



# Image Processing for Cultural Heritage Accessibility: Digitizing the Bayeux Tapestry

Matthieu Pizemberg, Abderrahim Elmoataz, Yvain Quéau

## ► To cite this version:

Matthieu Pizemberg, Abderrahim Elmoataz, Yvain Quéau. Image Processing for Cultural Heritage Accessibility: Digitizing the Bayeux Tapestry. 2022. hal-03684035

HAL Id: hal-03684035

<https://hal.science/hal-03684035>

Preprint submitted on 1 Jun 2022

**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.

# Image Processing for Cultural Heritage Accessibility: Digitizing the Bayeux Tapestry

Matthieu Pizenberg<sup>1</sup>, Abderrahim Elmoataz<sup>1</sup>, and Yvain Quéau<sup>1,\*</sup>

<sup>1</sup>Normandie Univ, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France

\*yvain.queau@ensicaen.fr

## ABSTRACT

This work showcases image processing and computer vision algorithms, in the light of accessibility to cultural heritage artifacts. We emphasize the potential of multimodal 2D image registration and fine-scale 3D-reconstruction techniques, with the aim to ease the work of historians and museum curators, as well as to make artifacts more accessible to the general public or to visually impaired people. This study focuses on the Bayeux Tapestry, a world-famous medieval wool embroidery included in UNESCO's Memory of the World register, and of fundamental importance for both the scientific community and the general public. This exceptional testimony on the society at the eleventh century is both a singular artwork and a historical source on a major event in the history of medieval Europe. Developing state-of-the-art image processing tools for its digitization will ease not only its access, but also its analysis, inspection and reproduction.

## Introduction

The Bayeux Tapestry is a unique historical document and a unique artwork. Indeed, it is the only embroidery of this importance that has reached us almost intact, crossing nine and a half centuries while others are known only from historical textbooks. This document of the eleventh century is a 70m long and 50cm high wool embroidery on a linen canvas. Throughout 55 scenes, it tells the epic of William, Duke of Normandy, who left Normandy with his armada in 1066 to conquer the kingdom of England (Figure 1), leading to the famous battle of Hastings where William defeated Harold (Figure 12) who had seized the throne in his place following the death of old King Edward (Figure 8).



(a) Overview of the first half of the Tapestry, as it is exhibited in the Bayeux Tapestry Museum.

Credits: Bayeux Tapestry Museum, City of Bayeux.

(b) William, Duke of Normandy, as depicted in scene 23.

(c) Close-up on William's armada crossing the Channel to conquer the kingdom of England (scene 37). Credits: Bayeux Tapestry Museum, City of Bayeux.

**Figure 1.** The Bayeux Tapestry in its exhibition hall.

This unique artifact of the Middle Ages is a contemporary testimony of the way people used to live in the eleventh century. It is also a historical documentation of events that changed the course of history for two nations (France and England), and for Europe. As such, it is of interest to textual historians, to art historians, to specialists in ancient textiles and architecture, to archaeologists, not to mention the general public. It is the first furniture to have been classified as a historical monument in 1840, and it was included in the UNESCO "Memory of the World" register in 2007. This artifact, which is the property of the French state, is part of the "Musée de France" collection of the city of Bayeux, and it has been exhibited at the Centre Guillaume le Conquérant (Bayeux Tapestry Museum) since 1983.

The Tapestry has been an object of investigation, research and hypothesis for almost three centuries. It was and it is still studied by historians, academics, novelists and journalists. It has been the subject of many questions about its sponsor, its date and place of manufacture, etc. However, researchers and scientists interested in the study of this artifact are confronted with problems related to temporal or geographical constraints i.e., to accessibility:

- The Tapestry is accessible to researchers only one month a year, in January, when the museum is closed to the general public. The rest of the time, research is complicated by the number of visitors (more than 400 000 in 2019).
- Studies rarely consider the artifact in its overall structure. The exceptional size of the document (70m long and 50cm high, see Figure 1a), the profusion of scenes (55), the variety of characters and decorations, the texts in Latin associated with the images: all of this forces specialists to focus on a precise portion of the artwork.
- As a furniture, it raises questions about the stages of the artistic process of its assembly and realization. For example, the possibility of detecting obvious technical changes between different parts of the Tapestry, may indicate that they were embroidered in several workshops. Some of these questions can only be answered by a careful examination of the face and reverse of the whole Tapestry. Yet, since the last digitization in 1983, this reverse side is protected and no researcher or curator can access it.
- There is a lack of a consensual reference system to find one's way around the work. Some authors choose to propose reproductions of the Tapestry in their works and to refer to it by illustrated images. Others try to rely on geolocalization systems more or less known by the readers in order to direct them (the scene numbering present on the canvas being part of the Tapestry).
- The embroidery being extremely fragile, it is protected by a glass (see Figure 1c), preventing its tactile exploration. Therefore, it remains inaccessible to visually impaired persons.

As part of a restructuration project of the Bayeux Tapestry Museum by 2024, a digitization of the face of the tapestry was jointly carried out in 2017 by La Fabrique du Patrimoine en Normandie and by the GREYC laboratory (Normandy University, UNICAEN, ENSICAEN, CNRS). These scans have generated a mass of data under the form of 86 high resolution, multispectral images corresponding to different modalities, as well as 3D point clouds. The objective of this digitization is to circumvent the aforementioned accessibility constraints, by providing open-access 2D and 3D digitizations for different modalities. Along with image processing tools, such data will allow scientists, curators and the general public to visualize the whole embroidery and navigate through it with a very high precision, while discovering elements that were previously inaccessible. One could then observe the artifact on their personal tablet or computer, or connect to a server at the Bayeux museum in order to study it without being subject to geographical and temporal constraints. These tools also have the advantage of increasing the possibilities in terms of cultural mediation and offering interactive ways to involve the public in the discovery process, making them more active in their approach to knowledge. For instance, the construction of a 2D panorama in daylight<sup>1,2</sup> may already serve as a background image for any indexing or information system, hence solving the geo-referencing problem. Lastly, the high-resolution 3D-reconstruction of the tapestry will allow to both virtually visualize the Tapestry in augmented reality applications, and to realize faithful 3D-printed mockups, which may be touched without damaging the original artwork.

