ARCHES

H2020 – 693229

Deliverable D5.2 Test setup – 1st Version

Lead Author: Andreas Reichinger
With contributions from: Cornelia Travnicek, Maria Schimkowitsch, Moritz Neumüller, Helena Garcia-Carrizosa, Joanne Wood, Laura Luidolt, Daniela Stoll
Reviewers: Anne Fay

Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (P)
Contractual delivery date:	09-2017
Actual delivery date:	09-2017
Version:	1.0
Total number of pages:	77
Keywords:	Blind, visually impaired, hard of hearing, art, museums, inclusion

Abstract

This is a report about the current stage of development of the first Test Setup and the tools developed to reach this goal: The tactile relief generation software, the context-sensitive tactile audio guide, and the produced artworks.

We show the progress we made towards reaching our milestones within WP 5, especially subtasks 5.1.1 “Integration of prototypes”, 5.1.2 “Extension of design tools”, 5.3.1 “Improvement of touch and gesture recognition”, 5.3.2. “Improvement, extension of user interaction methods and content authoring tool”, 5.4.1 “Design and production of reliefs for evaluation” and 5.4.3 “Creation of tactile audio guide for produced reliefs”. We also present the promising results of the first user studies within the participative research group and sum up the conclusions we drew from it.

Document Information

IST Project Number	**H2020 - 693229**	**Acronym**	**ARCHES**
Full Title	Accessible Resources for Cultural Heritage EcoSystems		
Project URL	http://arches-project.eu		
Document URL	http://arches-project.eu		
EU Project Officer	Luis García Domínguez		

Deliverable	**Number**	D5.2	**Title**	Test setup – 1st version
Work Package	**Number**	WP5	**Title**	On-site multisensory activities
Date of Delivery	**Contractual**	M12	**Actual**	M12
Status	Version 1.0		final ☒	
Nature	Prototype ☐ Report ☒ Demonstrator ☐ Other ☐ Ethics ☐ ORDP ☐			
Dissemination level	Public ☒ Confidential ☐			

Authors (Partner)	VRVis			
Responsible Author	**Name**	Andreas Reichinger	**E-mail**	reichinger@vrvis.at
	Partner	VRVis	**Phone**	+43 1 908 98 92 604

Abstract (for dissemination)	This is a report about the current stage of development of the first Test Setup and the tools developed to reach this goal: The tactile relief generation software, the context-sensitive tactile audio guide, and the produced artworks. We show the progress we made towards reaching our milestones within WP 5, especially subtasks 5.1.1 “Integration of prototypes”, 5.1.2 “Extension of design tools”, 5.3.1 “Improvement of touch and gesture recognition”, 5.3.2. “Improvement, extension of user interaction methods & content authoring tool”, 5.4.1 “Design and production of reliefs for evaluation” and 5.4.3 “Creation of tactile audio guide for produced reliefs”. We also present the promising results of the first user studies within the participative research group and sum up the conclusions we drew from it.
Keywords	Blind, visually impaired, hard of hearing, art, museums, inclusion.

Version Log			
Issue Date	Rev. No.	Author	Change
05-04-2017	0.1	AR,MS,HGC,JW	Test Results IAG
01-09-2017	0.2	AR,CT,MS	Review Version
17-09-2017	0.3	Andreas Reichinger	Submitted to the coordinator
18-09-2017	0.3	Javier Gutiérrez Meana	Forwarded to the reviewer
26-09-2017	0.4	Javier Gutiérrez Meana	Format
27-09-2017	0.5	Anne Fay	Suggestions and minor changes
28-09-2017	0.6	Andreas Reichinger	Update
29-09-2017	1.0	Javier Gutiérrez Meana	Final version

Table of contents

List of figures

List of tables

Definitions

BVI: Abbreviation for "Blind and Visually Impaired", typically used as BVI people.

C#: Programming Language. It is a general-purpose and object-oriented language, developed my Microsoft as part of the .NET framework.

F#: Programming Language. It is primarily a functional programming language, but fully compatible with C#.

GUI: Graphical User Interface. Type of user interface that allows the user of electronic devices to interact with the machine via visual indicators, like Buttons, Icons, etc.

HoH: Abbreviation for "Hard of Hearing"

IAG: Interactive Audio Guide, as alternative to Gesture Controlled Tactile Audio Guide

JSON: A common open-standard file format used for transmitting data in a human-readable form

.NET: Collective term for several software platforms distributed by Microsoft, which are used for the development and execution of application software, frameworks, programming languages and so on.

Tactile Image: Umbrella term for all pictures, artwork and photographs represented in touchable physical form to make them more accessible

Tactile Relief: The actual physical form of a tactile image, which can be made from different materials

V&A: Victoria and Albert Museum, London

VRVis: Short form for Zentrum für **V**irtual **R**eality und **Vis**ualisierung Forschungs-GmbH (English: Research centre for **V**irtual **R**eality and **Vis**ualization)

WC: The Wallace Collection, London

WPF: Windows Presentation Foundation. Framework for creating graphical user interfaces in Windows-based applications.

1 Introduction

In this report, we document the progress of the WP5 "On-site multisensory activities" that was performed in the 9 months between Month 3 and the delivery date Month 12. The goal of this work package is to complement the screen-based work (desktop and mobile apps) of the other work packages (especially WP3 and WP4) with on-site sensory activities, especially including the tactile sense. The target setup is therefore beyond screen-based interfaces and focuses on an interactive haptic installation in the museum, with a *touchable relief interpretation* of one of the exhibits, together with a *context-sensitive tactile audio guide*, that reacts on touch and other gestures and offers additional information about the exhibit, as well as guidance for people with special needs. These ideas originated from work performed at VRVis since 2004 targeted especially at BVI people. In ARCHES we develop these concepts and technologies further, and investigate how these can be extended to serve a wider target audience.

As most of our proposed work is based on earlier prototypes created at VRVis, the work in the first phase was dedicated to thoroughly review which parts can be re-used and developed further, to incorporate new tools and techniques, and to build a stable foundation for further developments in the second phase.

This work package is divided into four parts, all of which will be reported on throughout this document and the simultaneously published deliverable D5.1 "Relief printer concept document".

In Task 5.1 "Tactile image generation", the goal is to write software tools that help an artist to convert two-dimensional images, like paintings or photos, into touchable reliefs. This task is very complex, as a lot of different techniques have to be combined, and thus requires a stable software basis into which these tools can be bedded. Therefore, a large portion of the effort of this first phase was put into designing this basis framework, in order to prepare it for the development of more advanced tools in the second phase.

Task 5.2 "Relief printer", is concerned with the development of a novel rapid prototyping approach that allows the fast and erasable physical printing of the designed height field data of a relief. As there is a simultaneously published dedicated deliverable D5.1, we refer to this document for our current progress.

Task 5.3 "Context-sensitive tactile audio guide" is concerned with the development of a touch- and gesture-interface directly on relief surfaces, based on optical finger tracking. This work originated from a very early prototype made in a previous project, with an ad-hoc custom-made setup that was very un-robust, took very long to set up, and was only usable with a single relief as code and data was mixed. In this first phase, we managed to port this solution to a very stable commercially available platform, the HP Sprout, strongly decreased setup time with an automatic calibration method, and opened the program for multiple reliefs, by developing a dedicated file format and an authoring tool, improved the finger-tracking algorithms, and could gain a speedup of up to a factor of 2-3. We already could perform hands on testing with participants in London, see Task 5.4. The feedback was already turned into improvements and extensions for the various needs, and was positively appreciated. We even got one step ahead, which is way beyond the ideas we wrote in the Grant Agreement, but believe it is worth the effort: Since the HP Sprout has an in-built projector, we managed to project images onto the reliefs! Together with the additional touch-screen this has the possibility to extend the "Context-sensitive tactile audio guide" into a "multimedia guide". The basis for this is already laid with the already implemented projection calibration and image warping functions, ready to be explored in the second phase.

Task 5.4 "Organisation of multisensory activities and testing" is concerned with the actual design and production of a test setup that will be installed at each participating museum to be tested and improved. According to the Grant Agreement, these tools include one tactile relief for each museum of one of their exhibits, and the corresponding tactile audio guide setup and authored content, as well as relief printer mediums to facilitate testing of these. This is however in contrast with developments and discussions of the consortium. We early agreed that the artworks of the museums to be made tactile are to be chosen by the participative research groups, as it is essential for the methodology to involve the members of the participative research groups in the whole process. Thus, only artworks for the Wallace Collection and the Victoria and Albert Museum have yet been chosen, as London is the only place where the participative research group is already formed and active. As agreed with the consortium, work on the artworks of the other four museums will only begin, once the participative research groups are formed and the selection process is complete. Further, we agreed, that we do not want to rush the design for the sake of this deliverable. In contrast to the purely software-based tools of the other work packages, physical tools, like our reliefs, can only be built once, and then hardly be updated; only a new tool can be built. Consequently, we decided to place a high emphasis on the design process until all parties are satisfied. We therefore report on the current state of the design, and postpone publication of the realised reliefs until the next deliverable.

In order to facilitate testing already before these reliefs are ready, we got the opportunity to get copies of earlier reliefs, as a test-case until the new reliefs for the museums are ready. As at this stage the new tactility and technology is in the foreground, the depicted content was not much of a concern. It may even have given a broader view to the participants, and allowed them to peek into foreign museums’ collections. Especially Gustav Klimt’s painting “The Kiss” interested several participants, as this painting is famous beyond its borders. It would even be interesting to share the produced reliefs among the participatory research groups to broaden the experiences, as, in contrast to the developed apps in WP3 and WP4, only a single art piece can be worked on per museum. We will look into this possibility, once all participative research groups are in place.

Similarly, since most tasks performed in this stage were made before the artworks for the reliefs were selected, throughout this work earlier reliefs have been used for testing and are therefore in the images.

Concerning other technologies, T-coil loops have already successfully been tested together with the tactile audio guide. The transmitter could simply be plugged into the 3.5mm headphone port of the HP Sprout, and was successfully paired with participant’s hearing aid, which made the spoken texts much clearer to understand.

Finally we want to report, that parts of a preliminary version of this report have already been compiled into two scientific publications, one already accepted and awaiting the final publication process, one still in review.

2 Task 5.1 "Tactile image generation"

In 2004, VRVis started to create tactile images. Originally those images were developed for BVI people as an aid to better perceive visual art, than by descriptions alone. It quickly turned out though, that many more people enjoy the tactile versions, not only BVI people.

