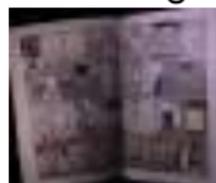


## Limitations:

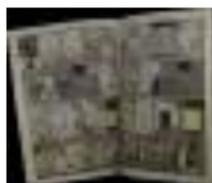
- May miss thin structures
- Only estimate shape, not reflectance

# 3D-scanning

3D-scanning = estimation of shape *and* color



Image



Reflectance



(a)



(b)



(c)



3D-model



(d)



(e)

- Geometric techniques (e.g. SfM, MVS) only recover shape (no color estimation)
- Photometric techniques (e.g., shape-from-shading, photometric stereo), which are based on **inverting the image formation model**, recover both ( $\approx$  *inverse rendering*)

# Photometric techniques

Top-performing for the recovery of very thin structures

Example: close-up on a 10 euros banknote



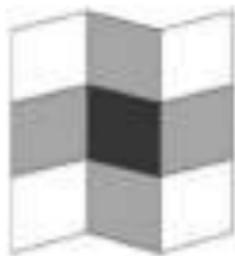
Durou et al., *Advances in Photometric 3D-Reconstruction*

# Shape-from-shading: a classic ill-posed problem

Given an image  $\mathbf{I} : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^m$ , *Shape-from-Shading* (SfS) consists in inverting the image irradiance equation

$$\mathbf{I} = \mathcal{R}(z, \rho, \ell) \quad (1)$$

with  $\mathcal{R}$  a *radiance* function depending on the unknown depth  $z : \Omega \rightarrow \mathbb{R}$ , surface reflectance  $\rho : \Omega \rightarrow \mathbb{R}^m$ , and incident lighting  $\ell : \Omega \rightarrow \mathbb{S}^2$ .



RGB image:  
 $\mathbf{I}$



Sculptor's  
explanation:  
 $z$



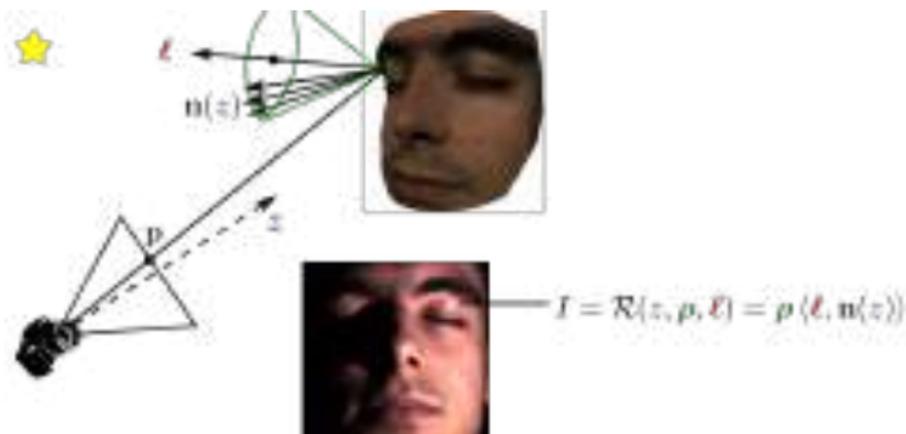
Painter's  
explanation:  
 $\rho$



Gaffer's  
explanation:  
 $\ell$

*Adelson and Pentland's workshop metaphor*

# Local ill-posedness in shape-from-shading



Even if  $\rho$  and  $\ell$  are known, the normal cannot be estimated directly by inverting the forward model, because infinitely many normals, lying on a cone, are equally admissible

→ A regularization mechanism is necessary. The most natural way to do so consists in considering a differential version of the forward model, assuming the surface is differentiable almost everywhere

# Illustration of ill-posedness

Even with known **surface reflectance**  $\rho$  and **incident lighting**  $\ell$ ,  
shape estimation by SfS is an ill-posed inverse problem  
(Horn, 1970)

Example: two solutions of  $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$  with  $\mathbf{I} := \text{Lena}$ , **white reflectance** ( $\rho \equiv 1$ ) and **frontal lighting** ( $\ell \equiv [0, 0, -1]^T$ ):



# Illustration of ill-posedness

Even with known **surface reflectance**  $\rho$  and **incident lighting**  $\ell$ ,  
shape estimation by SfS is an ill-posed inverse problem  
(Horn, 1970)