In the present paper, we describe the image processing tools which were developed for facilitating the access to the Tapestry to scientists and to the general public. We first elaborate on 2D imaging technologies allowing to navigate through the artwork in an offline manner, through a large-scale, multimodal panorama and semi-automatic tools for registering historical layers. Then, we demonstrate the potential of 3D computer vision tools for virtually relighting it in surface inspection or augmented reality applications, and for realizing high-detail 3D-printed mockups. A conclusion on this digitization work lastly follows, along with a discussion on accessibility of the softwares developed during this project.

## Registering Multimodal Digitizations

The Bayeux Tapestry is extraordinary in many regards, yet one of its specificities is its exceptional size (70m long and 50cm high). Providing a way to the general public to visualize it without direct access is therefore a nontrivial problem. To that aim, a digitization campaign was carried out by La Fabrique de patrimoines en Normandie in January 2017. This digitization had several goals: serve as reference for studying future evolution of the artifact; serve as basis for different research programs; feed a geo-referenced image database; and provide high-resolution images of the artifact for research, cultural, education and communication purposes. The campaign resulted in a series of 86 high resolution images (roughly 8000 × 5500 pixels), in multiple modalities such as daylight, infrared and ultraviolet (fluorescence and reflection). The daylight images enable a first high definition digital coverage of the Tapestry, by opening the way to the constitution of a readily available reference background<sup>1</sup>. Yet, other images obtained with different modalities bring complementary information which can help in the restoration or the study of the Tapestry, as previously done for instance in the context of wall paintings studies<sup>3</sup>. For instance,

the ultraviolet fluorescence images highlight some restoration materials. Other images such as those obtained by ultraviolet reflection, converted into false colors can highlight alterations of the support of the tapestry such as stains. In practice though, capturing the different modalities involves manual operations on the sensor, which induces spatial displacements of up to a few hundreds of pixels. For these reasons, creating a large-scale panorama of the artifact under different registered modalities, as illustrated in Figure 2, would be worthwhile yet nontrivial.



**Figure 2.** Multi-modal panorama of the beginning of the Tapestry. Five different imaging modalities are represented here with different colour codes. They are both spectrally registered at pixel accuracy, and spatially registered into a large-scale panorama which can be explored in a web interface.

Two questions thus arise: 1) how to spatially register (i.e., stitch together) the 86 images of each modality, and 2) how to spectrally register (at pixel accuracy) the different modalities. Existing panorama creation softwares fall short in our use case for two reasons. First, an image of the full embroidery length would amount to approximately  $680.000 \times 5500$  pixels, meaning around 3 gigapixels. Such sizes bring technical difficulties due to memory consumption of stitching algorithms and incompatibilities with many image formats (the JPG format, for example, cannot handle more than 65535 pixels width or height). The second reason is that existing panorama softwares do not handle registrations of multiple modalities. Consequently, we developed our own processing pipeline in order to create a multi-modal panorama from the  $N_{\text{frames}} = 86$  high resolution ( $8000 \times 5500$  pixels) images under  $N_{\text{modalities}} = 6$  different modalities. We considered the two following peculiarities of our problem:

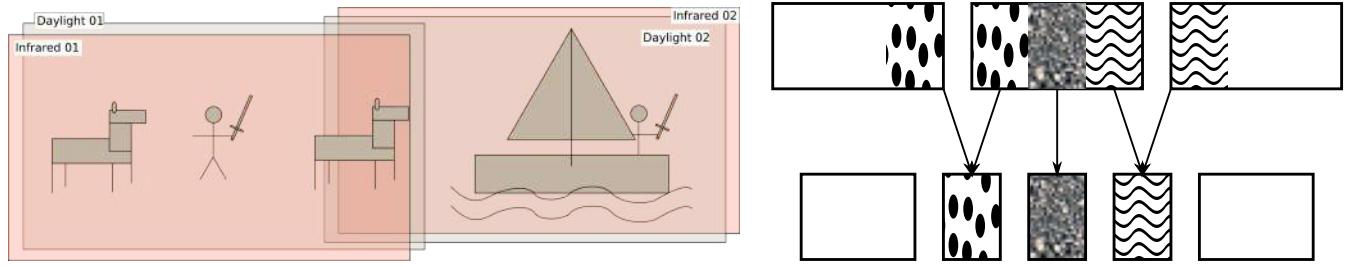
- Every image in one modality has approximately the same framing and contents as the corresponding images in the other modalities. Moreover, two corresponding images in two modalities only exhibit rigid deformations that can be modelled by small camera displacements. Indeed, during capture, the embroidery stayed still between the changes of filters on the sensors.
- Two consecutive images in one modality share roughly 30% overlap (see Figure 3, left), and the overlapping region between two consecutive images presents non-rigid deformations. This is because the camera setup was still, while the embroidery was slid on a rail, leading to folds and other non-rigid image deformations.

These observations tend to indicate that we can proceed in a two-steps procedure: 1) achieve spectral registration frame-by-frame, so as to register each scene independently in the same coordinates as the daylight modality, which serves as reference; 2) create a panorama in one modality (daylight), and apply the same panorama deformations to the other, already registered, modalities.