Many tactile images are created on swelling paper or with Braille embossers, and are therefore very flat, abstract simplifications and are often not easy to interpret. On the other hand, some institutions, notably in Italy hired sculptors who created very detailed relief interpretations of paintings. Since then we experiment with various computer aided design tools, to digitally create reliefs with a similar quality. A number of prototypes and software tools were created over time. In ARCHES, we now strive to create a single program, that unites many of these tools, in order to allow a more efficient work environment for the artist.

Tactile reliefs, as we create them, are supposed to be true to the originals. We therefore take a high-quality scan of a painting, and then try to recreate the height at each pixel, so that a geometrically consistent, plastic, three-dimensional relief-version is created. The task is, therefore, to create a so-called depth-map from a painting that encodes the height for each pixel. Such a depth-map can be displayed, e.g., as a gray-scale image as can be seen in Figure 1c, and directly corresponds to its 3D representation which can be also directly visualised on the computer. In order to physically produce such a relief, it is either converted into 3D geometry files and sent to a 3D printer, or it can directly be converted into machine paths for CNC milling machines, that carve the relief out of a solid block of various materials (see Figure 1d).

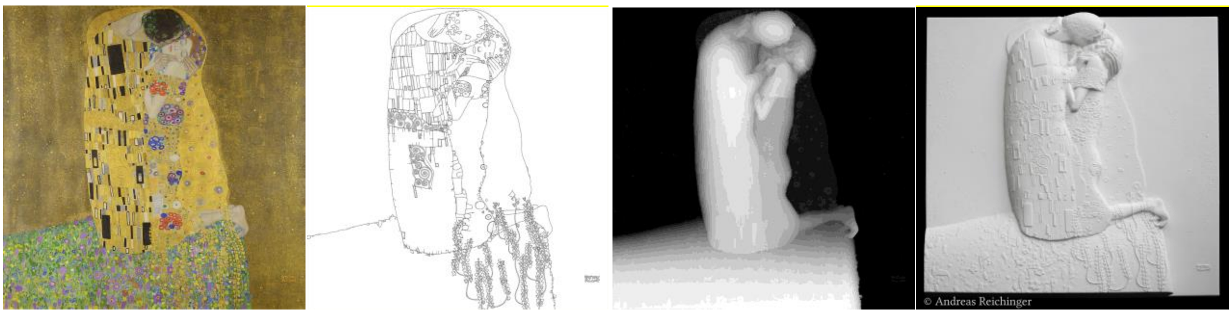

Figure 1: From painting to relief. From left to right: (a) Original scan of the painting; (b) Traced outlines to segment the painting into relevant parts; (c) Depth map (dark is low, bright is high); (d) Final relief.

The general idea of relief design is thus to take a 2D image as input and then generate a height field of said image by finding the appropriate height for each pixel in the picture. In most cases 1:1 mapping (e.g. brightness to depth) is not useful as the tactile relief needs to convey the content of the image, which means not a direct translation but rather a 3D re-interpretation of the depicted scene.

The design workflow we previously developed for this purpose can be explained best using one of the already existing reliefs as an example: "The Kiss" by Gustav Klimt [7].

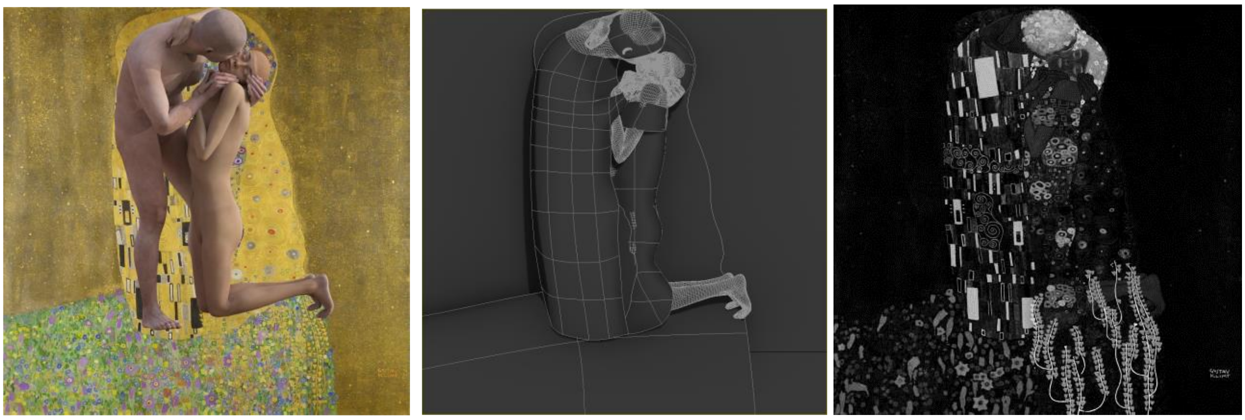

Figure 2: Stages of relief development. From left to right: (a) Study with digital mannequins to recreate the body poses; (b) Modelled meadow/cliff, and modelled clothes and hair over the figures using Bézier patches; (c) Final texture layer extracted from the original painting and manual enhancements.

In order to recover the body poses, we used digital mannequins in a 3D editing program, and iteratively adjusted those until they exactly resembled the pose in the painting (see Figure 2a). We then modelled the clothes over the naked bodies and the meadow hill under their feet using Bézier patches (see Figure 2b). To achieve an accurate resemblance, the individual parts of the 3D scene were rendered as depth map, then warped, cut and composited to exactly align with outlines (see Figure 1b) extracted from the painting.

After discussions with experts, minor revisions, and optimisation of the depth composition to maximise the available depth at important parts (e.g. face, feet, hands), we started to extract the texture, and added it as additional height variations over the base relief. In most parts we could extract a meaningful texture layer directly from a gray-scale version of the original painting (see Figure 2c) with various filters. Several important parts needed to be improved. These were segmented, either manually or colour-based, and included as corrections on the texture layer. These parts include the spirals on the male's coat, the different kinds of flowers on the meadow and in her hair, the wreath in his hair, the signature, and the tendrils on the right part of the meadow.

It is evident that the design process can be a tedious task that takes comparably long (about 2-6 weeks per art piece) and, in our previous workflow, required up to ten different kinds of software, which made it inherently difficult and time consuming, because of the various file formats and often required conversions between those, different user-interfaces and methodologies on which the use has to adapt for each change, and the need to go through this lengthy process before seeing whether the changes lead to the desired result.

2.1 Subtask 5.1.1 "Integration of prototypes"

In order to overcome the aforementioned limitations of the previous design process, the goal of this work is to develop a single program that potentially can incorporate all required functionality in a single consistent framework and user-interface, so that in the future the process of relief design takes less time, results in better quality relief data and can be offered as a competitive service. As this task is not trivial, a large part of time was reserved for the task of software development (as according to the proposal). In this first phase, we concentrated on building a stable software framework and extensible user-interface, to pave the way for the development of more powerful tools in the second phase. In the following sections, we describe the concepts we designed, and the current state of the software.

2.1.1 The software

The windows desktop software we develop is currently called DeepPictures, which should remind on the idea of giving images and paintings its depth as a relief. It is written in the .NET languages C# (an imperative language) and F# (a functional language) to support both programming philosophies, wherever one of the philosophies provides an advantage. It is based on the VRVis in-house developed Aardvark platform, which provides many libraries, especially in the fields of image processing and 3D graphics. The graphical user interface (GUI) is developed using the Windows Presentation Foundation (WPF).

An image can be imported, and then, based on the image, a corresponding tactile relief can be designed. The progress is stored as a .DPF file (DeepPictures File) which is a compressed JSON data stream. This was selected as it is human readable which could help prevent data-loss if file-format changes would make old files unreadable. Manual correction would then be possible.

At present, the software is only developed as an in-house tool, with no intention to publish it. This allows a faster development cycle, and enables us to quickly test various tools for their suitability. We do not rule out, that the software can be made public, but that will require further work and legal consultation or exchange of possibly used 3rd party libraries.

2.1.2 The graphical user interface

The user-interface (see Figure 3) and some functionalities of the developed GUI are loosely based on Adobe Photoshop since this was the central software used in earlier relief designs and has proven to be especially suitable for the requirements to the new software.

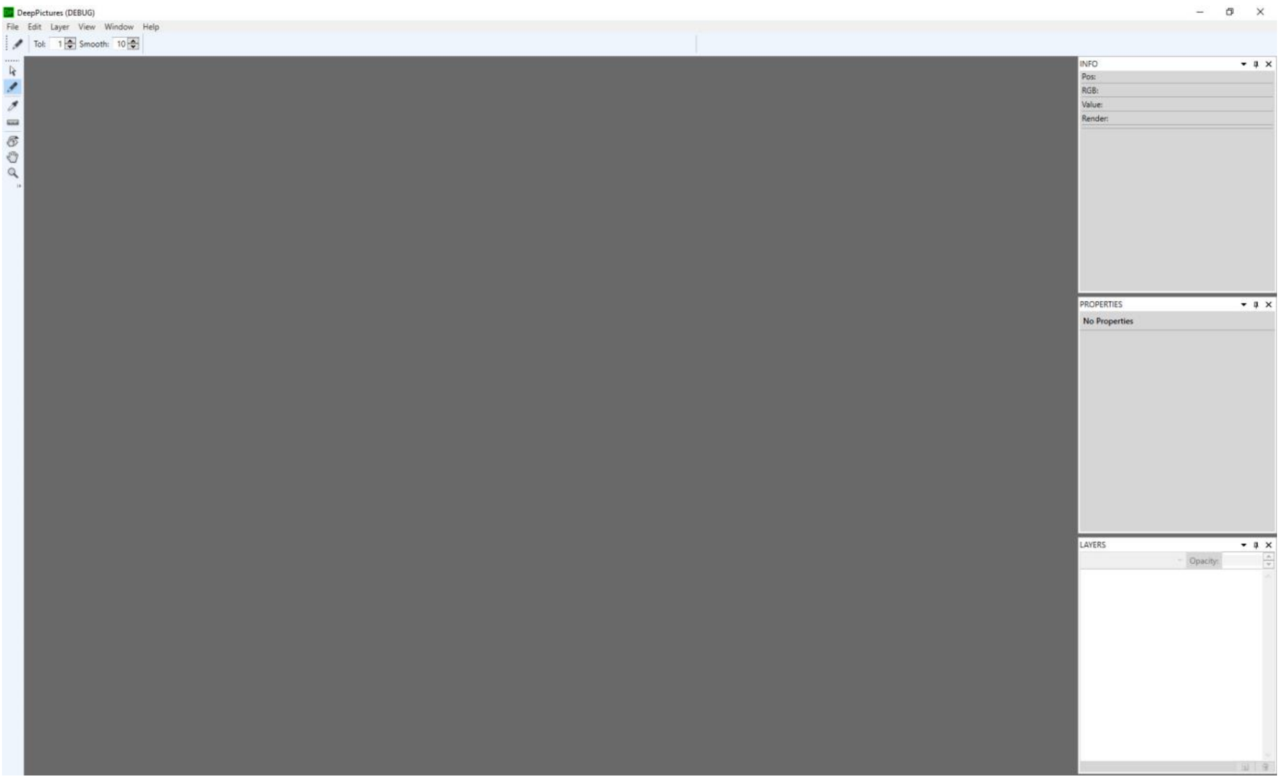

Figure 3: Layout of the graphical user interface.