Example: two solutions of  $I = \mathcal{R}(z, \rho, \ell)$  with  $I :=$  Lena, **white**  
**reflectance** ( $\rho \equiv 1$ ) and **frontal lighting** ( $\ell \equiv [0, 0, -1]^T$ ):



# Illustration of ill-posedness

Even with known **surface reflectance**  $\rho$  and **incident lighting**  $\ell$ ,  
shape estimation by SfS is an ill-posed inverse problem  
(Horn, 1970)

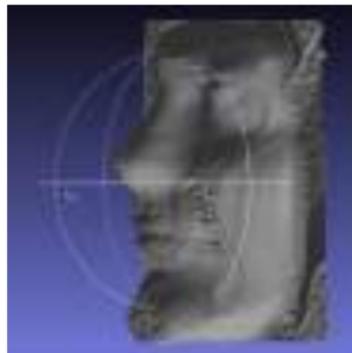
Example: two solutions of  $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$  with  $\mathbf{I} := \text{Lena}$ , **white reflectance** ( $\rho \equiv 1$ ) and **frontal lighting** ( $\ell \equiv [0, 0, -1]^T$ ):



# Illustration of ill-posedness

Even with known **surface reflectance**  $\rho$  and **incident lighting**  $\ell$ ,  
shape estimation by SfS is an ill-posed inverse problem  
(Horn, 1970)

Example: two solutions of  $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$  with  $\mathbf{I} := \text{Lena}$ , **white reflectance** ( $\rho \equiv 1$ ) and **frontal lighting** ( $\ell \equiv [0, 0, -1]^T$ ):



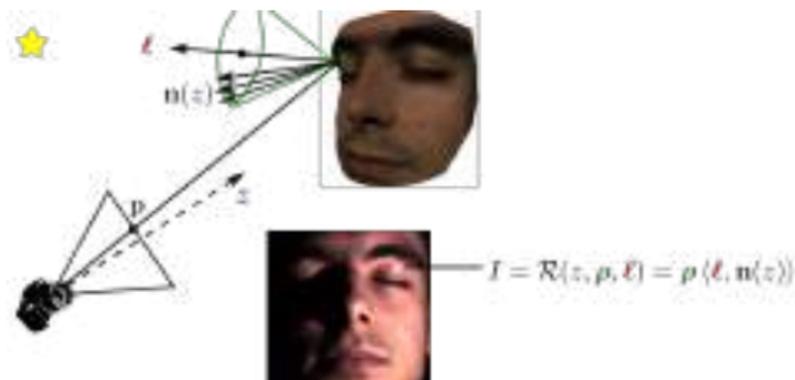
# Illustration of ill-posedness

Even with known **surface reflectance**  $\rho$  and **incident lighting**  $\ell$ ,  
shape estimation by SfS is an ill-posed inverse problem  
(Horn, 1970)

Example: two solutions of  $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$  with  $\mathbf{I} := \text{Lena}$ , **white reflectance** ( $\rho \equiv 1$ ) and **frontal lighting** ( $\ell \equiv [0, 0, -1]^T$ ):



# Differential image formation model



Assuming orthographic projection, the scene's surface is characterized by a depth map  $z : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$  such that

$$\mathbf{n}(z) = \frac{1}{\sqrt{|\nabla z|^2 + 1}} \begin{bmatrix} \nabla z \\ -1 \end{bmatrix}$$

And thus we get the following nonlinear PDE in  $z$ :

$$I = \mathcal{R}(z, \rho, \mathbf{l}) = \frac{\rho}{\sqrt{|\nabla z|^2 + 1}} \left\langle \mathbf{l}, \begin{bmatrix} \nabla z \\ -1 \end{bmatrix} \right\rangle$$

# Eikonal model

An important instance of the differential model

$$I = \mathcal{R}(z, \rho, \ell) = \frac{\rho}{\sqrt{|\nabla z|^2 + 1}} \left\langle \ell, \begin{bmatrix} \nabla z \\ -1 \end{bmatrix} \right\rangle$$

is that of a white surface ( $\rho \equiv 1$ ) under frontal lighting ( $\ell = [0, 0, -1]^\top$ ). This yields the celebrated *eikonal equation*:

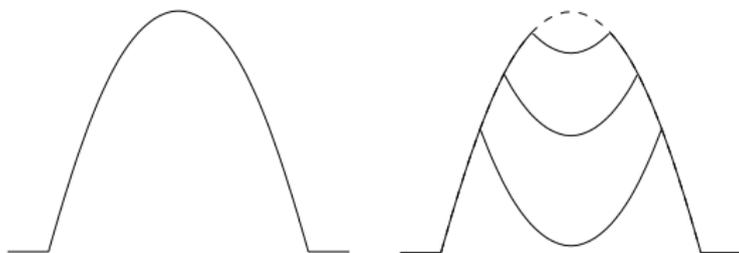
$$|\nabla z| = \sqrt{\frac{1}{\rho^2} - 1}$$

# The concave/convex ambiguity

The eikonal equation

$$|\nabla z| = \sqrt{\frac{1}{f^2} - 1}$$

characterizes the gradient's magnitude, but not its sign: this is the concave/convex ambiguity:



*Left: example of the 1D-surface  $z = 1 - x^2$ . Right: Under vertical lighting, three other solutions (amongst an infinity), which are differentiable almost everywhere, satisfy the same eikonal equation.*

# The concave/convex ambiguity

Humans solve the concave/convex ambiguity intuitively because we learnt that the world is mostly “convex”



*A volcano in Hawaii (copyright Whitman Richards), looking like a crater if the image is rotated by  $\pi$ .*

Among all the solutions to the eikonal equation, it is thus reasonable to look for the “maximum” one

# Solving the eikonal equation

The *maximum viscosity solution* of the discrete eikonal equation

$$|\nabla z_{i,j}| = f_{i,j} := \sqrt{\frac{1}{l_{i,j}^2} - 1}$$

can be numerically approximated using Fast Matching, semi-Lagrangian methods, or the following Lax-Friedrichs Hamiltonian scheme:

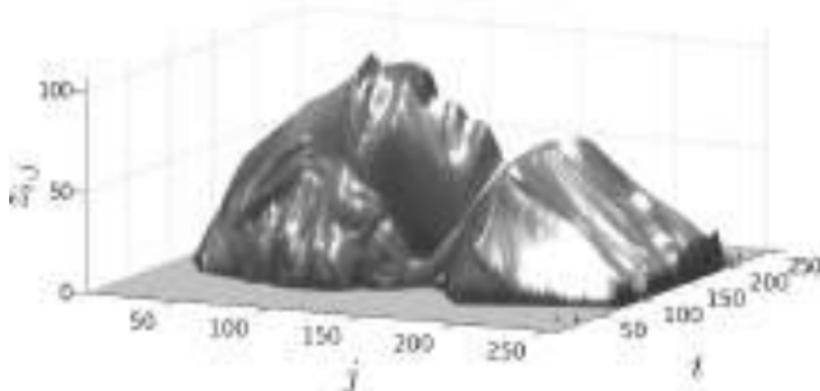
$$z_{i,j}^{(k+1)} = \frac{z_{i-1,j}^{(k)} + z_{i+1,j}^{(k)} + z_{i,j-1}^{(k)} + z_{i,j+1}^{(k)}}{4} - \frac{1}{2} \left( \sqrt{\left( \frac{z_{i+1,j}^{(k)} - z_{i-1,j}^{(k)}}{2} \right)^2 + \left( \frac{z_{i,j+1}^{(k)} - z_{i,j-1}^{(k)}}{2} \right)^2} - f_{i,j} \right)$$

with Dirichlet boundary condition  $z_{i,j} = g_{i,j}$  if  $(i,j) \notin \Omega$

## Example of SfS result

Example of result at convergence of the previous scheme, on a real-world dataset:

3D-reconstruction at iteration 512



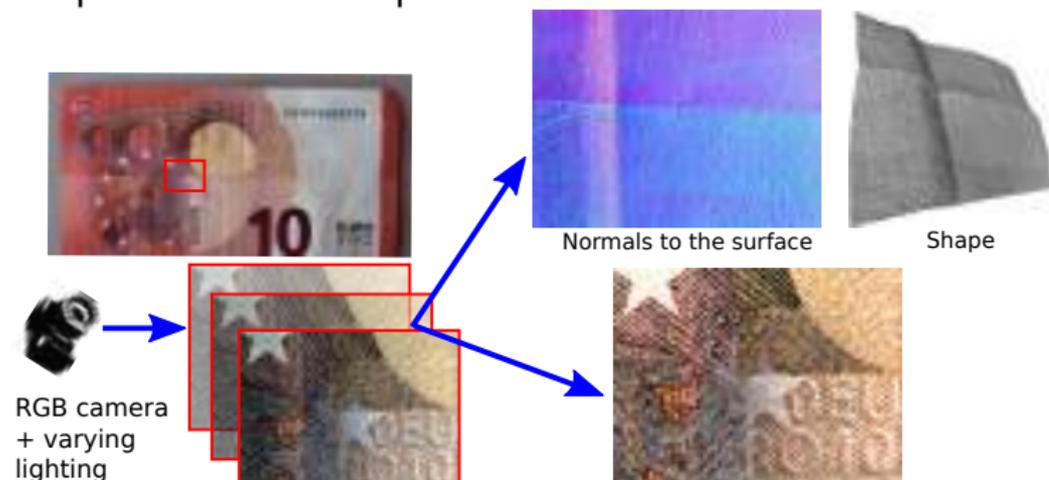
Subtle geometric details are recovered, yet there is a clear low-frequency bias, and most real-world scenes do not have uniform reflectance.

Most natural solution: capture multiple images, under varying illumination.

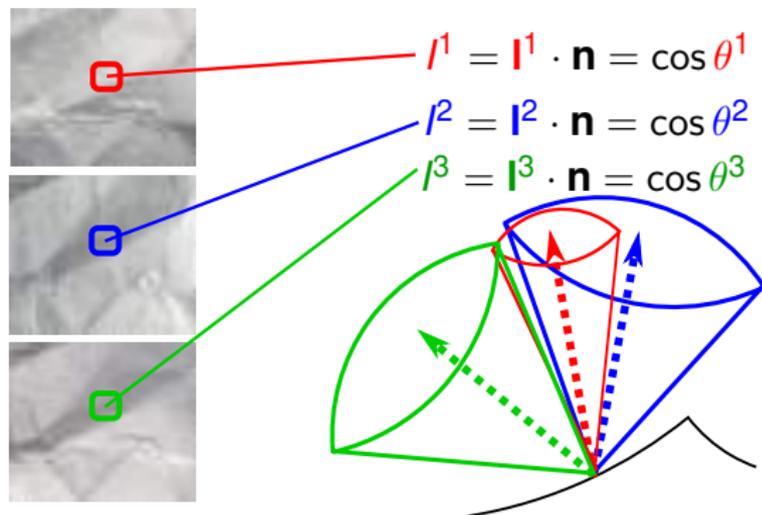
# Photometric stereo

When no prior knowledge is available, it is hopeless to achieve a reasonable 3D-reconstruction based on SfS

Photometric stereo (Woodham, 1978) is an extension of SfS which considers multiple images of the surface, taken from the **same viewing angle** but under **varying lighting**. It can unambiguously recover geometry, and it is the only shape-from-X technique which is able to **estimate reflectance**

