Tailoring x-ray tomography techniques for cultural heritage research

Bossema, F.G.

1
Introduction

Visualizing the internal structure is a crucial step in acquiring knowledge about the origin, state, and composition of cultural heritage artifacts. Among the most powerful techniques for exposing the interior of cultural heritage objects is computed tomography (CT), a technique that computationally forms a 3D image using hundreds of radiographs acquired in a full circular range. The diversity in materials, shapes and sizes of cultural heritage objects poses a challenge for the broad application of this technique, since each scan needs to be tailored to the object. Moreover, the lack of affordable and versatile CT equipment in museums, combined with the challenge of transporting precious collection objects, currently keeps this technique out of reach for most cultural heritage applications. In this thesis, we explore how CT imaging can be further integrated in cultural heritage applications by 1) integrating expert feedback into the data acquisition process, 2) tailoring acquisition methods to obtain specific information, 3) developing a low-cost method to use in-house X-ray facilities for 3D CT imaging and 4) providing tools to interactively visualise and inspect the CT data.

In this introductory chapter, we will first discuss the difference between 2D and 3D imaging (section 1.1). We then present the concepts underlying each of the steps of the CT research process (section 1.2). Next, we will outline the possibilities and challenges of the application of CT imaging to cultural heritage objects (section 1.3). We close with the research questions underlying the chapters in this thesis (section 1.4).

1.1 2D and 3D imaging

1.1.1 2D radiography

In 2D X-ray imaging or *radiography*, X-rays travel through an object and a detector measures the intensity of the X-rays after they leave the object. Based on the length of the object and the density of the materials the X-rays are absorbed. Thus, a 'shadow-like' image of the interior is formed on the detector. On a single X-ray image, all features along a line from the source to the detector are superimposed on the image and depth information is lost. This makes it difficult to distinguish between features when they are behind higher density materials, because these occlude lower density features. Inspection of individual features within the object can therefore be challenging. 2D X-ray imaging is used in museums for the live radiographic inspection of objects to investigate possible internal damage or to obtain information about how it was made. In figure 1.1 radiographs of a wooden sculpture from the Rijksmuseum are shown, which was scanned for the purpose of dating the wood by inspecting the tree rings.

Figure 1.1: a) *Woman with lantern*, Rijksmuseum collection [159], b) five radiographs that together show the entire object.

1.1.2 3D computed tomography

A 3D CT dataset is obtained from collecting radiographs in a full angular range around the object. These radiographs are then combined into a 3D reconstruction image. This digital representation of the object can be sliced open to obtain cross-sections in any direction. A CT reconstruction provides depth information and pulls apart the different features, which greatly enhances the knowledge that can be gained from the scan compared to the individual radiographs. An example of 3D CT imaging is shown in figure 1.2 for the same wooden statue as presented in figure 1.1. This shows the increased information gain of the 3D CT image compared to the 2D radiographs. The radiograph does not contain sufficient information to extract a horizontal cross-section of the tree rings, which would be needed for dating the wood. A horizontal slice from the 3D CT reconstruction, however, provides a sharp image of the tree rings, enabling dating of the object [63].

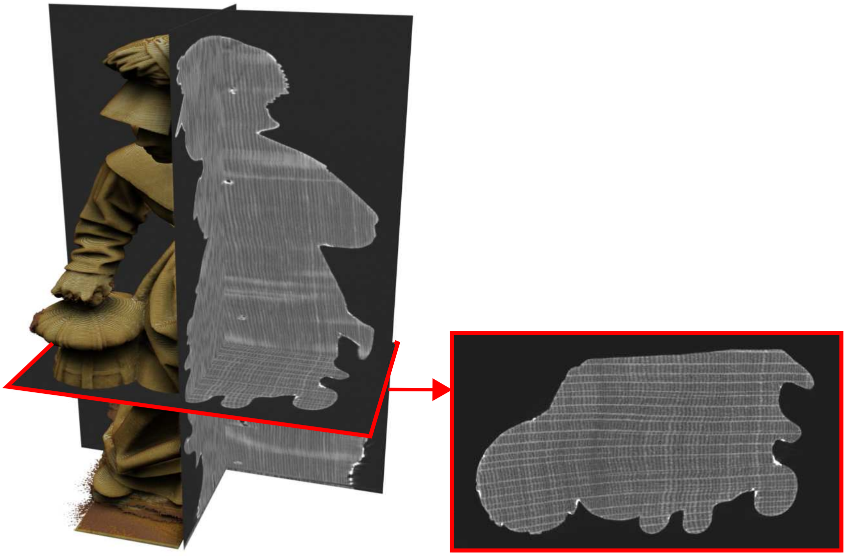

Figure 1.2: CT reconstruction of the *Woman with lantern*: 3D visualisation and orthogonal slices, horizontal slice showing the tree ring pattern.