The software has a simple layout, which can easily be reconfigured if necessary:

- In the centre there is the main view window, which can hold multiple documents views via a tabbed view interface.

- On top is a menu bar for general commands.
- On the left is the tool bar that allows selecting the current tool for mouse interactions in the document view.
- Right beneath the menu bar is a tool options bar which gives quick access to important options of the selected tool.
- On the right is space for multiple larger panels via which e.g. the layers and their properties can be accessed on the right

All these elements will be described in the following sections.

2.1.3 Main view

The main view window offers two views: (a) 2D view of the content (e.g. a picture file) and (b) a 3D preview of the relief. In the 2D view the user can interact via mouse or pen and can draw freely to segment the image into individual parts. To actually create the relief from the segments, the user switches into the 3D view in which the layers that compose the relief can be added and manipulated. Additionally the 3D view allows the height field to be rotated out of the plane to better investigate the current process of the relief. Figures 3-7 show examples of the application in the 2D and 3D view at various steps of the relief design process.

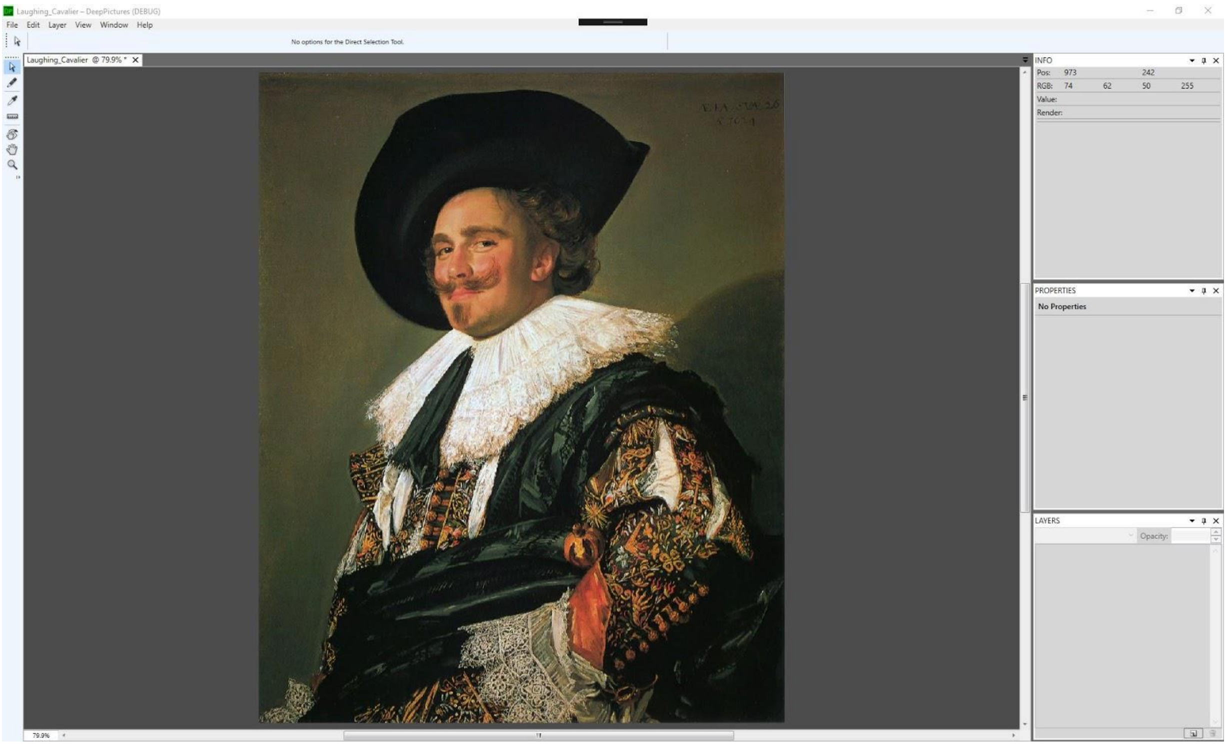

Figure 4: 2D View with an initially imported image.

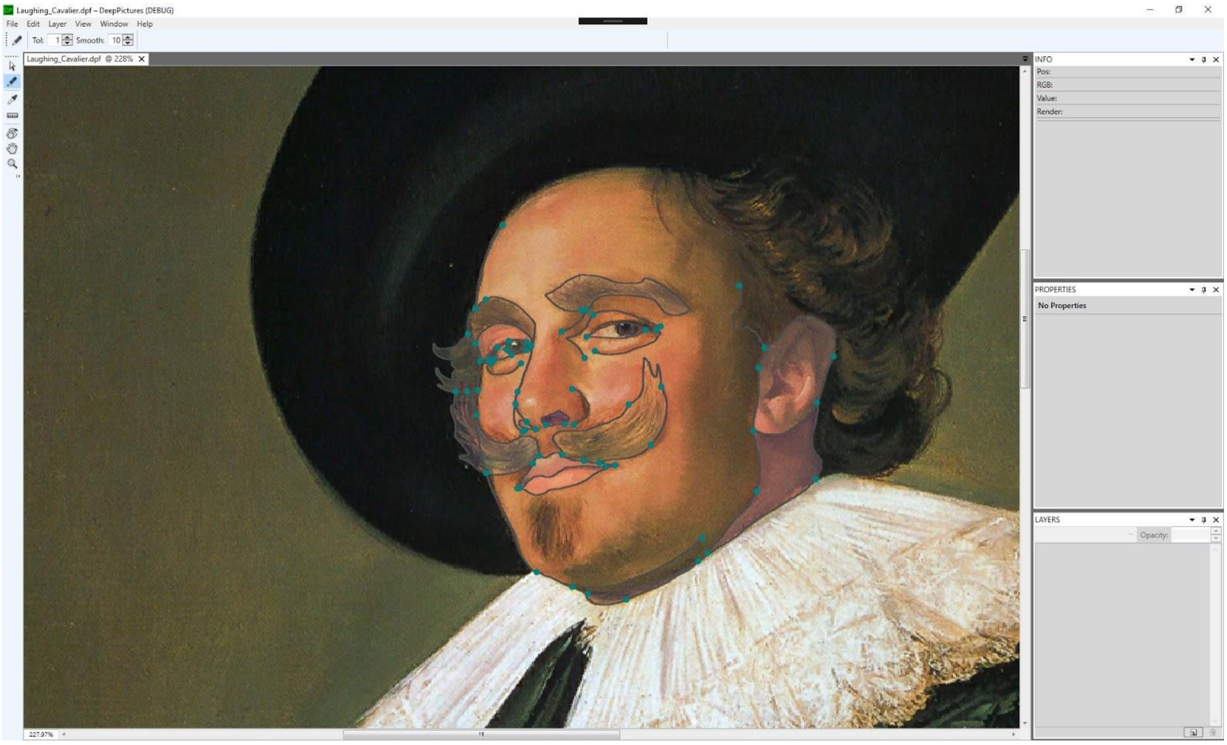

Figure 5: 2D View showing a segmented face created with the pencil-tool. Those segments are then further used for the relief generation process. Segments are automatically generated when closed contours are detected. To indicate them, segments are underlined with a random, semi-transparent colour.

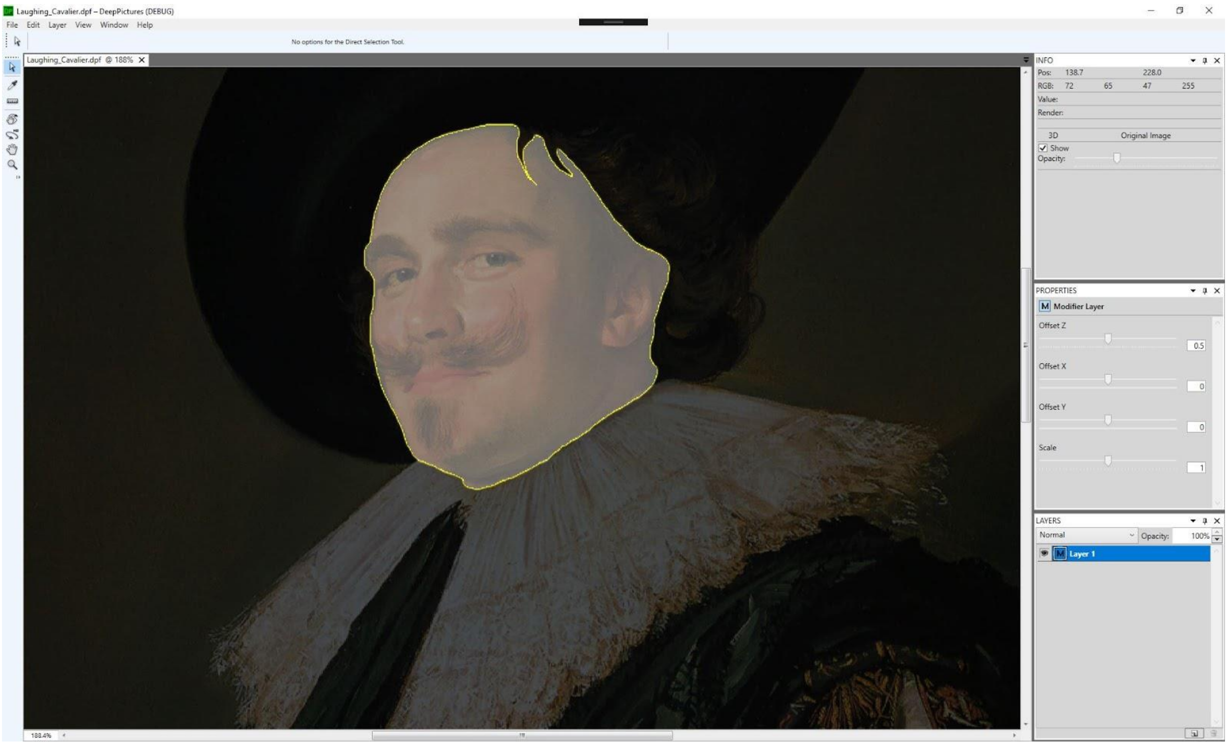

Figure 6: 3D View showing the original image in the background and a layer with a depth map representing the face. The height of a depth map is visible through its brightness, the brighter the higher a part is. The yellow highlighting shows the depth map associated with the selected layer. The segments that form the outer contour of the depth field were drawn and selected in the 2D view.