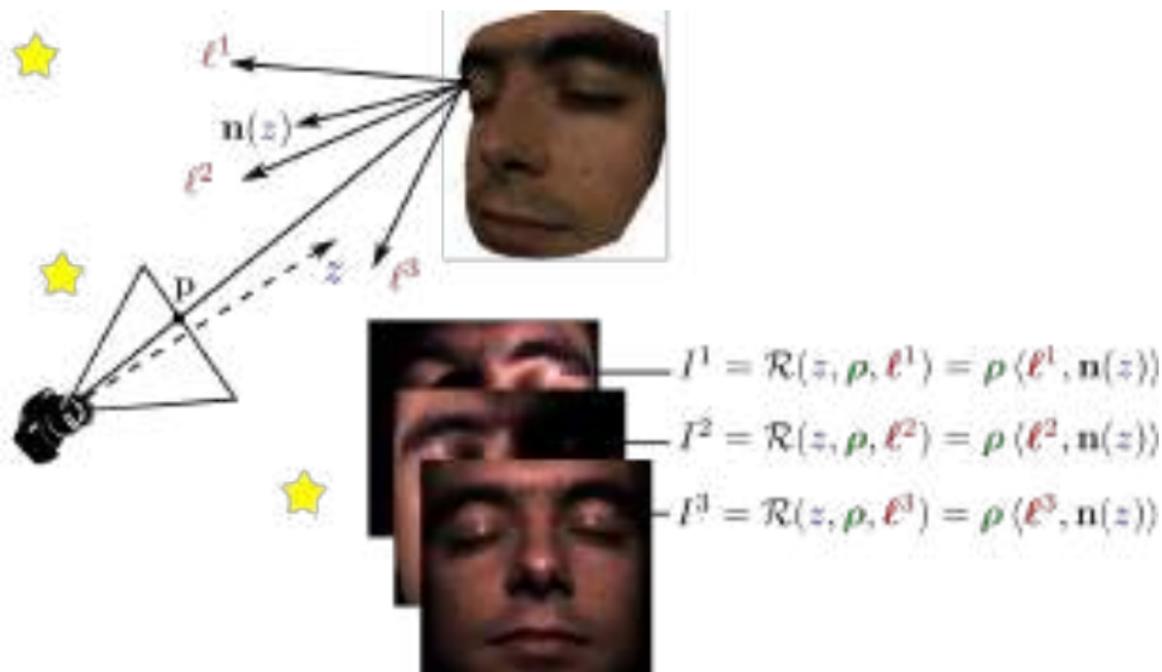


# Well-posedness of photometric stereo

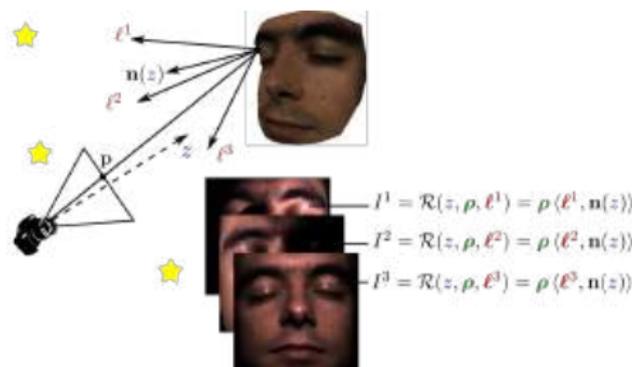


- $m = 1$  image (SfS), known albedo: infinitely many possible unit-length normals in each point
- $m = 2$  images, known albedo: up to two possible normals
- $m \geq 3$ : unique approximate solution - problem is over-constrained (thus albedo can be estimated !)

# Linear photometric stereo model



# Local estimation of shape and reflectance



Linear system in  $\mathbf{m}(\rho, \mathbf{z}) := \rho \mathbf{n}(\mathbf{z}) \in \mathbb{R}^3$ :

$$\begin{bmatrix} I^1 \\ I^2 \\ \vdots \\ I^m \end{bmatrix} = \begin{bmatrix} \ell^{1T} \\ \ell^{2T} \\ \vdots \\ \ell^{mT} \end{bmatrix} \mathbf{m}(\rho, \mathbf{z})$$

- If: a)  $m \geq 3$ , b) lighting vectors  $\ell^i$  are known, and c) they are non-coplanar,
- Vector  $\mathbf{m}(\rho, \mathbf{z})$  can be estimated in each pixel,
- Then we deduce the normal **and albedo** by

$$\rho = |\mathbf{m}(\rho, \mathbf{z})|$$

$$\mathbf{n}(\mathbf{z}) = \frac{\mathbf{m}(\rho, \mathbf{z})}{|\mathbf{m}(\rho, \mathbf{z})|}$$

- Eventually, normals  $\mathbf{n}(\mathbf{z})$  are *integrated* into depth  $\mathbf{z}$

# Integrated variational approach

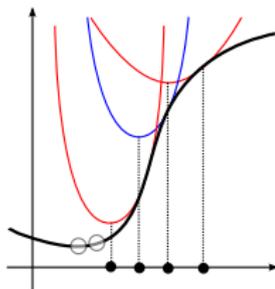
For robustness: direct estimation of depth  $z : \Omega \rightarrow \mathbb{R}$  and albedo  $\rho : \Omega \rightarrow \mathbb{R}$  solutions of

$$I^i = \rho \ell^i \cdot \frac{[\nabla z^\top, -1]^\top}{\sqrt{|\nabla z|^2 + 1}}, \quad i \in \{1, \dots, m\}$$

Linearisation  $\rho := \frac{\rho}{\sqrt{|\nabla z|^2 + 1}}$ , then

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

**Non-convex** problem: optimisation by alternating reweighted least-squares (alternating MM)



# 3D-reconstruction of metallic coins



1 euro (Italy)