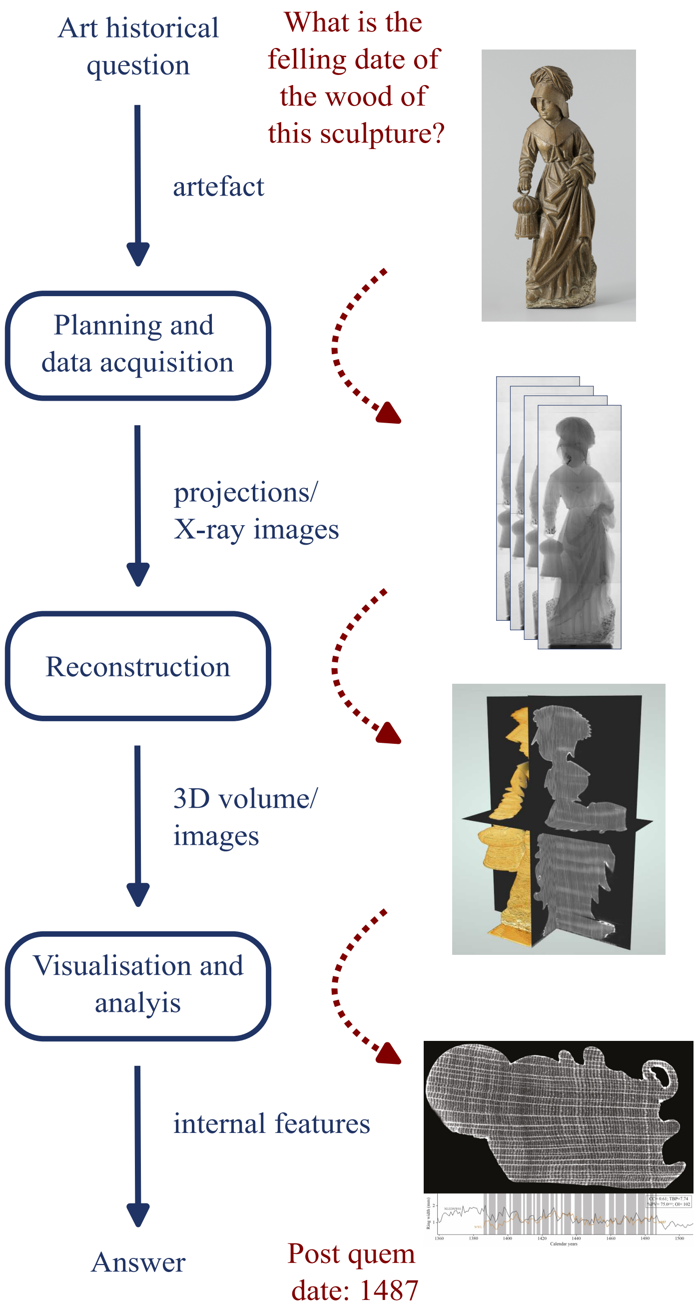

Figure 1.3: CT scanning research process from question to answer (in blue), with an example detailing the process of scanning the *Woman with lantern* (in red). A *post quem date* is an earliest date for the felling of the tree, such that the object must have been made after that date.

1.2 Computed tomography workflow

Since the invention of CT by Sir Godfrey Hounsfield in 1967 [151], this technique has been applied to many applications outside the original application in the medical domain [102]. Over the last decades, it has also been increasingly applied to the investigation of art objects. An important challenge for the general implementation of CT scanning for cultural heritage objects is that there is a wide variety of art objects, with different sizes, shapes and materials. Therefore, there is not one standard way to CT scan an art object. In contrast, for medical CT scans the position and general shape of the interior is known beforehand and the densities of the subjects are similar [16, 95]. The state-of-the-art medical scanners have been optimized with respect to these characteristics and standard acquisition protocols have been developed, which are repeatable for different patients. In the case of cultural heritage studies, it is unknown beforehand which features are present in the object and what scan specifications are needed to image them.

When CT scanning cultural heritage objects, the scanning process needs to be adapted to the object and the questions that are related to it. The research process is shown in figure 1.3. It starts from an object based question (or questions) by an object expert (e.g. 'what is the felling date of the wood of this sculpture?'). During the data acquisition phase X-ray images are collected from a range of angles, possibly in multiple scans if the object is larger than the detector. This series of X-ray images is then the input for a reconstruction algorithm, which produces a 3D image of the interior of the object. This reconstruction is visualised and inspected to extract features (e.g. tree rings) that can answer the original research question.

The main phases in the CT research process for cultural heritage objects are shown in figure 1.3: planning and data acquisition, reconstruction, and visualisation and analysis. On the right side these phases are illustrated by the steps in the scanning process of the *Woman with lantern*, from the Rijksmuseum collection [159]. Below we will discuss each of the research process phases in detail and outline the limitations, specifically as they apply to the scanning of cultural heritage objects.

1.2.1 Planning and data acquisition

The first step towards resolving questions using CT imaging is the data acquisition. Below we discuss the components of the CT system and how X-ray images are formed.

The CT scanner

X-ray imaging setups for scanning static objects typically consist of an X-ray source, a detector and a rotation stage in between, on which the object is mounted (see figure 1.4). For our workflow, we use a cone-beam X-ray source and a digital flat-panel detector. In dedicated CT systems, the location of each of these components are accurately known and the rotation stage can be controlled to obtain radiographs at precise angular steps. The distance from the source to the rotation axis is called the *source object distance* (SOD) and the distance from rotation axis to the detector

1

object detector distance (ODD). In order to be able to CT scan an object, the object needs to fit within the confines of the CT scanner and fully rotate. It also needs to be stable and not move in between X-ray images. When the object is wider or taller than the field of view of the detector, it is necessary to make *tiled* scans: multiple CT scans with the source and detector at different locations. The choice for the number of vertical and horizontal tiles and the tiling method, either moving detector only or both source and detector, depends mostly on the shape of the object. Additionally, the tiling method influences the reconstruction step, which is further described in section 1.2.2.

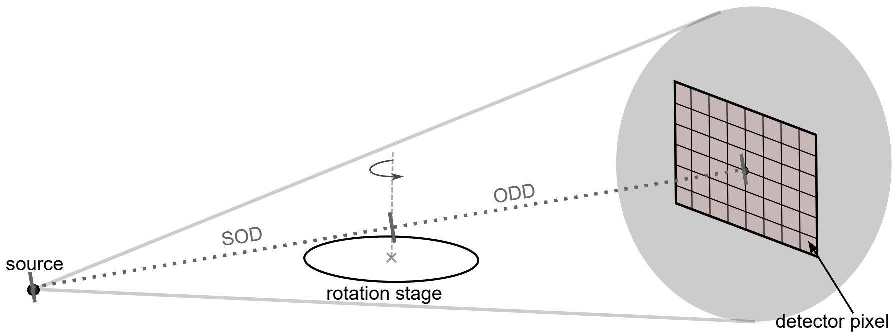

Figure 1.4: Schematic representation of a standard CT setup consisting of an X-ray source, a detector and a rotation stage in between.

X-ray absorption and projections

X-ray images are also called *radiographs* or *projections* and are formed by taking measurements of the intensity of the X-rays after they are attenuated by the object. This results in a projection of the object's internal structure. The X-ray absorption is material dependent [48, 50, 95]. The intensity $I(l)$ measured at a point on the detector depends on the length and material of the object, following Lambert-Beer's Law:

$$I(l) = I_0 e^{\int_l -\mu(x)dx}, \tag{1.1}$$

where $\mu(x)$ is the absorption coefficient of the material, l the line from the source to that point on the detector and I_0 the intensity of the X-rays when they leave the source.

Equation 1.1 can be rewritten to obtain the linearized photon count along l:

$$p(l) = -log(\frac{I(l)}{I_0}) = \int_l \mu(x)dx. \tag{1.2}$$

For a schematic illustration, see figure 1.5.