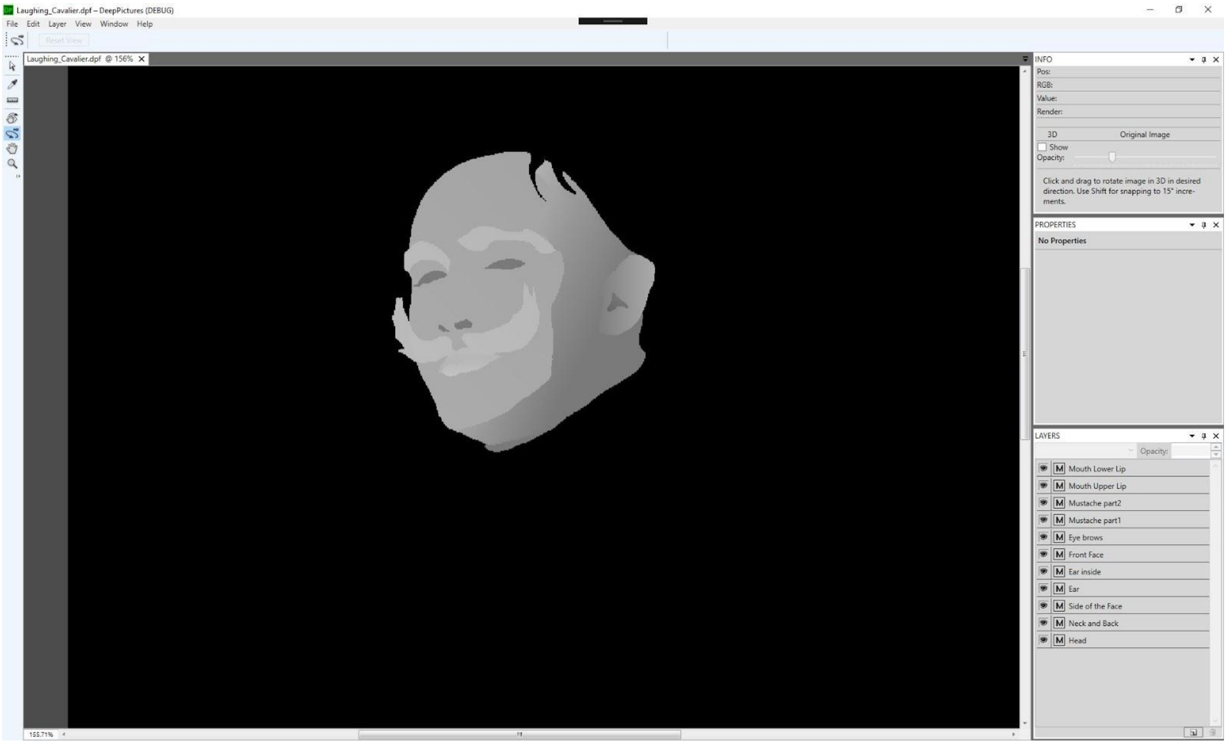

Figure 7: 3D View showing a depth map composed of several layers.

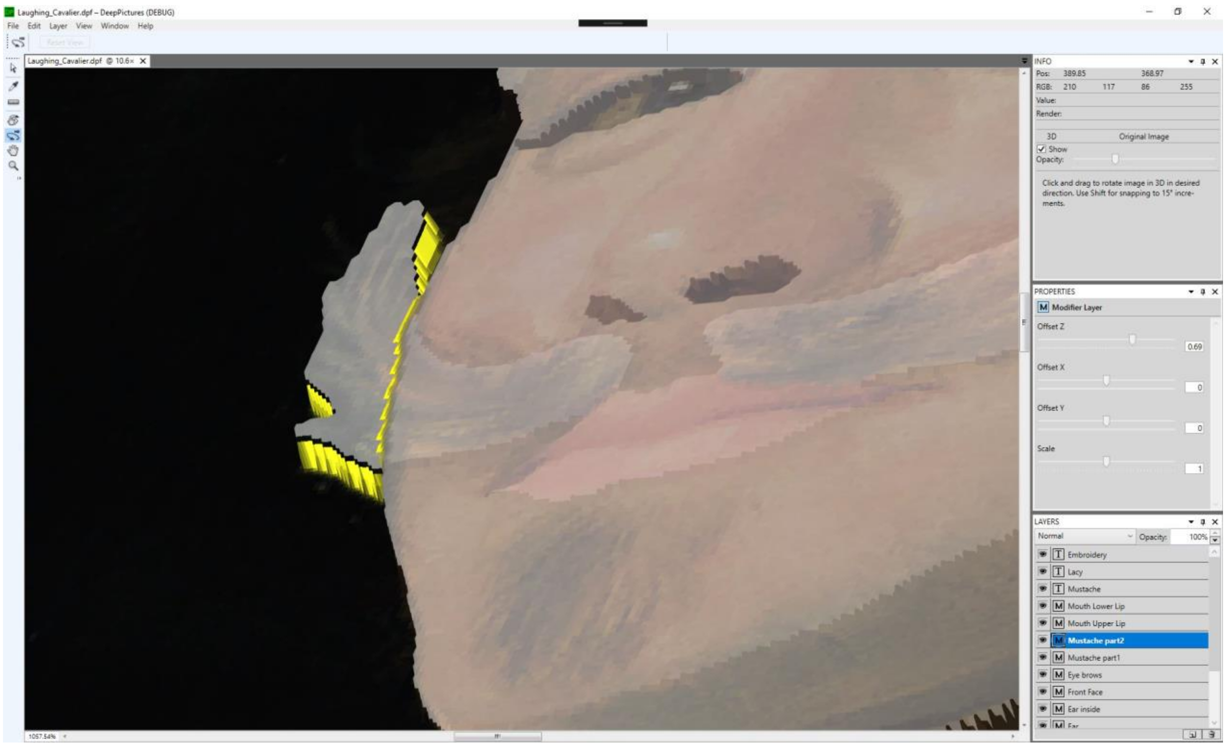

Figure 8: Example of a 3D view showing the depth map with the original image and at a further advanced state in a rotated view to facilitate the adjustment of the individual layers. The yellow outlining indicates which part of the depth map belongs to the selected layer.

2.1.4 Tools

Tools allow and facilitate the relief generation process, and describe the effect of user interactions in the document view with the mouse or stylus. Only one tool can be active at any time in the toolbox. Each tool has a quick access key to allow fast switching between tools without removing the mouse from the document view. If the user keeps pressing the quick access key for longer than 2 seconds, such a switch is only temporarily with automatically switch back to the previously selected tool when the user stops pressing the key. The toolbox may adapt depending on the currently selected view (see Figure 9). The difference is that in 3D the Pen tool for drawing is not available; instead there is an additional rotation tool for allowing 3D-rotations.

Further information and control over the individual tools are provided by the *tool options bar*, (see Figure 10, which shows as an example the toolbar of the Zoom Tool).

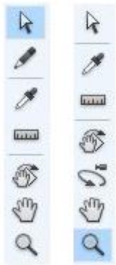

Figure 9: Toolbox of the application in the 2D view on the left side, and in the 3D-View on the right side. 2D-Toolbox from top to bottom: Selection, Pen, Pipette, Ruler, Rotate 2D, Hand and Zoom. 3D-Toolbox from top to bottom: Selection, Pipette, Ruler, Rotate 2D, Rotate 3D, Hand and Zoom.

Figure 10: Toolbar for the Zoom tool, which enables further control over the tools function.

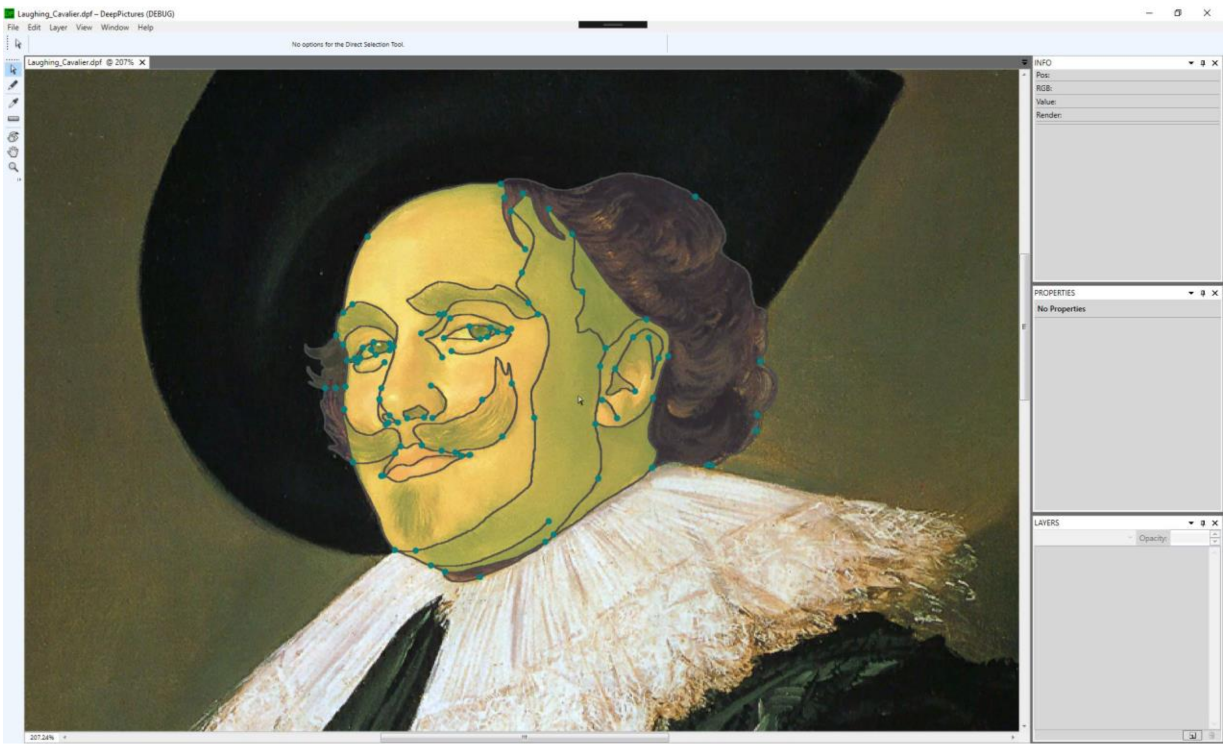

Figure 11: Multiple segments are selected. Selection is indicated by the yellow-greenish highlighted colour of the segments.