50 cents (Spain)



1 yuan (China)



3D-reconstructions

# 3D-reconstructions of the human skin

Image 3D de ride



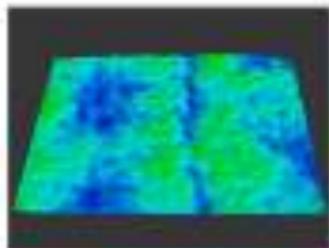
Image 3D de cicatrice



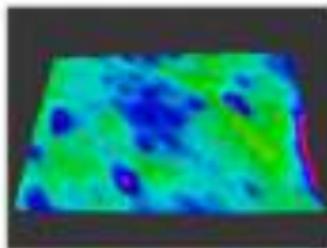
Image 3D d'acné



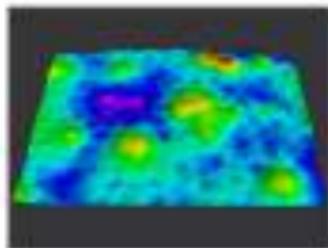
Carte d'élevation de ride



Carte d'élevation de cicatrice

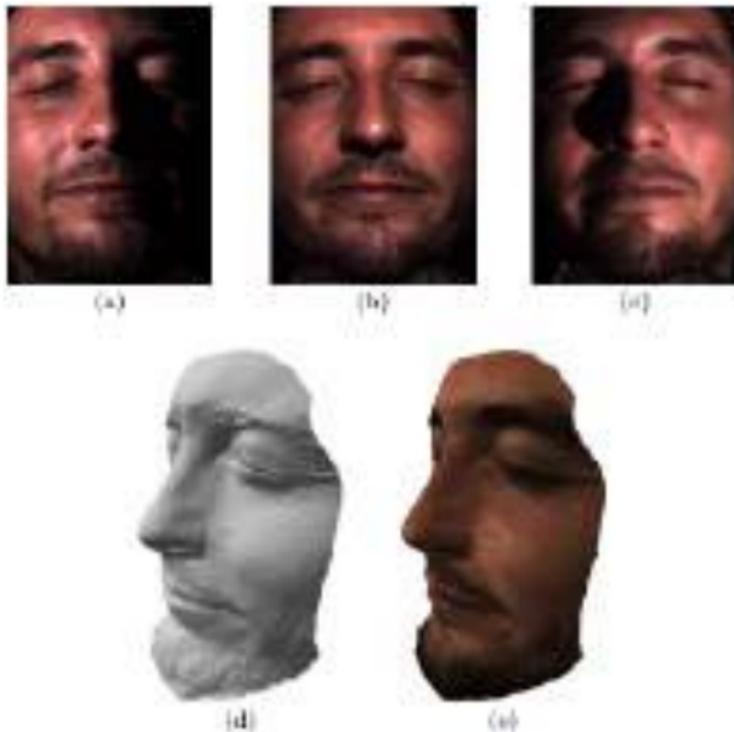


Carte d'élevation d'acné



Source : <http://www.pixience.com/produits-2/c-cube-recherche-clinique/module-3d/>

# 3D-reconstruction of a human face



Quéau, Durix, et al., "LED-Based Photometric Stereo: Modeling, Calibration and Numerical Solution", JMIV 18

# Outline

- 1 The Inclusive Museum Guide project
- 2 Shape-from-shading and photometric stereo
- 3 Acquisition campaign

# Recall our challenges



- 1) Thinness of the wool strings
- 2) Color changes
- 3) Glass creating reflections
- 4) Fragility to light exposure

Photometric stereo already provides an answer to 1)-2)

# Campaign preparation: handling the glass

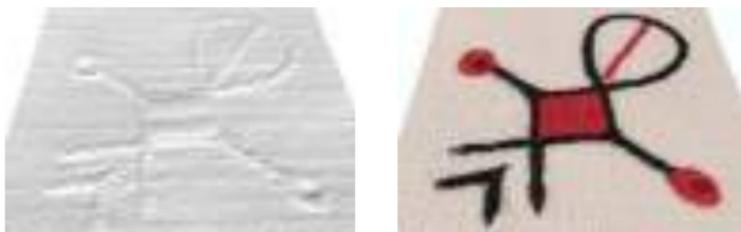


Proof of concept experiment by M. Pizenberg during lockdown # 1: self-made embroidery behind a glass, with stuck spheres for calibrating a moving flash light source