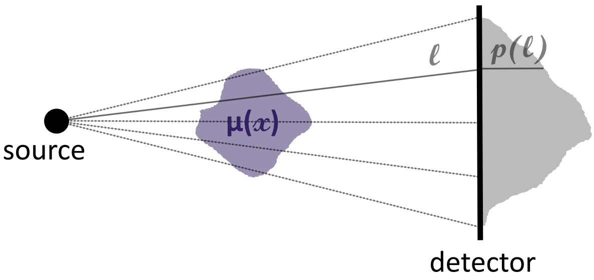

Figure 1.5: Schematic representation of X-ray absorption along a line l.

The values of the projection measurement depend on a number of parameters: source intensity, exposure time and the material composition and location of internal features within the object. What is seen on the projection image further depends on the precise location of the source, object and detector, as well as the orientation of the object and the acquisition angles. In figure 1.6 this is illustrated with two radiographs of a 17th century cutlery case taken at two different angles. Due to the cone shape of the beam, the object is moreover magnified on the projection. The magnification factor m is determined by $m = \frac{SOD+ODD}{SOD}$. The higher the magnification, the smaller the features that can be seen in the reconstruction. In order to obtain sufficiently high contrast in the features, the source settings (energy and power) need to be tailored to the object [106].

Planning and data collection

When it has been decided to perform a CT scan of an object, the first consideration is how to mount it on the rotation stage. Mounting the object in a fixed and stable manner is important to obtain accurate results and ensuring the safety of the object. This is often done by designing a foam holder specific for the object. The foam is almost transparent on the X-ray image, while it allows to stabilise the object and keep it in a position that is easiest for the scanning process. With the object in the scanner, the CT scans are designed. First, it is determined how many scans are necessary to obtain data of the entire object. Second, the source settings (energy, power) and exposure time that give the highest contrast in the X-ray image are investigated [106]. A CT dataset is then recorded, which consists of a set of projections, typically hundreds to thousands acquired across a full rotational range. Additionally, a darkfield (a projection with the source turned off) and a flatfield (a projection with the source on and no object in between) are collected to compensate for defects in the detector and background radiation that produces noise. In a metadata file the source settings, exposure time and the locations of the source, detector and rotation stage are recorded, as well as the rotation angles. Depending on the settings, the time for a single scan varies from minutes to hours. This needs to be taken into account when planning the scans of objects that have to return to the museum within a certain amount of time, usually the same day.

1

Figure 1.6: a) Cutlery case from the Rijksmuseum collection [158], b) and c) two radiographs of the lid of the cutlery case, with c) rotated ninety degrees with respect to b). Image adapted from [106].

1.2.2 Reconstruction

After data acquisition, a CT reconstruction algorithm computes a 3D volumetric image of the scanned object based on the acquired radiographs and the metadata. In this section we will first introduce the underlying inverse problem and then discuss reconstruction algorithms and tiled reconstructions.

The inverse problem

Tomographic reconstruction is an inverse problem, in which an image of the original object is obtained from a series of projections of that object. When scanning an object, we obtain a finite number of measurements, namely one intensity measurement per detector pixel per projection. This is the projection data $\mathbf{b} \in \mathbb{R}^{N_\alpha \times N_p}$ with N_α the number of projections and N_p the number of detector pixels. The vector $\mathbf{x} \in \mathbb{R}^{N_v^3}$ defines the cubic reconstruction volume, with N_v the number of voxels in each direction. In most cases, the object is assumed to be contained in this reconstruction volume. The resolution or *voxel size* of the reconstruction depends on the magnification m and the detector pixel size δ, and is given by $\frac{\delta}{m}$.

Due to the fact that we have a finite number of measurements, the forward projection can be formulated as a system of linear equations. The forward operator $\mathbf{A}$ depends on the acquisition trajectory and thus on the locations of the source, detector and object and the rotation angles. Each component of the matrix $\mathbf{A}_{ij}$ corresponds to the absorption of object voxel j of the ray that leads to measurement i. The vector $\mathbf{x}$ is the digital representation of the object that leads to the acquired projection data $\mathbf{b}$:

$$\mathbf{A}\mathbf{x} = \mathbf{b}. \tag{1.3}$$

Reconstruction algorithms

The goal of reconstruction algorithms is to find a vector $\mathbf{x}$ which leads to the acquired projection data $\mathbf{b}$. We distinguish between two types of reconstruction algorithms: filtered backprojection methods and iterative methods.

Filtered backprojection type algorithms have been developed to obtain fast reconstructions for specific acquisitions geometries. The reconstruction is obtained by first applying a well-chosen filter $\mathbf{h}$ to the acquired data and then backprojecting it into the reconstruction volume:

$$\mathbf{x} = \mathbf{A}^T(\mathbf{h} * \mathbf{b}). \tag{1.4}$$

A standard method for reconstruction of a circular cone beam geometry is the Feldkamp-Davis-Kress (FDK) algorithm [72]. Within the reconstruction algorithm, the acquisition trajectory is required to be circular and the angles equidistant. The flexibility of these methods is therefore limited.

Another approach to finding a solution to the inverse problem, are iterative reconstruction methods. These methods aim to minimize the difference between the forward projected image representation and the data, for example:

$$\min_{\mathbf{x}} |\mathbf{A}\mathbf{x} - \mathbf{b}|^2. \tag{1.5}$$

One widely used iterative reconstruction method is the Simultaneous Iterative Reconstruction Technique (SIRT) [87]. The SIRT algorithm operates by performing a gradient descent to minimize the residual, which is determined by forward-projecting the current estimate of the object representation and comparing it to the data. In each iteration, the current estimate $\mathbf{x}^i$ is updated to obtain the next estimate $\mathbf{x}^{i+1}$ as follows:

$$\mathbf{x}^{i+1} = \mathbf{x}^i + \mathbf{C}\mathbf{A}^T\mathbf{R}(\mathbf{b} - \mathbf{A}\mathbf{x})^i, \tag{1.6}$$

where $\mathbf{C}$ is a diagonal matrix containing the inverse of the sum of the columns: $c_{jj} = \frac{1}{\sum_i a_{ij}}$ and $\mathbf{R}$ for the rows: $r_{ii} = \frac{1}{\sum_j a_{ij}}$, compensating for the number of rays that hit each pixel and the number of pixels that are hit by each ray.

The SIRT method is more flexible than the before mentioned FDK algorithm, since it can be used with different acquisition geometries. A drawback is that it is relatively slow and computationally intensive.

Tiled scan reconstructions

There are a few methods to obtain tiled scans and reconstructions for objects that do not fit within the detector frame. Below we address the different methods and outline the advantages and challenges of the reconstruction. For each of these methods it is important to have sufficient overlap between neighboring radiographs. Which of these methods are possible depends on flexibility of the hardware and software that controls the scanner components.