The **Pen tool** allows the user to draw segments onto the image. These are used to segment the image into parts that require different treatment (see section 2.2.1). Mostly these lines indicate that there should be a depth transition in the relief. It can also be used to select different depth-creation methods for different parts. In any case, the exact outline ensures that the resulting relief will be true to the original. Figure 11 shows example drawings on top of the image. Each segment or a group of segments can then be selected, using the Selection tool. Selected segments can then be associated with a new layer. For more information on layers see section 2.2.2.

The **Selection tool** can also be further used to select curves and/or curve points (indicated by the blue filled circles), which can currently be used to delete them. In the future, it will further be possible to move the individual points or bend the curves to allow small adjustments of the segments. Also, future layer types might not only be just applicable on a segmented region, but might also use the individual points and the curves to allow a more detailed control over the depth map.

The **Pipette tool** is currently more a debugging aid and measures the colour value or other properties of the pixel beneath the Pipette tip. That information is then displayed in the Info Panel (see section 2.1.5.1).

To measure the distance of two points of the image, and thus of the final relief, the **Ruler tool** can be used. So far, it only provides information in the XY plane and ignores the relief’s depth. Future plans consider adding that feature, too.

The **Rotate 2D, Rotate 3D** and **Hand tool** are used for navigating in or out of the image plane. As a fluent navigation is very crucial to get a good plastic impression in order to estimate the tactility, a lot of effort was placed into these tools. The **Hand tool** is similar to Photoshop’s Hand tool with which the user can drag and thus move the image to a desired position in the XY-Plane of the user’s view. The **Rotate 2D tool** is also similar to Photoshop’s tool, and allows rotating the image in the XY-Plane of the current view’s centre (see Figure 12). Therefore, this control is available in both the 2D and 3D views, which we found very important. The **Rotate 3D tool** covers the remaining degrees of freedom and is used to rotate the relief or depth map out of the view’s XY-Plane, which facilitates the proper adjustment of the individual layer depth maps.

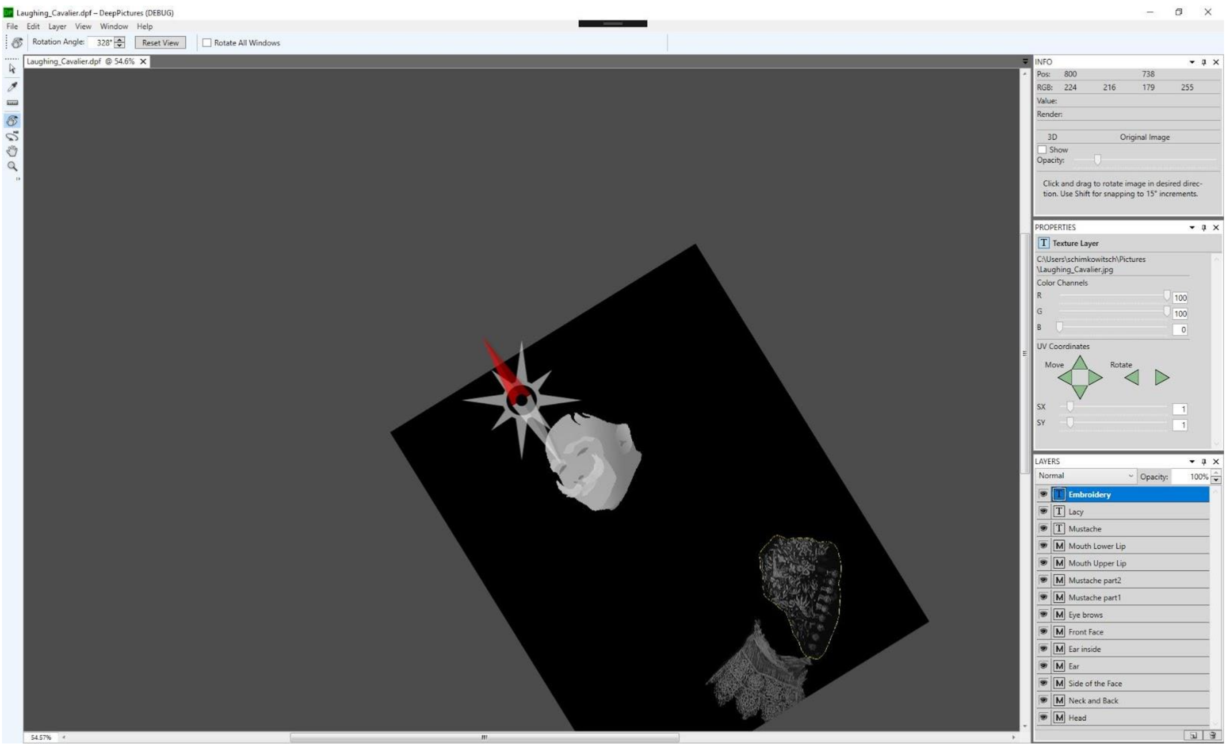

Figure 12: Using the Rotate 2D tool to rotate the image in the view’s XY-Plane. The displayed compass rose shows the current angle of the rotation, and is only visible during the rotation process.

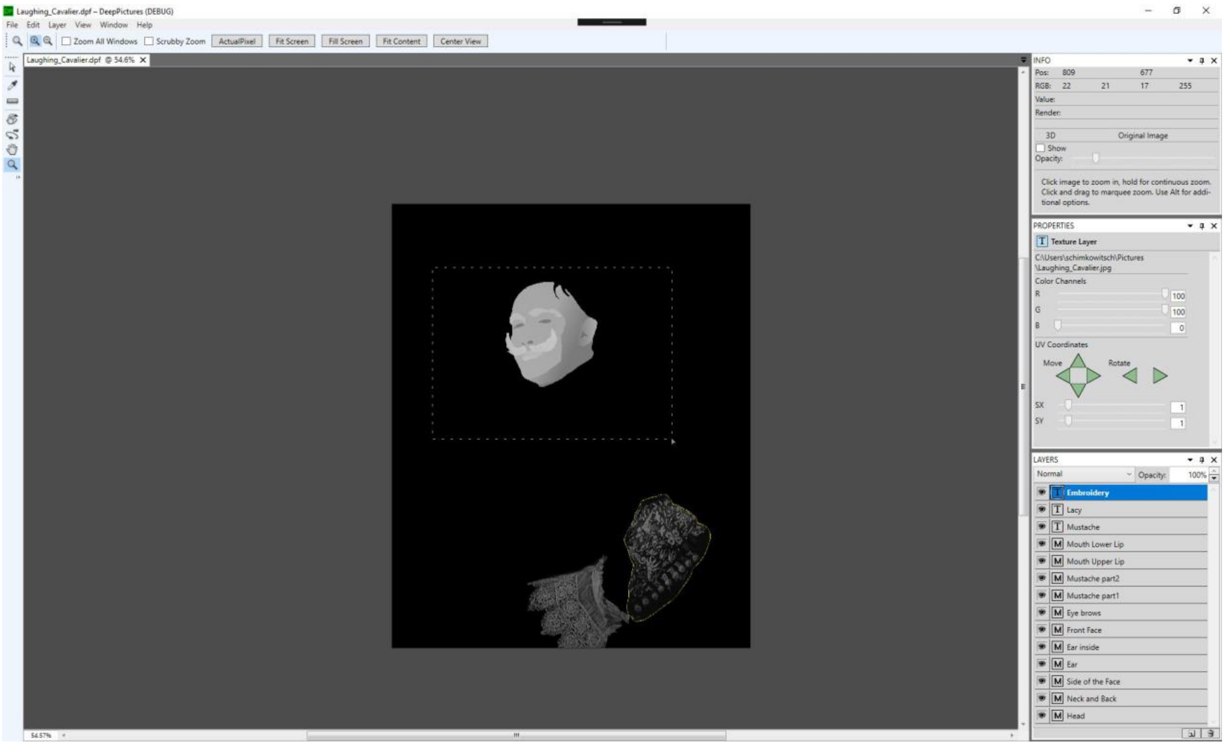

Figure 13: One way to use the zoom tool: drag a dashed rectangle around the region that shall be zoomed in. The size of the rectangle decides the zoom factor.

The **Zoom tool** is used to zoom in or out, either to get a better look unto the details of image or of the relief, or to get an overview. Multiple modes are supported. In normal mode, the user can either just press the left mouse button on a point in the image, and the view is smoothly zoomed in or out with the clicked point used as the zoom centre. Or the user can drag a rectangle around the region that shall be zoomed in, see as an example Figure 13. In scrubby zoom mode, the user presses the mouse at the desired zoom centre, and then drags left or right to adjust the zoom level. Further, ALT-scroll wheel can be used to zoom in predefined increments. In addition, several short cuts are defined to set special views, like 100%, fit all or fit width. Providing multiple modes helps the designer to quickly get the desired view.

2.1.5 Panels

Currently three panels on the right side offer access to, and further information on the content of designated parts of the tactile relief data file: an info panel, a layers panel and a layer properties panel that changes its contents according to which layer is currently in use.