# Campaign preparation: handling the glass



Captured images



3D-reconstruction

Fine-scale structures are revealed. Glass does not perturbate 3D-reconstruction, as long as 1) the camera is fronto-parallel, 2) light directions deviate enough from the viewing direction, and 3) robust algorithm is used

## Campaign preparation: limiting light

The tapestry is fragile, so we need to limit the luminous flux we will project onto it → very short camera exposure time of  $1/200s$ , and use of a synchronized flash light source



## Campaign preparation: measuring light

To get access to the tapestry, we had to obtain agreement from authorities, announcing how much light would be projected → need to measure the luminous flux emitted by the flash, by calibrating the camera response and using a lux-meter



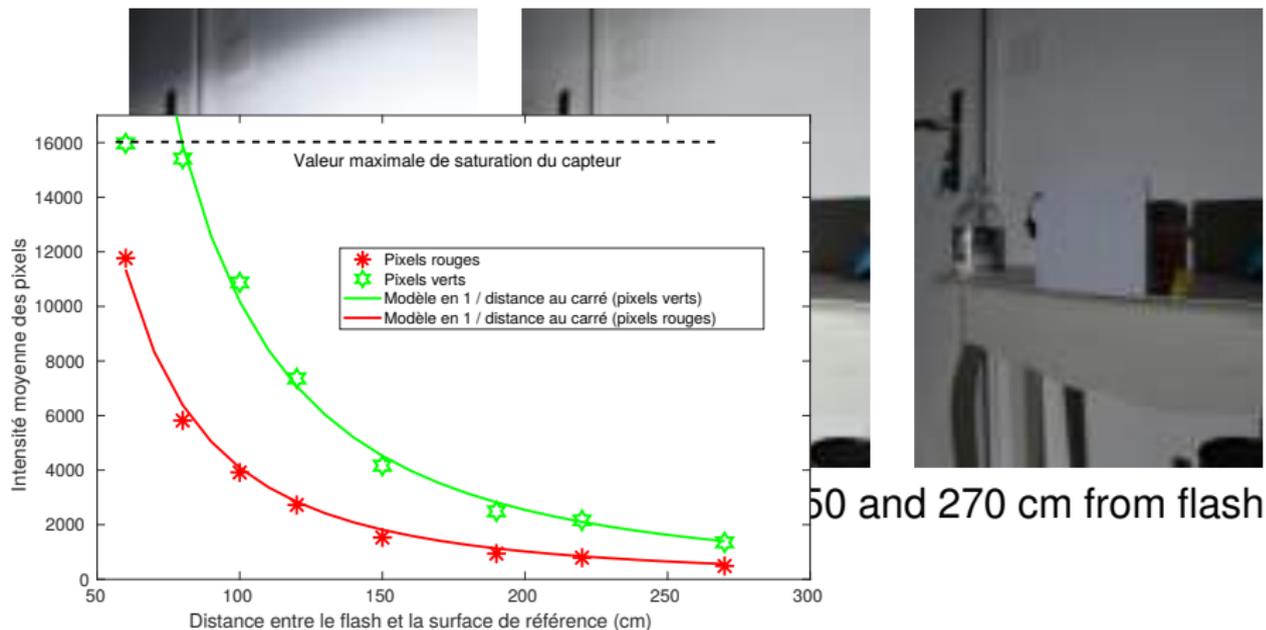
Instant light intensity measured by a lux-meter, to verify the sensor linearity

## Campaign preparation: measuring light



Images of a Lambertian plane at 60, 150 and 270 cm from flash

# Campaign preparation: measuring light



→ total surfacic energy at 1 m, for 12 shots per scene: 312 lux.s  
( $\approx$  6 sec. of the already in-place lighting system)

⇒ acceptable for authorities

# Data acquisition

Selection of 12 parts of the tapestry for digitization, and in-situ image acquisition in January 2021:

Acquisition of the death of Harold sequence

# The death of Harold sequence

Acquired data and 3D-reconstruction:



# The death of Harold sequence

Acquired data and 3D-reconstruction:



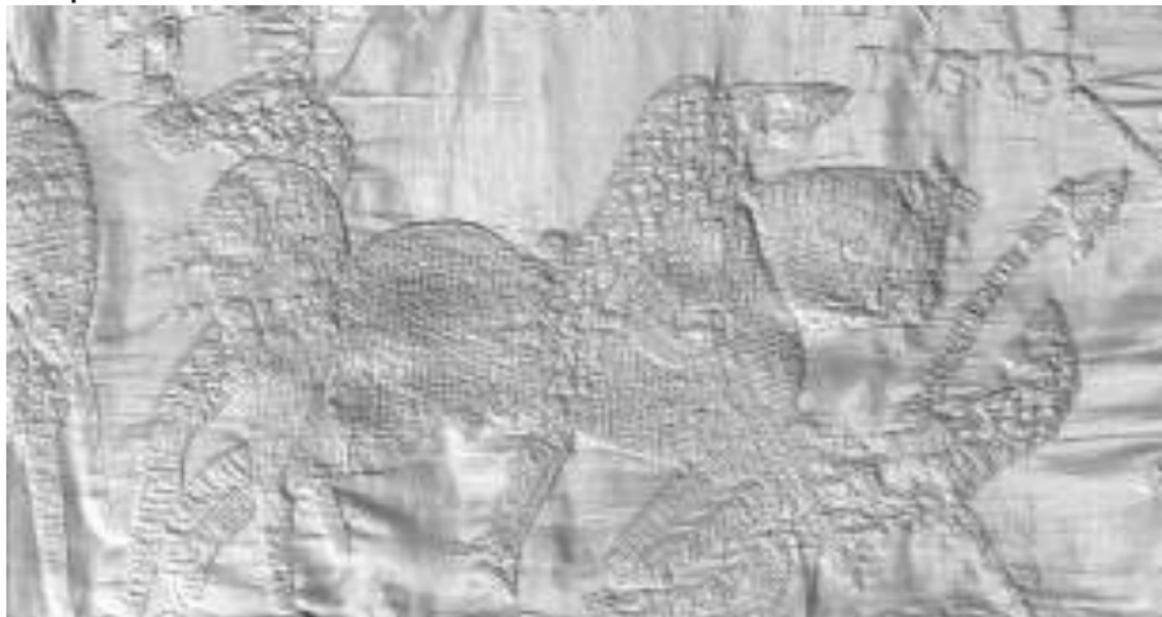
# The death of Harold sequence

Acquired data and 3D-reconstruction:

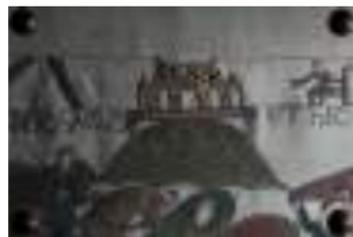
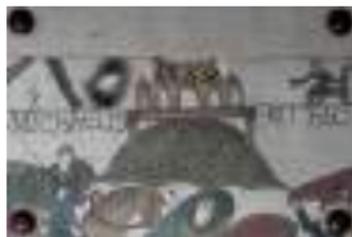
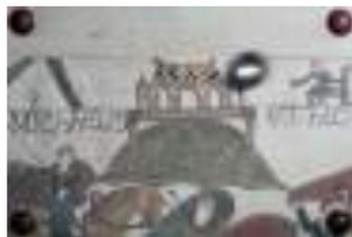


# The death of Harold sequence

Acquired data and 3D-reconstruction:



# The Mont St-Michel sequence



# The Hastings sequence



From left to right: images (close-up), 3D-reconstruction, and estimated albedo (not used for the project)

What's next ?

We're now ready for 3D-printing :) !

Thank you for your attention !

## References

- Durou, J.-D., Falcone, M., Quéau, Y., and Tozza, S. *Advances in Photometric 3D-Reconstruction*. Springer, 2020.
- Quéau, Y., Durix, B., Wu, T., Cremers, D., Lauze, F., and Durou, J.-D. "LED-Based Photometric Stereo: Modeling, Calibration and Numerical Solution". In: *Journal of Mathematical Imaging and Vision* (2018).
- Quéau, Y., Mecca, R., and Durou, J.-D. "Unbiased Photometric Stereo for Colored Surfaces: A Variational Approach". In: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- Quéau, Y., Wu, T., Durou, J.-D., Lauze, F., and Cremers, D. "A Non-Convex Variational Approach to Photometric Stereo under Inaccurate Lighting". In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.