- **Vertical tiling: moving detector only**
  If only the detector is moved, the images are sampled from the same cone beam. Therefore, the radiographs can be stitched together for each projection angle to form one large radiograph, from which the reconstruction can be computed as if it were a single scan. This makes it possible to use any reconstruction algorithm. For this method, the height is however limited to the height of the cone beam at the detector distance, since if we move the detector too far up or down, the beam will no longer fall on it and part of the detector will be in

shadow. Moreover, on the edge of the beam, it may be possible that cone angle artefacts appear. These artefacts are distortions in the reconstructed image that are caused by the larger angle between the source and detector.

- **Vertical tiling: moving detector and source**
  When moving the detector and source together vertically, the images are sampled from different cones. Because the rotation axis is projected onto the center of the detector, the region of interest (ROI) section of the object is captured in all the projections and therefore the different scans can be reconstructed independently and merged afterwards.

- **Horizontal tiling: moving detector only**
  In general horizontal tiling is more challenging than vertical tiling, because the rotation axis is no longer projected onto the detector for all tiles. Similar to vertical tiling, if only the detector is moved, the radiographs can be merged before reconstruction and any reconstruction algorithm can be employed. Again, a drawback is the limited movement range due to the size of the cone beam.

- **Horizontal tiling: moving detector and source**
  By moving detector and source together, the drawback of the previous method is avoided. Now, it is however not possible to merge the radiographs before reconstruction, because they are sampled from different cones. A reconstruction can be obtained by using an interactive algorithm, such as the SIRT algorithm described above, providing the different datasets and metadata to the algorithm all at once.

Which of the tiling methods is chosen is highly object and question dependent. The size and shape and how the object is positioned can influence the tiling method and number of tiles. The tiling method moreover influences the image quality and obtained resolution. For a CT scan of a painting from the Rijksmuseum collection *Cadmus sowing dragons' teeth*, we tested two options for mounting the object and tiling the scans on a mockup plank (see figure 1.7). Since the aim of the scan was to obtain a cross section of the tree rings, only a ROI along the width of the painting was needed. The horizontal tiling needed significantly fewer tiles. The vertical tiling needed more tiles and material outside the ROI moved through the field of view, but the final image quality was higher than the horizontal tiling. The vertical tiling could also be reconstructed faster and more easily. Due to the very small width of the tree rings, a high resolution and image quality were needed and thus the vertical tiling mode was used for the final scan of the panel painting.

Another example is the *Woman with lantern*, which was scanned with five vertical scans moving both source and detector, see figure 1.1. The FDK reconstruction method was applied to each of the scans, with were afterwards merged using the Flexbox toolbox [108].

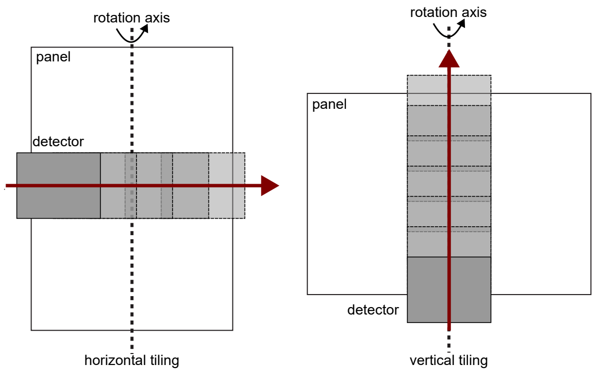

Figure 1.7: Schematic indicating the two tiling modes that were investigated for the scan of a panel painting, the panel is outlined in white and the detector in grey, with lighther grey indicating subsequent detector positions. Image adapted from [60].

1.2.3 Visualisation and analysis

When the reconstructions have been obtained, these need to be inspected to obtain the answer to the original question. Below we first discuss the visualisation options and then the search for internal features that answer the questions about the object.

2D slices and 3D representation

There are multiple options for investigating and visualising CT reconstruction data. Within the medical world, specialised software tailored to the scanner and purpose, is regularly supplied by the manufacturer of the hardware. Within laboratory settings, the use of free and open source software ImageJ [166] is widespread. While it is easy to view the data as a sequence of slices in ImageJ, its 3D capabilities are limited. It can be challenging to look at 2D cross sections of a 3D object as the relationship between what can be seen on the outside and the internal features is not always clear. A 3D visualisation can assist in making these connections clear.

For 3D visualisations of CT data there are several options, amongst which Slicer3D [71] is an open source solution that is used for both medical and academic purposes. There are also commercial options available, such as the versatile 3D visualisation programs VGStudio [201] and Avizo/Amira [192] but their cost limits their availability and suitability in cultural heritage institutions lacking sufficient funding. Another non-commercial option for volume rendering is Dragonfly [66], which offers a non-commercial licence for academic research.

Internal features

The last step in the CT research workflow is to look for the internal features that can answer the art historical question that started the workflow. These features are very diverse and specific to the object and the question. They can for example be glue lines, cracks, tree rings or layers and patches of different materials.

In the case of the *Woman with lantern* for example the features of interest were the tree rings. The radiographs in figure 1.1 do not give sufficient information, since for tree ring measurements a cross section through the wood is needed. The CT reconstruction does provide this cross section. The horizontal slice shown in figure 1.2 can be used for tree ring measurements to obtain a date for the when the tree was cut (see figure 1.8).

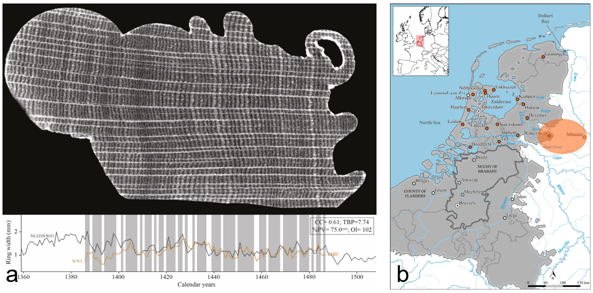

Figure 1.8: A slice from the CT scan from which the tree ring series could be measured (a) to obtain a *post quem date* (an earliest date for the felling of the tree) of 1487 and a provenance region (b). Image adapted from [63].

1.3 Computed tomography in cultural heritage

1.3.1 Applications in cultural heritage

Although CT imaging was originally developed for medical purposes, it has been applied to cultural heritage as well over the past decades. The investigation of interior features of an art object can give valuable information about the date and provenance of the object, the current conservation state and the making process.