2.1.5.1 Info Panel

Additional information like the position of the mouse cursor relative to the image and colour values at that position are displayed in the info panel. Also, it shows a usage description for the selected tool, if available. See Figure 14 as an example.

In the 3D-view it also offers control over an image-overlay on top of the relief. It is a general purpose place that can be used to add further information or controls in the future, and is especially useful during debugging.

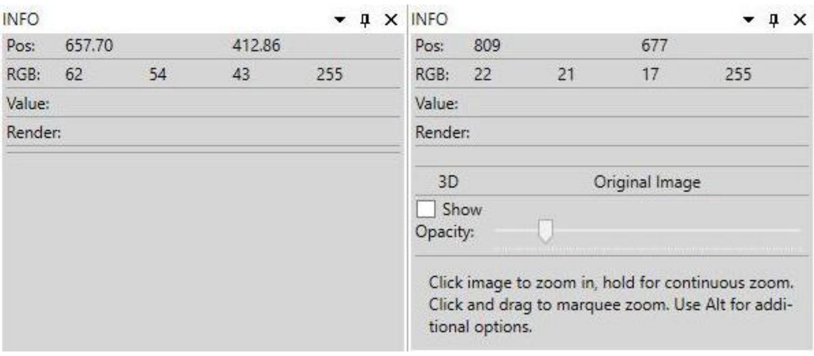

Figure 14: The Info Panel as it is displayed in 2D (left image) and in 3D (right image), showing further information about the image, like the colour value at a given position. In 3D, it also controls whether the original image should also be shown on top of the relief and, if so, with which intensity.

2.1.5.2 Layers panel

As explained in the beginning, the designed relief is represented by its depth map. In order to help designing this depth map, it can be composed of so-called Layers, each representing a part of the image and rules of how that part is to be processed or rather interpreted as a height field. Thus, all layers define a depth map in itself, which are then combined to create the final depth map of the tactile relief. There are currently four different kinds of layers (see section 2.2.2).

How the layers together compose the relief is determined by four factors: the order of the Layers, their assigned Opacity value, their Layer Composition Rule, and their Visibility. To combine the layers, the program starts with the bottommost layer as the base and then applies each consecutive layer on top of that to result in the final relief. The Opacity can be set in the upper right corner of the layer panel (see Figure 15), and controls how strong that layer’s influence on the composition is. The Layer Composition Rule (in the upper left corner of the layer panel) determines how that layer is combined with its preceding layers. So far, there are five rules:

- Normal, which simply overwrites the existing values with the layer’s non-zero height values;
- Additive, which adds the layer’s height-values to the existing ones (see Figure 16, left);
- Subtractive, which subtracts the layer’s height-values from the existing ones (see Figure 16, right);
- Minimum, which compares the existing values with the non-zero height values of the layer and lets only those values through which are either equal or lower than the layer’s height values;
- Maximum, which is similar to the Minimum rule, but which only lets values through that are either equal or higher than the layer’s height values.

The Visibility of a layer is indicated by an eye symbol on the left side of each layer (see also Figure 15). Similar to Photoshop’s layer system, the eye shows whether an layer is “visible”, in this case active or rather used in the relief composition, or “hidden” and thus not a part of the relief. This is useful to try different variations, or to disable some effects in order to better see the effects of underlying layers.

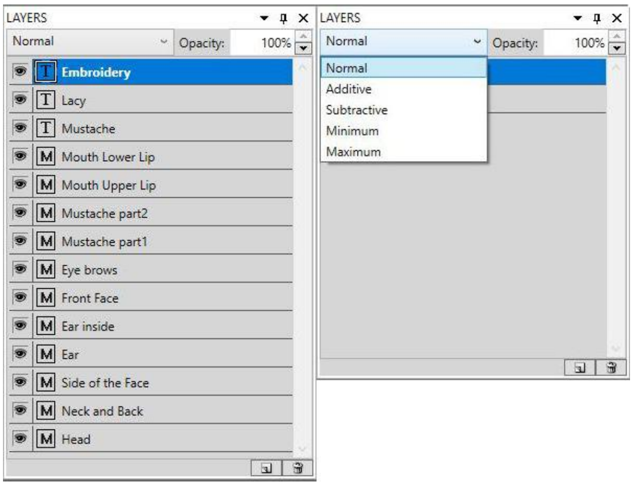

Figure 15: Two examples of the Layers Panel. The one on the left shows several layers that compose the final relief, and the one on the ride side shows the dialog for choosing a layer composition rule that can be applied for the currently selected layer.

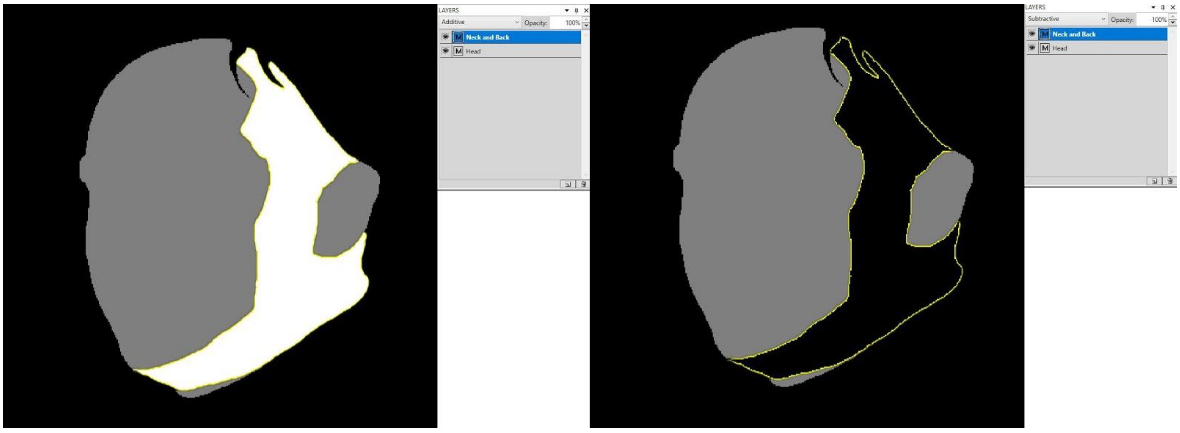

Figure 16: Example showing the effects of two Layer Composition Rules on the depth map. The head layer has a constant height, shown as gray. For display purposes, the layer above has a strong constant value. Left: Additive. The value is added resulting in a high (bright) value. Right: Subtractive. The value is subtracted resulting in a low (dark) value.

2.1.5.3 Layer Property Panel

Each layer has different properties, or settings, which control the processing. Three of these settings have already been discussed in the previous section, Visibility, Opacity and Layer Composition Rule. These are the same for each layer type and can therefore be directly placed in the Layers Panel. For the type specific properties, the Layer Property Panel is available that displays the adjustable parameters of the selected layer (or layers, if the multiple layers of the same type are selected). For the currently four different layer types, there are four different Properties Panels available. These are described together with their layers in section 2.2.2.

2.2 Subtask 5.1.2 “Extension of design tools”

2.2.1 Image segmentation

A crucial part of our design process is to analyse the image and to segment it into semantically coherent parts. These parts typically belong to the same object depicted in the image, and should have its distinct depth, which is determined by the spatial interaction of the different parts, e.g. if some object A overlaps an object B, A should be higher (nearer to the viewer) than B, at least at the area of overlap and its immediate surroundings. Segments could also be used to indicate different areas, where different textures should be applied, or they could be used to manually highlight important features, like important patterns and so on. As this is such a fundamental part of our design process, we decided to make this a corner-stone of the developed software: A Segment is a part that requires different algorithms or parameters to process.

The user defines these segments by tracing them on the 2D image with the pen tool (see for example Figure 5 and Figure 11). Closed contours are then recognised as valid segments, which are indicated by underlying them with a semi-transparent random colour. Those segments can be dynamically changed, by splitting or joining them, or by drawing holes later on in them. See as example Figure 17, in which the front face segment is yellowish highlighted and in which the mouth, nose, eyes and eyebrows, and the moustache are further segmented.

In our previous workflow this task was usually done in Adobe Photoshop, drawing thin black lines as the separation lines. This made subsequent modifications tedious, had no real notion of segments, and the pixels occupied with the black line could not belong to any segment. All these problems had to be managed manually throughout the design process. This new integration into the dedicated software and the underlying segment data-structure promises to overcome all these limitations, and offers a lot more flexibility.

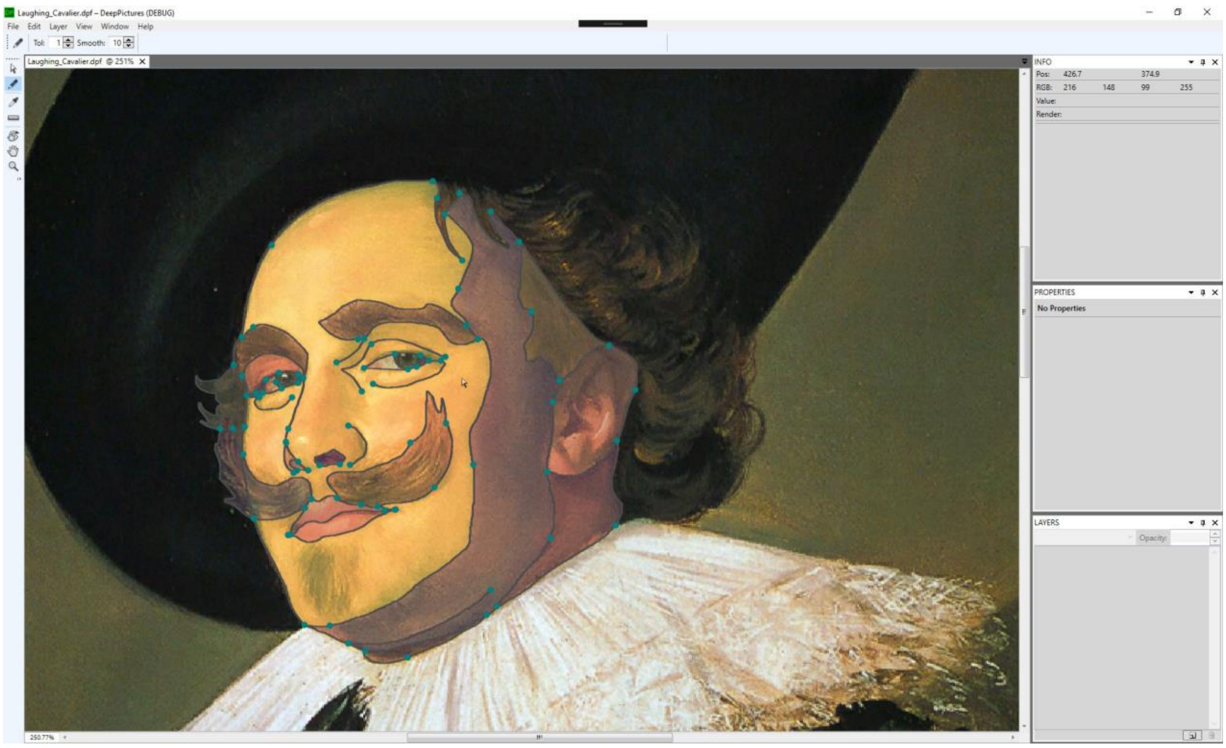

Figure 17: Face segments showing the decomposition of the face. The front face is highlighted to indicate that the mouse cursor hovers above it.