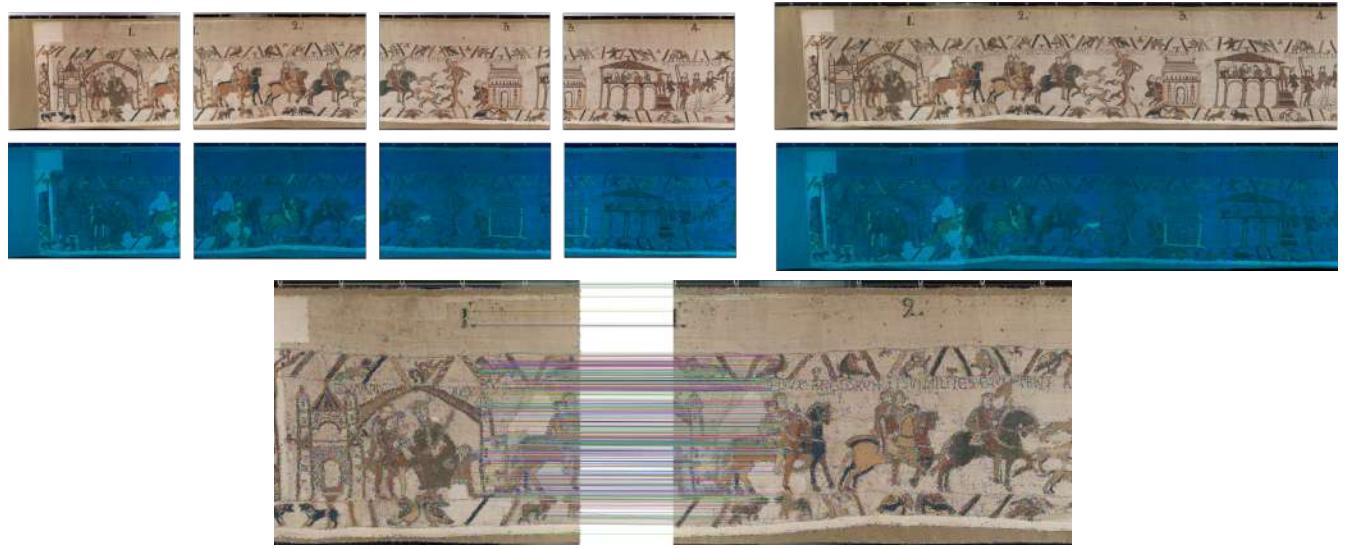
**Spectral Registration** For each frame, the problem consists in registering a series of  $N_{\text{modalities}} = 6$  images of different modalities together, where these images have roughly the same frame and the displacement between frames can be assumed rigid. In some sense, these images have been captured under unknown, varying lighting, as it is done for instance in the context of photometric stereo<sup>4</sup>. Let us denote by  $N_{\text{observations}}$  the number of pixels of the reference (daylight) frame, by  $\mathcal{T}$  the sought transformation, and by  $M \in \mathbb{R}^{N_{\text{observations}} \times N_{\text{modalities}}}$  a matrix stacking all the data. The registration problem can be recast as a rank minimization problem, and then relaxed as a convex optimization problem<sup>5</sup>:

$$\text{rank}(\mathcal{T}(M)) \rightarrow \min, \quad \underset{\text{Relaxation}}{\Rightarrow} \quad \min_{A, \mathcal{T}} \|A\|_* + \lambda \|A - \mathcal{T}(M)\|_1,$$

where the latter can be solved efficiently using augmented Lagrangian methods. After this spectral registration step, every coordinate  $(x, y)$  in each modality, should now correspond to the same location on the embroidery, as illustrated in the top-left of Figure 4.



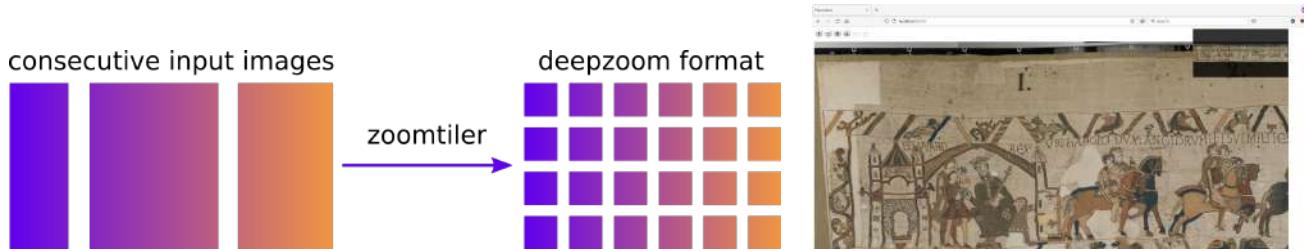
**Figure 3.** Left: Diagram representing two consecutive overlapping scenes, imaged under two modalities picturing the same contents with a small rigid deformation. The modalities need first to be registered (spectral registration), then the overlapping frames may be stitched into a panorama (spatial registration). Right: sketch of the panorama procedure. The input overlapping sequence is turned into a new sequence of non-overlapping images which may be concatenated into a panorama.



**Figure 4.** Top, left: four overlapping images, of size  $8000 \times 5500$  pixels, of the beginning of the Tapestry under daylight, and their registered fluorescence UV counterparts (revealing the restored parts). Top, right: reconstructed panoramas for the two modalities. Bottom: detected keypoints (circles) in the first and second daylight images, and keypoint matching (lines).

**Panorama Creation** Now that all images are aligned with the reference daylight modality, we may compute the full panorama in that reference modality. The algorithm presented here converts the input, overlapping sequence into a new sequence of images which may be smoothly concatenated into a panorama (see Figure 3, right). Each frame is split into roughly overlapping (30% of sides) and non-overlapping areas (central part), and the latter are copy-pasted in the output (cf. the “sand” texture in Figure 3, right). Then, the problem reduces to merging each right-side overlapping area with the left-side of the following scene as seamlessly as possible (cf. the “dots” and “wave” textures in Figure 3, right), which requires computing the deformation mapping from the left image to the right one. Since our scene is well textured, we follow a keypoint detection and matching approach (bottom row in Figure 4). On both images, we compute Akaze keypoints<sup>6</sup>, then compute pair matches between them, and filter the pairs using RANSAC<sup>7</sup>. Then, we create a new virtual coordinates frame and transform keypoint pairs into keypoint triplets, where the third virtual keypoint is on the line segment between the first and second keypoints with the following expression:  $p_3 = (\beta * p_1 + \alpha * p_2) / (\alpha + \beta)$  where  $\alpha$  is the distance from the left on the first image, and  $\beta$  is the distance from the right on the second image. Finally, we warp both images into this new virtual target image using the moving least squares algorithm (MLS)<sup>8</sup> with respectively the  $(p_1, p_3)$  and  $(p_2, p_3)$  pairs of keypoints. To avoid seams, the resulting image is a weighted mean of both images, such that it equals the left (resp., right) image on its left (resp., right) side, and blends both images in its central part. Once the panorama fusion has been computed on the reference daylight modality, we can re-apply exactly the same transformations on other modalities, since they have all been previously aligned with daylight images. For this to work we simply need to save the keypoints and warping parameters used for the MLS daylight warps. An example of concatenation obtained with this procedure is presented in the top right of Figure 4.