Dating and provenancing objects

CT can be used for the dating and provenancing of objects. In particular it is well suited for the investigation of tree ring patterns for dendrochronological measurements, as it provides access to the inner structure of the wood when the tree-ring patterns cannot be retrieved by direct inspection on the surface [19, 58, 135, 182].

The *Woman with lantern* is an example of successful dating and provenancing based on CT imaging. A horizontal cross section from the CT scan such as the one shown in figure 1.2c, could be used for tree ring measurements (see figure 1.8a) which led to a post quem date (an earliest date for the felling of the tree) and moreover a provenance for the wood (see figure 1.8b). This provenance could be used to draw conclusions about historical Dutch timber trade.

Assessing the current conservation state

CT imaging can yield additional insights into the conservation state of an object and previous restoration interventions [47, 134, 149, 150]. It can be used to assess the current condition of the object for restoration and conservation purposes, revealing information on internal damage such as cracks, gaps and filler material [140, 178].

In the case of a 17th century cornett (figure 1.9a), a single X-ray image shows that the upper half is much less dense than the lower half, indicating that the top half is more damaged by woodworm (figure 1.9b). In the ROI CT scan (figure 1.9c and d), however, the exact structure and severity of the tunneling could be investigated.

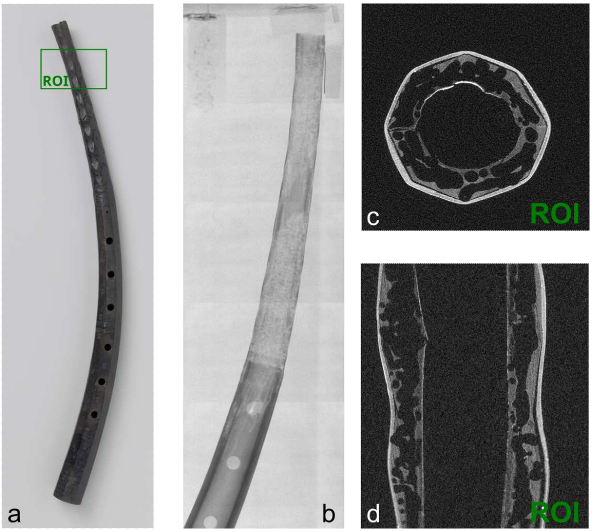

Figure 1.9: a) Cornett [157] with ROI region indicated, b) X-ray image of the top half of the cornett, c) horizontal and d) vertical slices through the CT reconstruction of the ROI.

Determining the making process

An important application of CT imaging for cultural heritage is the determination of the making process, since this can shed light on the creative process of the artist. For example, searching for lines in a terracotta statue can shed light on whether the clay was press molded or freely worked by hand [81] and scanning glass beads led to a better understanding of the production process and tools used to create them [132].

The example shown in figure 1.10 is the plaster sculpture *Python killing a Gnu* from the J. Paul Getty museum collection. The original object was smaller and more condensed. Later, the object was reconfigured to its current shape. In the horizontal slice through the base, clear evidence of the reconfiguration can be seen (figure 1.10c) as the original square base is enveloped in the current, more rocky base.

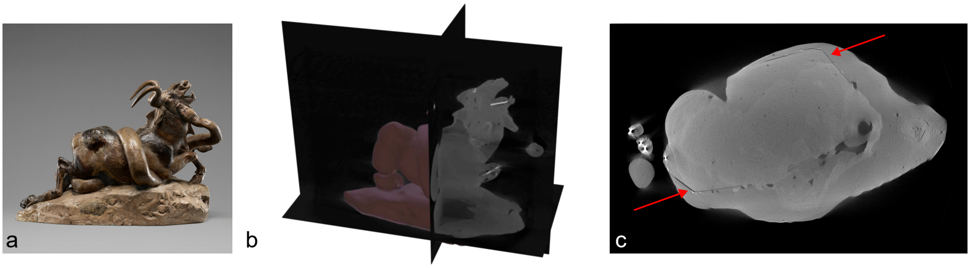

Figure 1.10: a) *Python killing a Gnu* [190], b) 3D CT reconstruction volume representation with orthogonal slices and c) horizontal slice showing the original square base enclosed in the current base.

1.3.2 Challenges for the application of CT to cultural heritage

Despite the known advantages of CT as a technique to obtain information about the interior of objects, it has not been widely implemented in museum research facilities. There are a couple of reasons why this technique has been difficult to apply to museum collection objects at large, which we will discuss in this section.

Shapes and sizes of objects

The variety in shapes and sizes of cultural heritage objects poses a challenge due to the often confined space in the scanner. First, they need to fit and be able to fully rotate within the scanner. Second, when the object is longer in one direction than another, it is difficult to find source settings to have sufficient contrast over the full rotational range. This is for example the case for paintings, which are long in one direction and thin in the other. In figure 1.11 the wide range of shapes, sizes and materials of objects that were investigated in the course of the research for this thesis is shown.

Figure 1.11: Overview of the objects scanned in the course of the research for this thesis and their sizes:
a) Green velvet purse with gold thread, after 1580, Rijksmuseum collection nr. BK-KOG-29 [153],
b) Jaguar figure, wood covered with mosaic, 1400-1521, British Museum, collection nr. Am,+.165,[187], ©The Trustees of the British Museum.
c) Cutlery case, purple velvet embroidered with pearls and gold thread, c. 1600-1625, Rijksmuseum collection nr. BK-NM-3086, [158],
d) Cornett, boxwood and leather, c. 1600-1650, Rijksmuseum collection nr. BK-AM-62-B [157],
e) Mummy mask, cartonnage, 1stC BC-1stC, British Museum, collection nr. EA 29472 [188], ©The Trustees of the British Museum,
f) Woman with lantern, oakwood, c. 1500-1550, Rijksmuseum collection nr. BK-NM-9253 [159],
g) Cadmus, Guided by Minerva, Observes the Spartoi Fighting, Peter Paul Rubens (after), oil on paper laid down on panel, before 1747, Rijksmuseum collection nr. SK-A-4051 [156],
h) Bottle in the shape of a shoe, leather, brass and wood, c. 1675-1700, Rijksmuseum collection nr. BK-KOG-1382 [155],
i) Book chest of Hugo de Groot, wood, leather and metal, c. 1600-1615, Rijksmuseum collection nr. NG-KOG-1208 [154],
j) Python Killing a Gnu, Antoine-Louis Barye, plaster, wax and metal, c. 1840-1860, The J. Paul Getty collection nr. 85.SE.48 [190].