2.2.2 Layers

Different parts of an image may depict different types of content. For some of these special modelling or processing tools may be better suitable than others, and different regions may require different settings and parameters. In order to manage these differences in the software, we took the concept of Layers, similar to Adobe Photoshop's Layers.

Each Layer represents a specific algorithm, together with its settings. As multiple Segments may belong to the same object and need to be treated in the same way, multiple Segments can be assigned to the same Layer. Similarly, as the same region may require different processing steps, the same Segment can be assigned to multiple Layers. For different algorithms, different types of layers exist. Currently four types of Layers are already implemented: modifier layer, human layer, mesh layer and texture layer. Having this framework in place, new modelling techniques or algorithms can be easily added, and existing layer types modified and extended, in order to satisfy the requirements we find during designing the reliefs in the second phase.

In the following sections, we briefly explain the currently existing layers, together with their settings, as exposed in their layer property panel.

2.2.2.1 Modifier Layer

The Modifier Layer is a simple type of Layer that puts a flat surface, initially parallel to the image plane, into the Segment, see Figure 19.

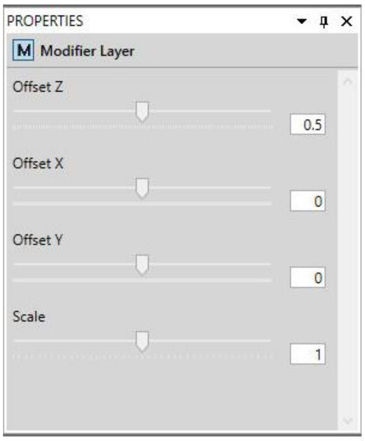

Figure 18: Properties Panel of a Modifier Layer, with the adjustable parameters Offset Z, Offset X, Offset Y and Scale. Offset Z directly translates into the height values of depth map. Offset X and Offset Y both apply a linear gradient over the whole depth map, in X and in Y direction respectively, thus can tilt the plane. Scale stretches or flattens the height values.

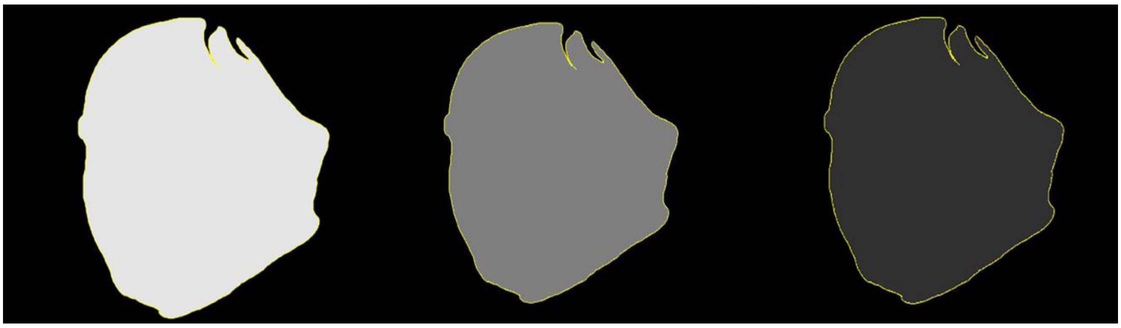

Figure 19: Modifier Layer example with a homogenous flat brightness or height value over the whole segment, shown with 3 different height values. Left: very high values; Middle: medium height values; Right: low height values.

That flat values can be changed by applying a linear gradient in X and in Y direction over the segment (see Figure 20). How strong that gradient is, depends on the values set via the corresponding Properties Panel (see Figure 18).

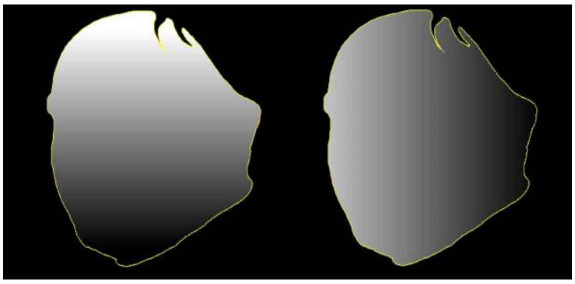

Figure 20: Modifier Layer with gradients in Y-Direction on the Left, in X-Direction on the right.

Furthermore, the scale property of a modifier layer additionally determines the contrast, for instance of the gradient. Figure 21 shows the scale properties influence on a gradient in XY-Direction.

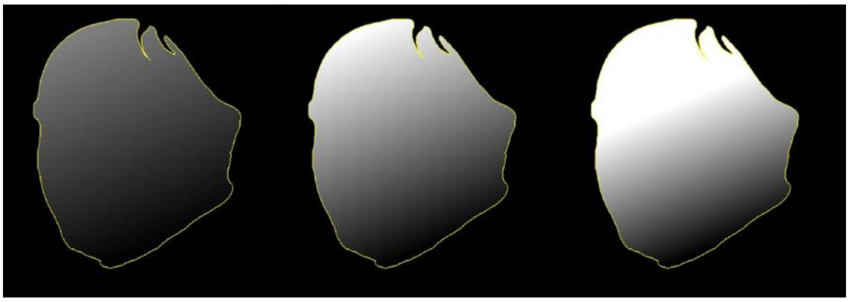

Figure 21: Modifier Layer: Changing a determined gradient in XY- Direction in the middle, either by flattening it with a low scaling value (left), or by stretching it with a high scaling value (right).

2.2.2.2 Human Layer

Images often contain humans, which are very difficult to model, and easily look alienating if not modelled properly. To simplify the modelling procedure, the software framework makes use of an integrated open source human base model (basically some kind of digital mannequin, see Figure 23), which can then be altered (modelled, posed) to fit the human in the given 2D image. For this purpose, the base model features a skeleton with manipulable bones which can be rotated about all three axes to achieve the desired pose, and morph targets on its body which allow the user to change the overall appearance of the digital model. This is our currently most complex layer type and is still work in progress.

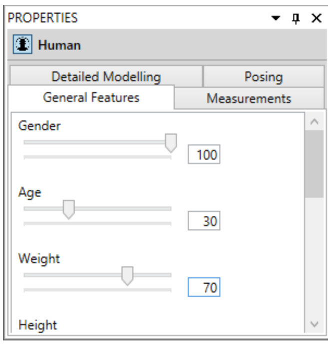

Figure 22: Properties Panel of a Human Layer. It has four tabs that are organising various functionalities: (a) Parameters for the general characteristics of the body, the “General Features”; (b) Body measurements can be modified under “Measurements”; (c) Parameters for bodily attributes in detail, which are summed up under “Detailed Modelling”; and (d) Handles for rotating the bones of the underlying skeleton, which can be found under “Posing”.

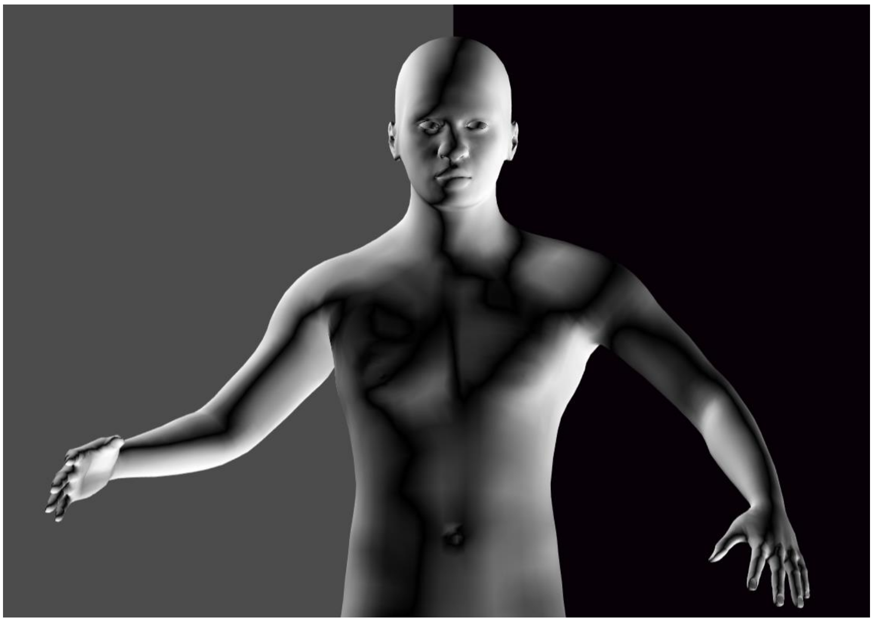

Figure 23: Human model.

2.2.2.3 Mesh Layer

It is sometimes easier to generate a model of an object or an image feature in a different modelling tool, or to take an existing model from one of the many 3D databases on the internet, and then to use that model in this tool. This is possible using the Mesh Layer. The Mesh Layer uses an associated segment as a window or a mask through which alone the used mesh is visible and hence influences the relief. Figure 25 shows an example of a mesh, where a Segment is used as its mask.

If no segment is selected, the whole image is considered as segment and functions as mask for the mesh. Figure 26 and Figure 27 show two examples of a Mesh Layer which uses the model of a human head. The first shows the head from the front, the second shows how that head might be used to recreate the Cavalier's face.

Further properties will be included in the future to allow easier alterations of the underlying mesh.