**Visualizing the Panorama** Concatenating the whole set of images would lead to a single image of 3 gigapixels. This is not pragmatic, both since very few image formats can hold such an image (one notable exception being TIFF), and since most image viewers choke while loading it. An alternative strategy is to use a zoomable format, and a dedicated viewer able to display such a format. Many options exist, and we chose the combination of the Deepzoom format<sup>9</sup> and the OpenSeadragon web visualization tool<sup>10</sup>. Deepzoom is composed of a pyramid of image tiles (typically  $512 \times 512$ ) where each level of the pyramid halves the image resolution, resulting in roughly four times less tiles than the previous level. All the tools we found to generate the Deepzoom tiles need as input a single image, but at the same time struggle with images as big as the one we can generate for the Bayeux Tapestry. We thus developed a custom made program called Zoomtiler<sup>11</sup> able to take as input our collection of consecutive images and output a hierarchy of image tiles following the Deepzoom format (see Figure 5). Finally, once a Deepzoom image has been created for each modality, we may view them with the OpenSeadragon Web viewer. It enables panning and zooming with dynamic loading of the relevant tiles depending on the current zoom level and viewing area, while also switching between the different modalities and preserving the current location.



**Figure 5.** Left: Zoomtiler, a simple CLI program to generate zoomable tiled images from a horizontal succession of images. Right: OpenSeadragon Web image viewer used on a small section of the Bayeux Tapestry. Users may navigate through the entire content of the embroidery with their mouse, and also easily switch between the different imaging modalities, as illustrated in Figure 2.

**Registering the Back of the Tapestry** In the previous paragraphs, we described how to create a large-scale, multimodal digital copy of the Tapestry as it can now be seen when actually visiting the museum. On the other hand, other images of the artwork are not available to the general public, yet of infinite value. One example is the back of the Tapestry, which has never been exposed to the general public, and remains unavailable even to scientists since its last study in 1983. During the previous digitization campaign, the opportunity arised to take pictures of this back. The same panorama generation algorithm could then be applied, allowing the researchers to also navigate through this usually unavailable piece of data. Navigating through this new panorama can be very useful for understanding the creation process of the artwork. Indeed, while the front of the Tapestry (see Figure 6, left) tells the History, its back informs us about the story of its manufacture. As illustrated in Figure 6, right, one can for instance visualize the different restorations carried out through the ages.

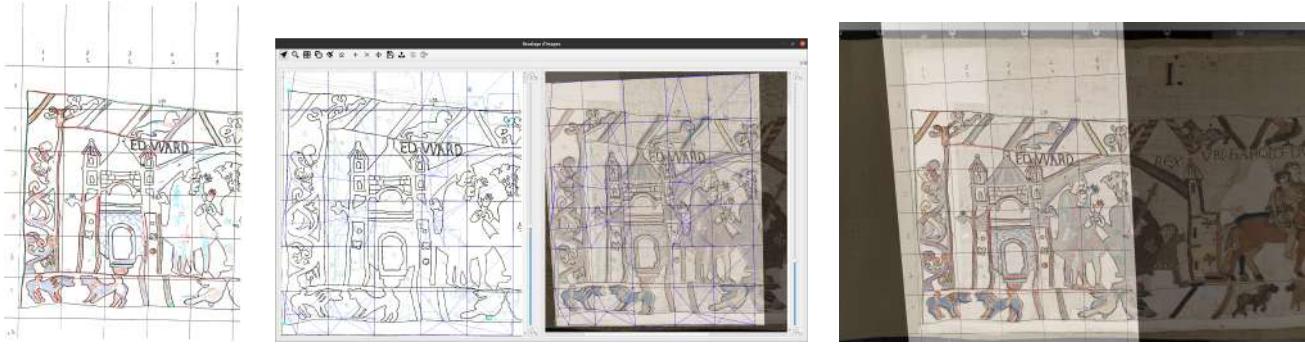


**Figure 6.** Reconstructed panoramas of the front (left) and back (right) sides of the Tapestry under daylight, visualized at different scales. Navigating through the front of the artifact allows one to follow the epic of William; on the flip side, navigating through its back allows one to follow the different stages of the manufacturing process, for instance by exploring parts which have been restored.

**Semi-automatic Registration of Historical Layers** During the winter of 1982-1983, the Tapestry had been transferred from its former location (hôtel du Doyen) to the current one. On this occasion, a condition assessment of the artwork was carried out and a series of layers of the back side of the artifact was manually created by the curators. They first drew a grid pattern forming a tracking system scaled such that each square of the layer's grid corresponds to a 10 cm by 10 cm square on the Tapestry. Thanks to this tracking system, curators were able to precisely manually draw the back side of the Tapestry. Finally these layers were used to annotate a large number of elements such as restoration pieces, holes, tears, seams, samples, or stains. Those annotations constituted a reference of the state of the artifact, which was intended to serve as basis for the evaluation of future deteriorations. A registration of these layers on the digital reference background will allow a more detailed analysis. Therefore, a specific tool has been developed in order to proceed to its registration on the previously discussed reference panorama. For each layer:

1. The grid is automatically detected using image binarization followed by Hough transform line detection;
2. The image of the layer is projected onto the corresponding scene in the panorama. This operation is performed semi-automatically, the user choosing four noncolinear points on the layer image and their counterparts on the panorama image;
3. The nonrigid transformation between the images is calculated and the two images are aligned;
4. We apply the same transformation to the grid that we have detected.