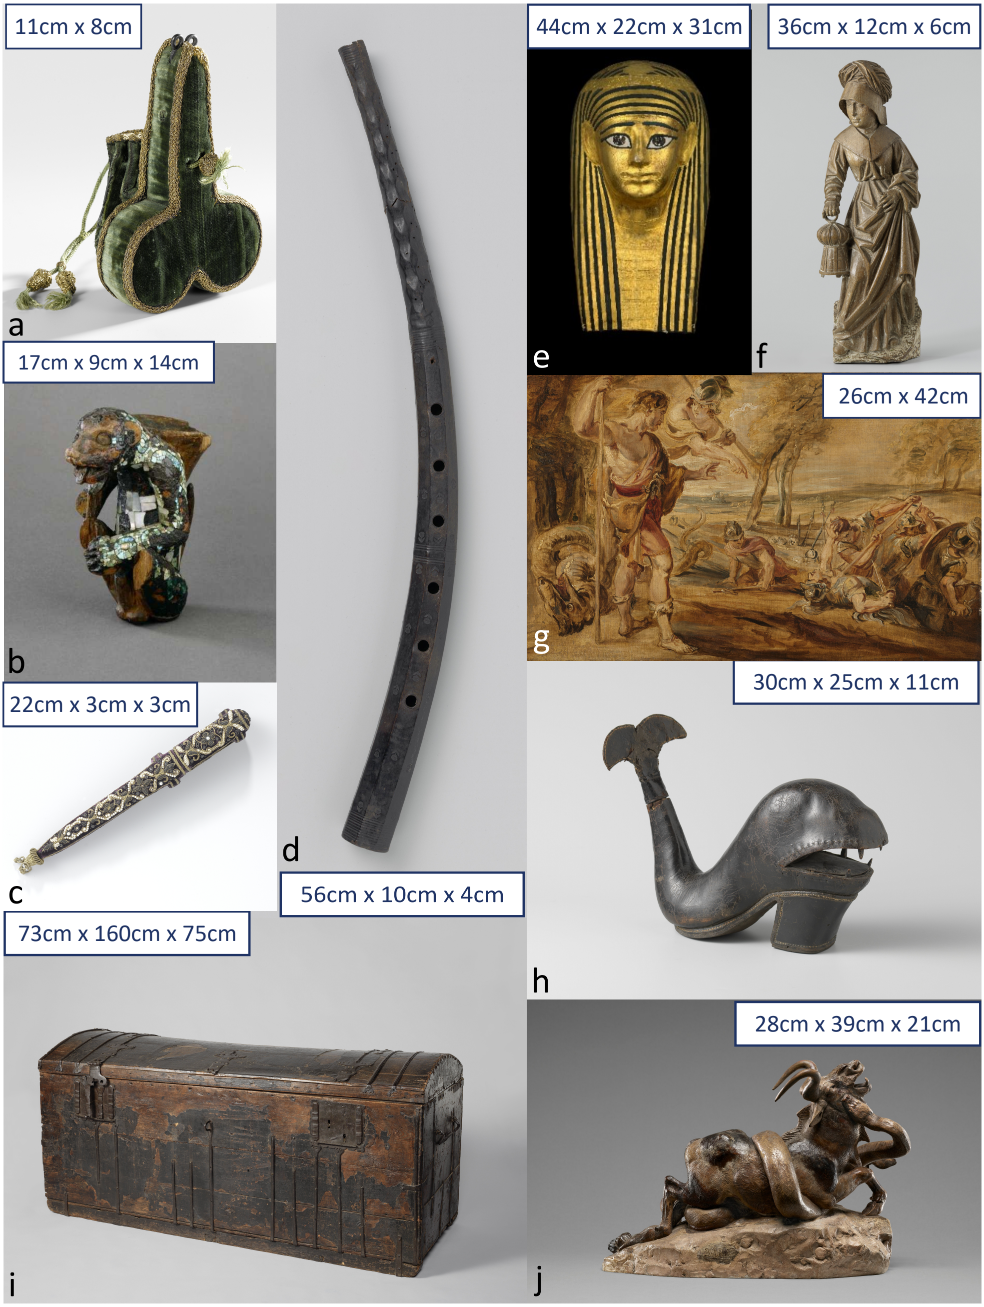
11cm x 8cm
a
17cm x 9cm x 14cm
b
22cm x 3cm x 3cm
c
d
56cm x 10cm x 4cm
44cm x 22cm x 31cm
e
36cm x 12cm x 6cm
f
26cm x 42cm
g
30cm x 25cm x 11cm
h
73cm x 160cm x 75cm
i
28cm x 39cm x 21cm
j

1

Variety of materials

Cultural heritage objects are made of a large variety of different materials, making the choice of acquisition settings a challenge, especially when the object contains both high and low density materials. This is most pronounced when one of the materials is a metal. This is for example the case for the cutlery case (figure 1.12a). See figure 1.12b for an example, where the gold thread surrounding the cutlery case causes streaking artefacts in the reconstruction. By carefully choosing the source settings and beam filtration, these artefacts can be reduced (figure 1.12c).

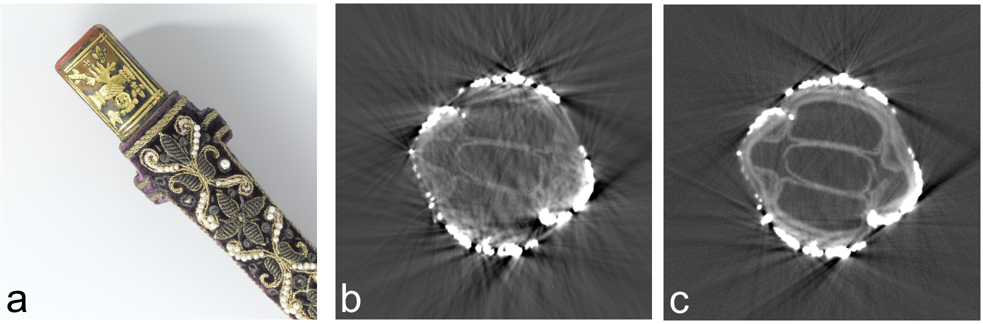

Figure 1.12: a) Detail of the cutlery case [158], b) and c) horizontal reconstruction slice of the widest section of the cutlery case with b) an untailored acquisition scheme, c) an acquisition scheme for which the source settings and beam filtration have been carefully chosen. Image adapted from [106].

Accessibility of scanning facilities

Most museums do not have CT scanning facilites in-house, making it necessary to move precious collection items to external scanning facilities. CT scans carried out at commercial scanning facilities can moreover be expensive. In this thesis, multiple facilities have been used, each with their own specific characteristics with regard to hardware and software. The most-often used facilities are the FleX-ray lab, a laboratory micro-CT scanner situated at the Centrum Wiskunde & Informatica in Amsterdam (figure 1.13a), and the in-house X-ray facility of the Rijksmuseum (figure 1.13b). In this thesis, we also use datasets recorded at the X-ray facilities in the British Museum (London) and the J. Paul Getty Museum (Los Angeles).

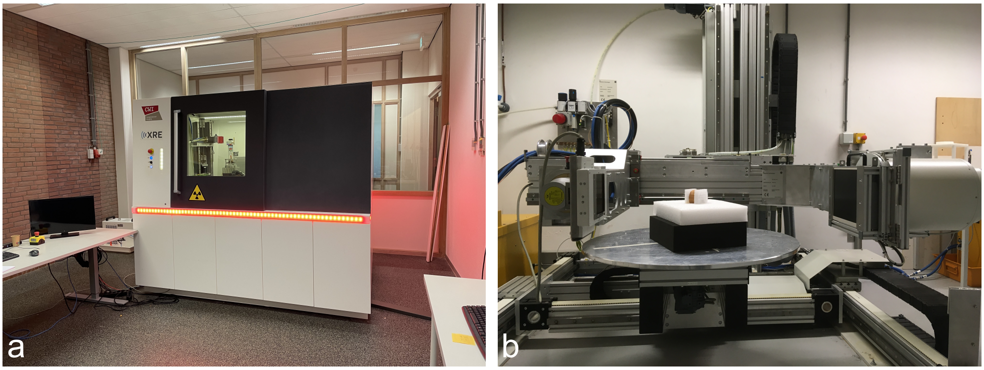

Figure 1.13: a) The Flex-ray lab at CWI, b) the in-house X-ray facility at the Rijksmuseum.

Data interpretation

The usual way of displaying the 3D reconstruction as 2D slices hampers the interpretability of the scans, because it can be a challenge to relate the features in a 2D cross-section to the three dimensional object. A 3D visualisation can assist in overcoming this issue, but a challenge for the interpretation of such 3D visualisations is that these do not reflect the actual colors and textures of the object. This is especially challenging for cultural heritage experts, who are trained to look at the original objects closely and move around and handle them to inspect the exterior details and not at a greyscale digital representation.

1.4 Research questions

In this thesis, we address the challenges outlined above by answering the research questions outlined in this section. Below, we indicate how the chapters in this thesis deal with the research questions and how the newly developed methods were applied to case studies from the cultural heritage domain.

Research question 1.
How can expert knowledge be integrated in the CT acquisition process to facilitate efficient answering of heritage questions?

In chapter 2 we discuss how active involvement of object experts can guide the acquisition process to efficiently answer questions about the objects. In many CT scanning facilities a CT scan is performed and the data is given to the object expert after the object has left the scanning facility. This means that the process needs to be repeated if the answer cannot be found by analysing the dataset. We argue that it is important to actively involve the object expert in the scanning process and adapt the scanning approach based on intermediate feedback. Not only can we collect information to answer the original question, but the collaborative inspection of intermediate results can lead to additional information because new features of interest can be identified and captured.

In chapter 2 we demonstrate the process by detailing the approach taken to scan a 17th century cornett (see figure 1.9) from the Rijksmuseum collection with active participation of the object experts and imaging experts. The scanning process was steered to obtain ROI scans answering questions that were inspired by intermediate results (see figure 1.14).

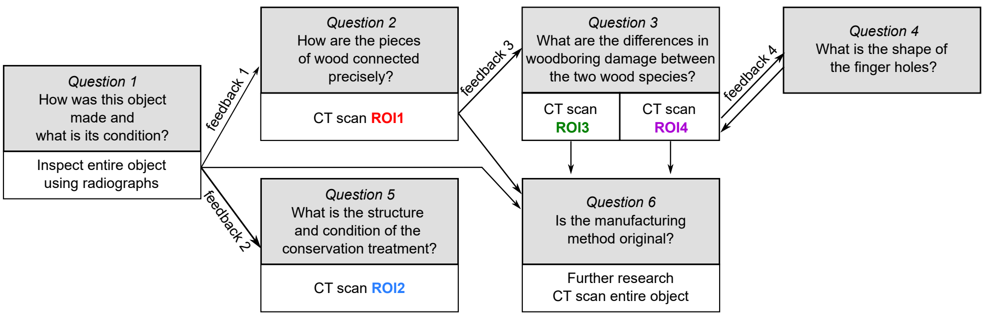

Figure 1.14: Example schematic of the workflow employed during the investigation of a 17th century cornett, indicating the steering of the scanning process by expert feedback.

Research question 2.
When a full CT is not feasible due to the size or shape of the object, can we design object tailored acquisition schemes to obtain the specific information that answers the art historical research question?

In chapter 3 we explore the possibility of obtaining specific information from a series of radiographs, when a full CT acquisition is not possible due to the large size of the object. When the object fits in the confines of the CT scanner, usually a full CT acquisition by taking images while rotating the object 360° is obtained. Using reconstruction algorithms, this data is then combined into a 3D reconstruction of the object. This approach however excludes objects that are too large to fit in the CT scanner. However, to answer some questions, no full 3D reconstruction is required. One of these cases is dendrochronological research, in which tree ring series are measured and compared to reference chronologies to determine the age and provenance of the wood. For this purpose a cross-section of the tree rings is needed. We developed an acquisition and reconstruction approach that uses X-ray data by moving an object on a linear trajectory between the source and detector to obtain clear images of the tree rings (see figure 1.15).

The approach for obtaining images of tree rings is first demonstrated on test planks in chapter 3. For these, the tree rings were visible on the outside and therefore the obtained tree ring measurements from the reconstructed X-ray images could be compared to the traditional method of measuring on photographs. Consequently, this approach is applied to the Hugo de Groot bookchest from the Rijksmuseum collection [154].