By iterating this process, we can build a panorama of layers and grids for the whole tapestry, as illustrated in Figure 7. This example shows that, logically, the grids are no longer perfectly rectilinear: it is an illustration of the difficulty of these different image processing operations, due to the very nature of the work, which is moving, and for which there are no perfect and fixed illustrations. Tensions due to the embroidery and its mode of presentation (vertical suspension, posed on a horizontal support) as well as its important length, have always induced deformations that all imaging operations on the Tapestry are subject to. Indeed, no Bayeux Tapestry imaging campaign to date has ever been able to deviate from this rule; to do otherwise would require the Tapestry to be "stabilized" in a room at least 70 m long, with the security conditions that are necessary for the preservation of this major artwork.



**Figure 7.** From left to right: one of the digitized layers from the 1982-1983 campaign; overview of the registration program, where the detected grid and matching points are highlighted; registration of the layer onto the panorama.

## Fine-scale 3D Digitization

So far, we have shown how to use image processing algorithms to enable visual inspection of the Tapestry on a multimedia device, so as to guarantee accessibility to the artwork for the general public and scientists. However, these tools only reveal what is “visible”, while the geometry of the surface might also be of interest. Indeed, allowing one to touch the Tapestry could allow visually-impaired people to “feel” the wool string and linen canvas composing the artifact, without actually seeing it. Obviously, this is not allowed because it would cause serious damages, hence a glass is placed before the artwork to ensure its protection (see Figure 1c). A workaround consists in creating highly-detailed 3D-printed copies of some scenes of the Tapestry, which could be touched by anyone or used in relighting applications, as illustrated in Figure 8. This, on the other hand, raises a nontrivial challenge: how to reconstruct in three dimensions the very thin geometric details of the surface (wool strings), without damaging it and without a direct access to it (presence of the protective glass) ?



**(a)** Two images of King Edward (scene 1), captured under varying lighting. Reflective spheres are stick onto the protection glass to calibrate the lighting directions.

**(b)** Printable 3D-model computed using photometric stereo.

**(c)** Relighting of the 3D-model, applying a gold reflectance to the surface.

**Figure 8.** From a set of photographies captured under the same viewing angle but varying lighting directions (left), the photometric stereo technique recovers a highly detailed representation of the surface geometry. This geometry can then either be printed in three dimensions to enable tactile explorations (center), or used in augmented/virtual reality applications (right).

**Existing Techniques for Surface 3D-reconstruction** Digitization of real-world assets in three dimensions is an open problem where choosing a reconstruction approach requires balancing the advantages and drawbacks of each method. The main criteria are the precision of the reconstruction, the subject scale and fragility, the setup cost and difficulty, as well as specific accessibility constraints of the subject. We are especially interested in digitizing the Bayeux Tapestry, which, as previously mentioned, is a thousand year old wool embroidery on a linen canvas. Preservation of this piece of European history is the main constraint, so coordinate measuring machines, with their probe-mounted moving arms, are obviously out of the equation. Other traditional 3D-scanning approaches for small and medium size objects include laser scanning<sup>12</sup>, structured light projections<sup>13,14</sup>, and photogrammetry<sup>15</sup>. However, for conservation purposes, we were not authorized to take the embroidery out of its protective glass and exhibition room for our 3D-digitization in January 2021. Therefore, laser and structured light projections could not be used, as the glass prevents these techniques to perform correctly. On the other hand, photogrammetry had already been used for the 3D-reconstruction of small oil paintings<sup>15</sup>. The setup is quite straightforward since one just needs a uniform lighting, a textured subject, and to take many images with a significant overlap. For a bigger subject however, it is highly recommended to create custom hardware such as the dedicated rack designed by Pamart et al. for the digitization of wall paintings<sup>3</sup>. The precision obtained by photogrammetry can reach subpixel accuracy for the control points matched in multiple images. For a sufficiently dense sensor, and with images close to the subject, this can result in high resolution 3D scans. In addition, many open-source photogrammetry consumer software exist, such as Alice Vision Meshroom<sup>16</sup>, making the technique easily available for non experts. We experimented with photogrammetric pipelines<sup>17,18</sup>, and found that photogrammetry was very efficient at recovering the overall structure of the scene's geometry (see Figure 9, yet it was not able to reveal the high-frequency details we are interested in for tactile experiments. Therefore, we turned our attention to another technique which has recently been shown to excel at revealing the finest details of cultural heritage artifacts<sup>19–27</sup>: photometric stereo.



**(a)** Dense point cloud representing the whole photogrammetric 3D-reconstruction.



**(b)** Close-up on a small part of the photogrammetric 3D-reconstruction.

**Figure 9.** The photogrammetric 3D-reconstruction of the Tapestry allows to virtually navigate through the entire Tapestry as if visiting the Museum, yet this 3D-reconstruction still lacks the finest geometric details.

**Photometric Stereo for Fine-scale 3D-reconstruction of the Tapestry** The photometric stereo technique was developed in the late 70s by Woodham<sup>4</sup>. It is a photometric 3D-reconstruction technique<sup>28</sup> i.e., a technique which aims at inverting a physical model describing the image formation process, rather than matching geometric features as in photogrammetry. In photometric stereo, 3D-reconstruction is achieved from a series of photos taken under the same viewing angle, while varying the lighting directions so as to create shading effects on the images. The camera being fixed, the precision of the 3D scan is mostly dependent on the sensor resolution: indeed, this approach provides geometric clues under the form of per-pixel surface normal estimates, which can then be integrated into a depth map<sup>29</sup>. The resolution of recent cameras tending to dramatically increase, one may have good hopes to recover structures as thin as the wool strings of the Tapestry. Besides, the setup for photometric stereo is quite straightforward and relatively cheap, as it only requires a remotely triggered camera, a tripod, and small spherical objects used for calibrating the light directions. Moreover, cultural heritage digitization teams are already very familiar with Reflectance Transformation Imaging (RTI) setups<sup>30</sup>, which are similar in spirit and can thus be used for the purpose of photometric stereo<sup>27</sup>. Another advantage of photometric stereo over photogrammetry is that photometric stereo enables better relighting options for the 3D scan, since it not only estimates the scene's geometry, but also its reflectance i.e., it decouples the intrinsic colour of the subject from the the lighting phenomena, as illustrated in Figure 10.



**Figure 10.** Left: 3 (out of 12) images of William revealing his face (scene 55). Right: using photometric stereo in the area highlighted in red, the input image sequence (one image of the sequence is recalled on the left) can be decomposed between reflectance (top, right) and normals (shown in false colors on bottom, right). Reflectance encodes the scene's color, independently from shape and illumination: see how the folds on the linen canvas have been eliminated. Normals, on the other hand, encode the scene's micro-geometry through the surface's orientation, which reveals thin details such as wool strings.

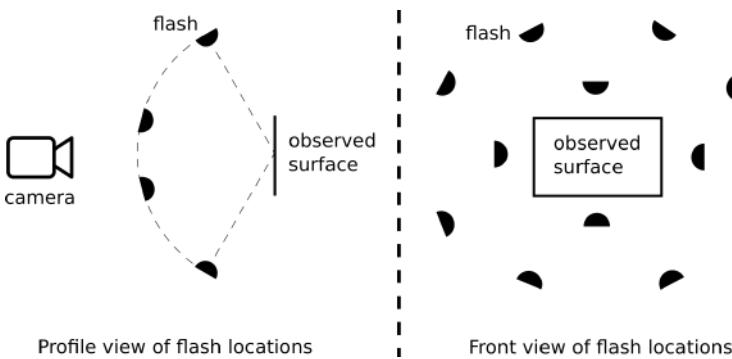
**Digitization Pre-study and Protocol** Photometric stereo infers geometry from the brightness variations arising in the images as the light direction changes. As such, getting sharp images without noise is key to the reconstruction quality. We thus decided to use a DSLR camera with high light sensibility, which can be remotely triggered and connected to a flash light (Figure 11d). An important point to mention is that a pre-study of the digitization protocol was performed and submitted to the "Direction régionale des affaires culturelles" (DRAC Normandie) to ensure there was no deterioration risk. In that report, we established that the energy brought by our flash light amounts to 6 seconds of the Bayeux Tapestry normal lighting conditions in its exhibition room. Apart from choosing an appropriate camera and light source, it is also extremely important to keep the exact same point of view between images. However, at the scale of a pixel, any sensor vibration due to manual trigger or floor vibrations due to walking around can create misalignments in the sequence of images. A strong and stable tripod, as well as remote triggers are thus very much recommended. In order to calibrate the location of the flash light for each image of a sequence, we placed small reflective spheres at the corner of the images. Since the artifact was kept behind its protective glass, we designed cheap, hand-made calibration spheres, by joining plastic door knobs and bathroom wall suction cups (Figure 11b), enabling easy displacements of the calibration spheres. Once the equipment is gathered, the digitization steps are the following: 1) place the camera on its tripod and adjust the frame; 2) place the reflective spheres on the image corners; 3) focus on the subject; 4) remotely launch a sequence of shots while changing the flash location between each shot. A minimum of 3 shots with non-co-planar light sources is required for photometric stereo but we went with 12 shots to enable robustness to outliers such as cast shadows and caustics induced by the spheres and glass panel (see Figure 12a, left). As illustrated in Figure 11c, eight of the sources were uniformly placed at an elevation angle of roughly 50° (the optimal value for the digitization of quasi-planar artifacts<sup>21</sup>), while four additional shots were taken with a smaller elevation angle to increase the number of outlier-free images.



**(a)** Digitization of the scene 55 (see also Figure 10). The camera is fixed on a tripod, and the flash is held on the side, directed toward the captured scene.



**(b)** Sphere stick on the glass panel.



**(c)** Distribution of flash locations on a hemisphere around the subject.



**(d)** Equipment for the 3D scan: a DSLR and a flash connected via a wireless radio command.

**Figure 11.** Equipment and protocol for the 3D-digitization of the Bayeux Tapestry by photometric stereo.

**Algorithm for 3D-reconstruction** Our goal is to turn a sequence of photometric stereo images into a highly detailed 3D-model, as depicted in Figure 12. To this end, we first estimate the incident lighting directions using the reflective spheres. The flash source being non-directional, each sphere provides a different estimate: we simply average them to have a reliable estimate at the center of the scene. Then, we simultaneously compensate the slight sensor's displacements and remove the outliers, by using a low-rank approximation procedure<sup>5</sup>. After this preprocessing, we iteratively estimate the reflectance, the normals and each source intensity, by semi-calibrated photometric stereo<sup>31</sup>. This provides us with a normal field which is integrated into a depth map using discrete cosine transform<sup>29</sup>. Since we wrongly assumed directional lighting for simplicity, the depth map exhibits a well-known low-frequency bias (“potato cheap” effect<sup>19</sup>). We found that this bias could be removed simply by fitting and subtracting a low-dimension polynomial to the depth map, which is simpler to implement than, e.g., resorting to photogrammetry<sup>32</sup>. Next, we recompute the surface normals using finite difference and integrate them again using discrete sine transform<sup>29</sup> to force constant values on the boundaries in order to ease 3D-printing. Lastly, the depth map is meshed and a small support is added, resulting in a printable 3D-volume (see Figure 8b). This pipeline is summarized in Figure 13.