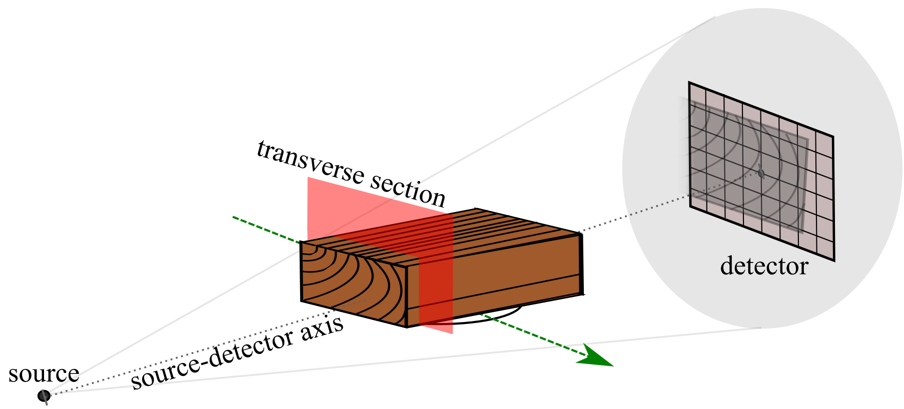

Figure 1.15: Line trajectory scanning.

1

Research question 3.
How can we optimise the use of the in-house X-ray facilities for live radiography by developing methods to use these setups for 3D imaging of museum collection objects?

In chapter 4, we propose a novel approach for creating accurate CT reconstructions using only standard radiography equipment already available in most larger museums. Specifically, we demonstrate that a combination of basic X-ray imaging equipment, tailored marker-based image acquisition protocols, and tailored data-processing algorithms can achieve 3D imaging of collection objects without the need for a costly CT imaging system. Our work paves the way for adoption of CT technology across museums worldwide.

We implemented our marker-based approach in the British Museum (London), the J. Paul Getty Museum (Los Angeles), and the Rijksmuseum (Amsterdam). In chapter 4, we demonstrate the method by scanning a small wooden test object in all three scanning facilities to compare the resulting 3D reconstructions. We moreover applied the method to a case study from the J. Paul Getty museum: *Python killing a Gnu* by Antoine-Louis Barye (French, 1796 - 1875), see figure 1.16. The CT reconstruction provided considerably more insight into the making process of this interesting sculpture than previously acquired radiographs. It showed the different layers of plaster and the parts where the object had been remodelled, as well as where metal rods are keeping the different pieces together.

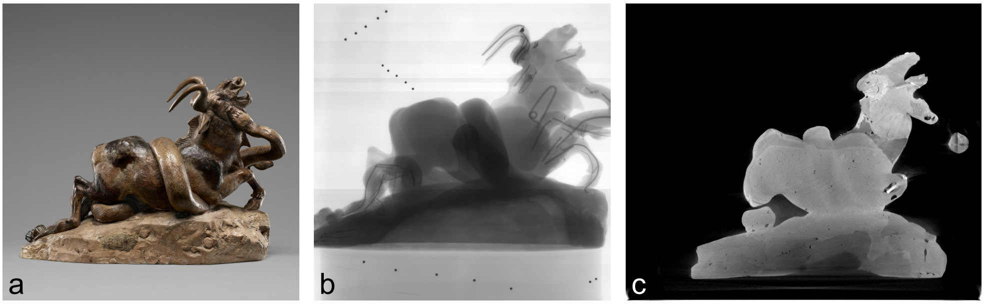

Figure 1.16: a) *Python killing a Gnu* [190], b) radiograph containing markers, c) cross section of the CT reconstruction.

Research question 4.
How can the data be visualised to improve the interpretability of CT scan data for conservators, restorers and art historians?

One of the challenges for the broad implementation and adoption CT imaging for cultural heritage research is the interpretation of the data. The 3D data is usually inspected as 2D slices or cross-sections through the object. It can be difficult to relate what is seen on the outside of the object to the features shown by the 2D cross-sections, especially for museum professionals, who are trained to look closely at the objects and handle them. 3D surface scans have been employed for the visualisation and digitization of the exterior of 3D objects. These scans give the colours and textures of the object and the information gained is thus complementary to CT scans, which give a greyscale density-based image of the interior. Therefore, in chapter 5, we provide an interactive visualisation plugin for the open-source software Blender, to combine and inspect two complementary 3D imaging modalities: CT images, which capture the interior; and surface scans, which capture the exterior (see figure 1.17).

The plugin workflow was applied to four case studies from the collections of the Rijksmuseum, Amsterdam, and the British Museum, London. Each of the objects presented their own challenges for data analysis and visualisation. The simultaneous visualisation of the exterior and interior led to knew insights into the making process of these objects.

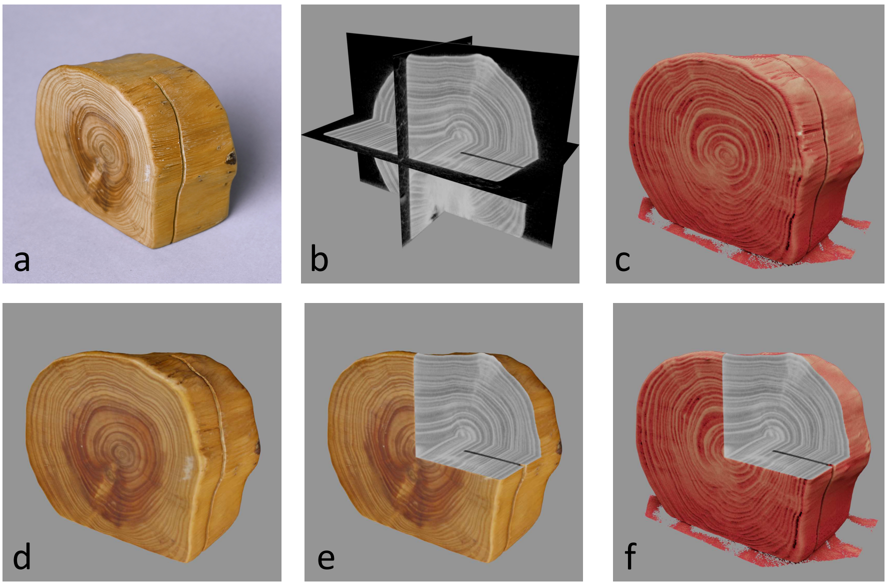

Figure 1.17: Combining CT scans and surface scans of a small wooden block: a) photograph of a small wooden block (h 5cm x w 6cm x d 3cm), b) X-ray CT scan represented as orthogonal slices, c) CT 3D volume render, d) surface scan, e) combined image modalities showing surface scan and slices, f) CT 3D volume render and slices.