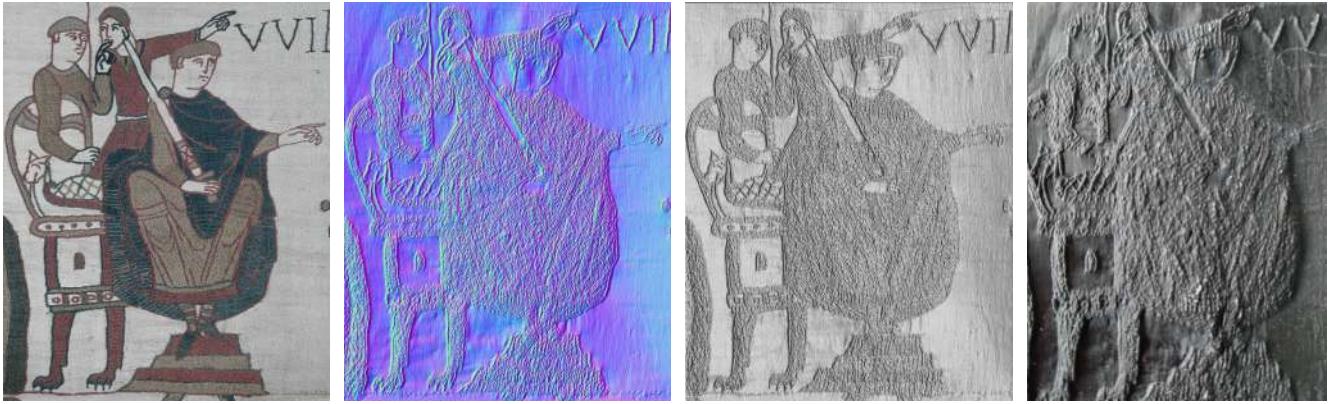


**(a)** Two images from the death of Harold sequence (scene 57). Note the shadows and caustics induced by the spheres and the glass.



**(b)** 3D-reconstruction obtained using photometric stereo, superimposed on one of the input images.

**Figure 12.** Highly detailed 3D-reconstruction using photometric stereo.



**(a)** Reflectance map estimated by photometric stereo. **(b)** Normal map (in false colors) estimated by photometric stereo. **(c)** Depth map obtained by integration of the normals. **(d)** Photography of a 3D-printed mockup.

**Figure 13.** Summary of the photometric stereo pipeline for creating 3D-printed mockups of the Tapestry, illustrated on the Duke William sequence (scene 27). From a series of pictures taken under varying lighting (one of these images is shown in Figure 1b), a reflectance map and a normal map are computed using photometric stereo. The latter is then integrated into a depth map, which can further be meshed and 3D-printed, in view of tactile experiments for visually-impaired people.

## Conclusion

In this article, we described a few image processing techniques which could be used to reinforce the accessibility to a unique cultural heritage artifact, namely the Bayeux Tapestry. We first introduced a 2D multimodal registration technique which enables the virtual visit of the artwork on a multimedia device, for the general public or for scientists interested in the study of the manufacturing process. Then, we elaborated on 3D-reconstruction techniques permitting to create 3D-models which may be used in relighting applications, or 3D-printed in view of tactile explorations by visually-impaired people. Future work will explore the possibility to bring such 3D-models onto force-feedback tablets<sup>33</sup>, so as to enable tactile experiments without the need for an actual 3D-printed mockup.

Beside the accessibility to the public and to scientists interested in this exceptional artifact, we are genuinely convinced that pushing for accessibility and reproducibility of research is important. Therefore, throughout our work, we have focused on the accessibility of the results and the reproducibility of the methods. For the multimodal panorama, we carefully described the algorithms, released code for some of the steps<sup>11,34</sup>, and are planning to release the full pipeline with even more details provided in its documentation. Similarly, for the 3D digitization pipeline, we are releasing a suite of web applications as well as their code to the public. The source code will enable further collaboration, and the web applications will provide re-usability for the broader cultural heritage community, less familiar with programming. Currently we are making two web applications available<sup>35,36</sup>. The first aims at performing light calibration of a sequence of photometric stereo shots by identifying bright spots on the reflective spheres and computing the light direction for each shot. The second aims at performing the 3D reconstruction when provided with a photometric stereo image sequence and its light calibration parameters, obtained with the first application. We hope that in the future, such tools will help others digitizing other unique cultural heritage artifacts.

## Data availability

The source codes of the softwares which were developed to produce the results on multimodal panorama creation and visualization, as well as fine-scale 3D reconstruction, are available at <https://github.com/mpizenberg>. Datasets are available from the corresponding author on reasonable request.

## References

1. Official digital representation of the Bayeux Tapestry - XIth century. Credits: City of Bayeux, DRAC Normandie, University of Caen Normandie, CNRS, ENSICAEN. Photos: 2017 – La Fabrique de patrimoines en Normandie. . <https://www.bayeuxmuseum.com/en/the-bayeux-tapestry/discover-the-bayeux-tapestry/explore-online/> (2021). [Online; accessed June 16th, 2021].
2. Berthelot, C. et al. La reconstruction du panorama de la «Tapisserie de Bayeux» comme fond de référence d'un système d'information documentaire spatialisée. Algorithmes de traitement d'images dédiés à la construction d'un panorama et au recalage d'images sur le panorama. *In Situ. Revue des patrimoines* (2019).

3. Pamart, A. *et al.* Multispectral photogrammetric data acquisition and processing for wall paintings studies. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, 559–566 (2017).
4. Woodham, R. J. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Image understanding systems and industrial applications I*, vol. 155, 136–143 (1979).
5. Pizenberg, M., Quéau, Y. & Elmoataz, A. Low-rank registration of images captured under unknown, varying lighting. In *International Conference on Scale Space and Variational Methods in Computer Vision*, 153–164 (2021).
6. Alcantarilla, P. F. & Solutions, T. Fast explicit diffusion for accelerated features in nonlinear scale spaces. *IEEE Transactions on Pattern Analysis Mach. Intell.* **34**, 1281–1298 (2011).
7. Raguram, R., Frahm, J.-M. & Pollefeys, M. A comparative analysis of RANSAC techniques leading to adaptive real-time random sample consensus. In *European Conference on Computer Vision*, 500–513 (2008).
8. Schaefer, S., McPhail, T. & Warren, J. Image deformation using moving least squares. In *ACM SIGGRAPH*, 533–540 (2006).
9. Microsoft. Deep Zoom. [https://docs.microsoft.com/en-us/previous-versions/windows/silverlight/dotnet-windows-silverlight/cc645050\(v=vs.95\)](https://docs.microsoft.com/en-us/previous-versions/windows/silverlight/dotnet-windows-silverlight/cc645050(v=vs.95)) (2011). [Online; accessed November 10th, 2021].
10. OpenSeadragon, an open-source, web-based viewer for zoomable images, implemented in pure JavaScript. <https://openseadragon.github.io/> (2011). [Online; accessed November 10th, 2021].
11. Pizenberg, M. Zoomtiler, a simple CLI program to generate zoomable tiled images. <https://github.com/mpizenberg/zoomtiler> (2021). [Online; accessed November 10th, 2021].
12. Levoy, M. *et al.* The digital Michelangelo project: 3D scanning of large statues. In *Conference on Computer Graphics and Interactive Techniques*, 131–144 (2000).
13. Zhang, Z. Microsoft kinect sensor and its effect. *IEEE Multimed.* **19**, 4–10 (2012).
14. Zaman, T., Jonker, P., Lenseigne, B. & Dik, J. Simultaneous capture of the color and topography of paintings using fringe encoded stereo vision. *Herit. Sci.* **2**, 1–10 (2014).
15. Chen, C., He, S., Chen, G. & Cao, H. Photogrammetry-based 3d printing reproduction method for oil paintings. *Int. J. Pattern Recognit. Artif. Intell.* **32**, 1854007 (2018).
16. Griwodz, C. *et al.* AliceVision Meshroom: An open-source 3D reconstruction pipeline. In *ACM Multimedia Systems Conference*, 241–247 (2021).
17. Lozes, F., Elmoataz, A. & Lézoray, O. PDE-based graph signal processing for 3-D color point clouds: Opportunities for cultural heritage. *IEEE Signal Process. Mag.* **32**, 103–111 (2015).
18. Elmoataz, A. & Lozes, F. Un cadre unifiant le traitement d’images et de nuages de points 3D : applications à la tapisserie de Bayeux. In *27ème Colloque GRETSI sur le traitement du signal et des images*, 793–796 (2019).
19. Huang, X., Walton, M., Bearman, G. & Cossairt, O. Near light correction for image relighting and 3d shape recovery. In *Digital Heritage*, vol. 1, 215–222 (2015).
20. Cheveau, A. 3D Scanning Solution for Textured Object using Photometric Stereo with Multiple Known Light Sources. In *Archiving Conference*, 1, 6–9 (2018).
21. Brenner, S., Zambanini, S. & Sablatnig, R. An investigation of optimal light source setups for photometric stereo reconstruction of historical coins. In *Eurographics Workshop on Graphics and Cultural Heritage*, 203–206 (2018).
22. Valach, J., Bryscejn, J., Fíla, T., Vavřík, D. & Štefcová, P. 3D Digitization of Selected Collection Items Using Photometric Stereo. In *Digital Cultural Heritage*, 31–36 (2018).
23. Vanweddenen, V. *et al.* A status quaestionis and future solutions for using multi-light reflectance imaging approaches for preserving cultural heritage artifacts. In *Euro-Mediterranean Conference*, 204–211 (2018).
24. Salvant, J. *et al.* Photometric stereo by UV-induced fluorescence to detect protrusions on Georgia O’Keeffe’s paintings. In *Metal Soaps in Art*, 375–391 (2019).
25. Xu, B. *et al.* 3d surface measurement and analysis of works of art. In *Asilomar Conference on Signals, Systems, and Computers*, 1779–1782 (2019).
26. Abzal, A., Saadatseresht, M., Varshosaz, M. & Remondino, F. Development of an automatic map drawing system for ancient bas-reliefs. *J. Cult. Herit.* **45**, 204–214 (2020).

27. Dulecha, T. G., Pintus, R., Gobbetti, E. & Giachetti, A. SynthPS: a Benchmark for Evaluation of Photometric Stereo Algorithms for Cultural Heritage Application. In *Eurographics Workshop on Graphics and Cultural Heritage*, 13–22 (2020).
28. Durou, J.-D., Falcone, M., Quéau, Y. & Tozza, S. A comprehensive introduction to photometric 3d-reconstruction. In *Advances in Photometric 3D-Reconstruction*, 1–29 (Springer, 2020).
29. Quéau, Y., Durou, J.-D. & Aujol, J.-F. Normal integration: a survey. *J. Math. Imaging Vis.* **60**, 576–593 (2018).
30. Malzbender, T., Gelb, D. & Wolters, H. Polynomial texture maps. In *Conference on Computer Graphics and Interactive Techniques*, 519–528 (2001).
31. Cho, D., Matsushita, Y., Tai, Y.-W. & Kweon, I. S. Semi-calibrated photometric stereo. *IEEE Transactions on Pattern Analysis Mach. Intell.* **42**, 232–245 (2018).
32. Karami, A., Menna, F. & Remondino, F. Investigating 3d Reconstruction Of Non-Collaborative Surfaces Through Photogrammetry And Photometric Stereo. *The Int. Arch. Photogramm. Remote. Sens. Spatial Inf. Sci.* **43**, 519–526 (2021).
33. Gay, S. L., Pissaloux, E., Romeo, K. & Truong, N.-T. F2T: A Novel Force-Feedback Haptic Architecture Delivering 2D Data to Visually Impaired People. *IEEE Access* **9**, 94901–94911 (2021).
34. Pizenberg, M. Rust implementation of image deformation using moving least squares. [https://github.com/mpizenberg/rust\\_mls](https://github.com/mpizenberg/rust_mls) (2021). [Online; accessed April 4th, 2022].
35. Pizenberg, M. & Vincent, F. Light calibration from reflective spheres. <https://github.com/mpizenberg/light-calibration-reflective-spheres> (2021). [Online; accessed April 4th, 2022].
36. Pizenberg, M. & Vincent, F. Calibrated quasi-planar photometric stereo in Rust and on the Web. <https://github.com/mpizenberg/calibrated-quasi-planar-photometric-stereo> (2022). [Online; accessed April 4th, 2022].

## Acknowledgements

This work was funded by the RIN grant Guide Muséal, and the ANR grant Inclusive Museum Guide (ANR-20-CE38-0007). A.E is partially supported by ANR Grant SUMUM (ANR-17-CE38-0004). Multimodal images were acquired by La Fabrique de patrimoines en Normandie. Access to the Tapestry was made possible by the Bayeux Tapestry Museum and the City of Bayeux, under the supervision of Clémentine Berthelot.

## Author contributions statement

All authors contributed to the preparation, writing and revisions of the manuscript.

## Additional information

The authors declare no competing interests.