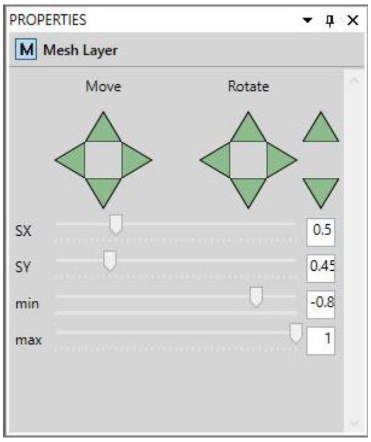

Figure 24: Properties Panel of a Mesh Layer. The Move buttons allow a positional adjustment of the mesh in x and y direction, the Rotate buttons allow a rotation around the x and y axis, and additional around the z axis. SX and SY allow individual scaling of the mesh in X and Y dimension. Min and max factors adjust the height values, i.e. stretch or flatten the resulting relief part.

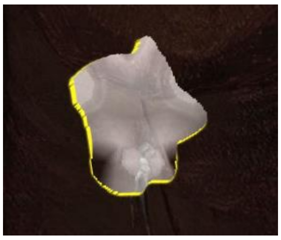

Figure 25: Mesh Layer in combination with a segment. The Segment functions as mask for the underlying mesh. Hence only those parts of the mesh that overlap with the segment are visible

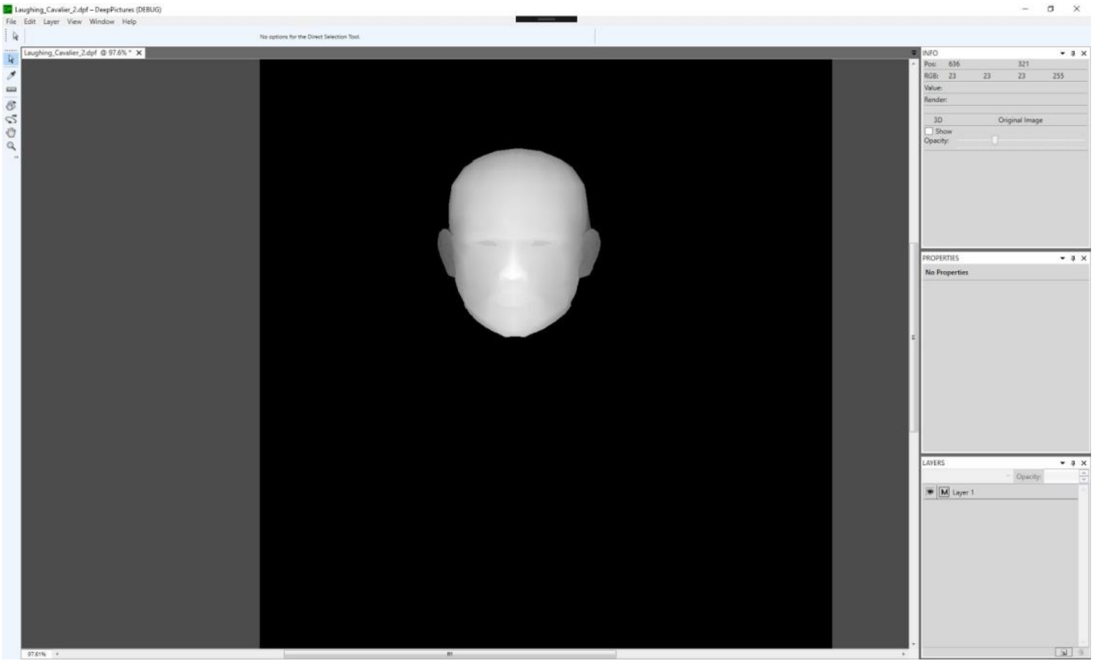

Figure 26: Single Mesh Layer showing the model of a human head in its initial form.

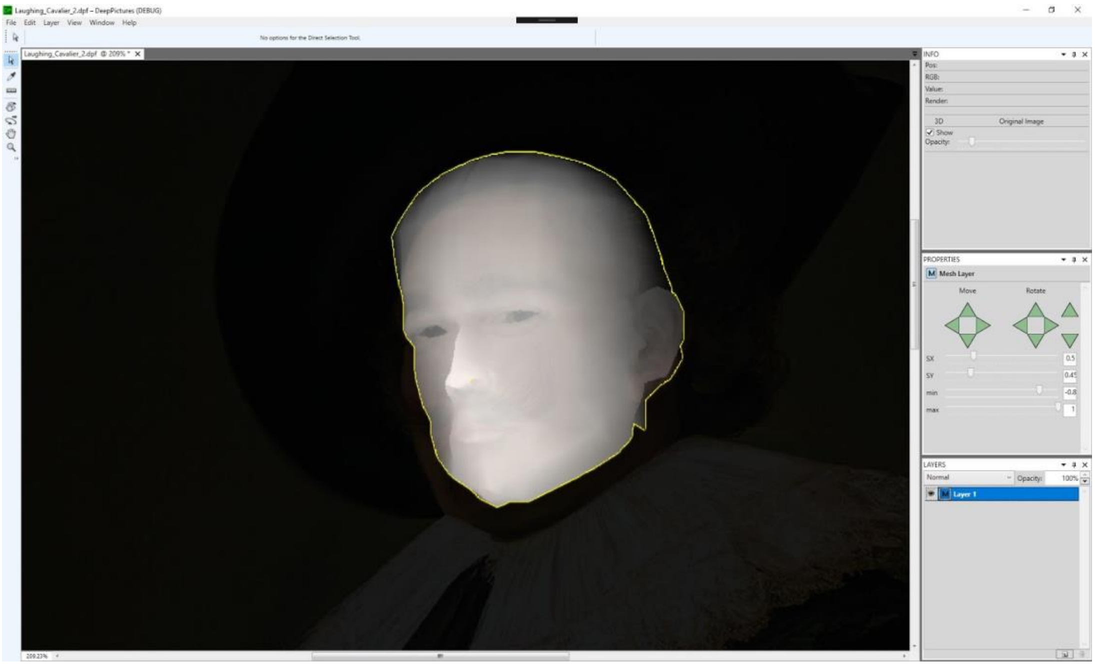

Figure 27: Showing the same human head as in Figure 23, but already positioned, rotated and scaled to approximately fit for example the Cavalier’s pose (see Figure 17).

2.2.2.4 Texture Layer

To recreate more elaborate features it is often sufficient to use and filter a part of the image (a texture) to get the required height values. Such features are for instance decorative patterns like the embroidery or the lace on the Cavalier’s sleeve, see Figure 28.

For that, the provided texture image is converted into a greyscale image, from which the brightness values can be directly used as depth values. For the conversion three handles are provided that determine how strong the influence of the individual colour channels are in the final composition.

In future work it is planned to include more filter functions that reduce the noise in the textures, see Figure 30.

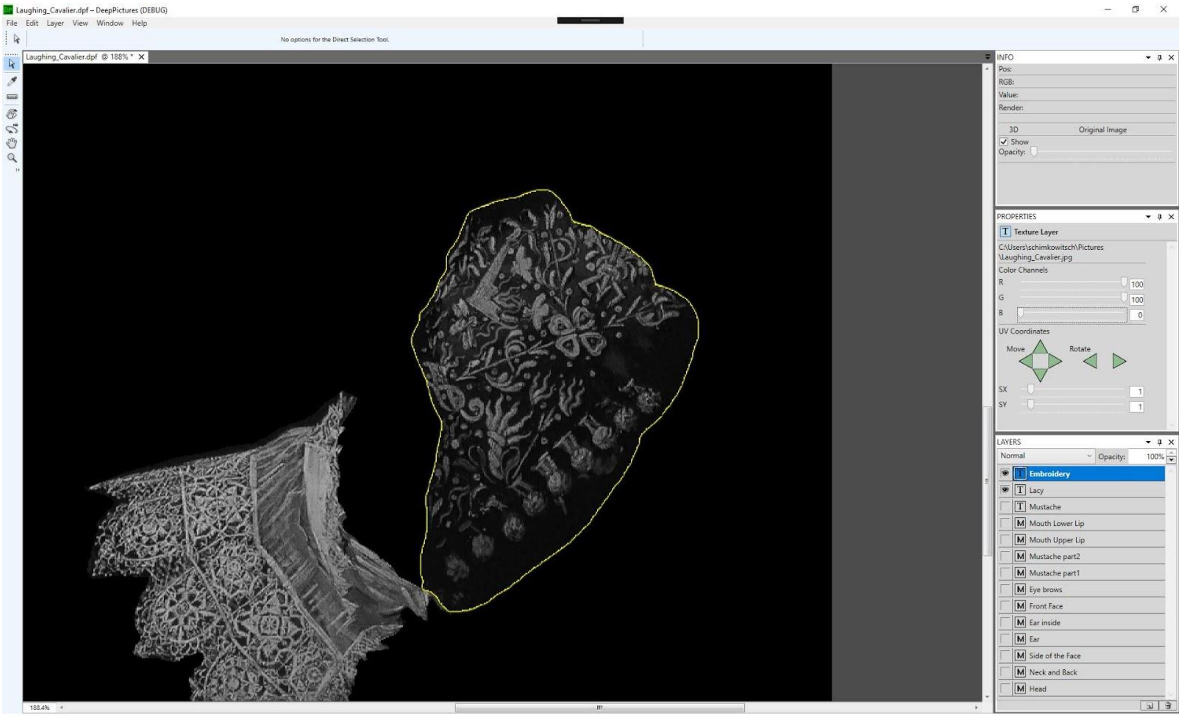

Figure 28: Two Texture Layers that use the information provided in the image itself to generate the layer’s depth map. This can be easily used for textures or simple, rather flat patterns.

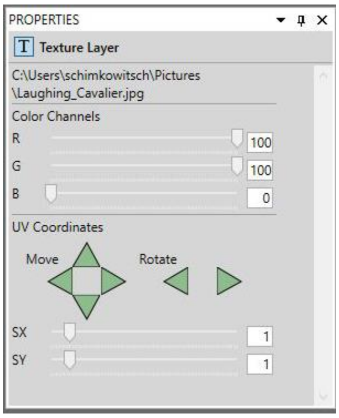

Figure 29: Properties Panel of a Texture Layer. It consists of two parts, in the first part the weights per colour channel can be adjusted, and the second part allows the relative positioning of the Texture compared to the underlying image.

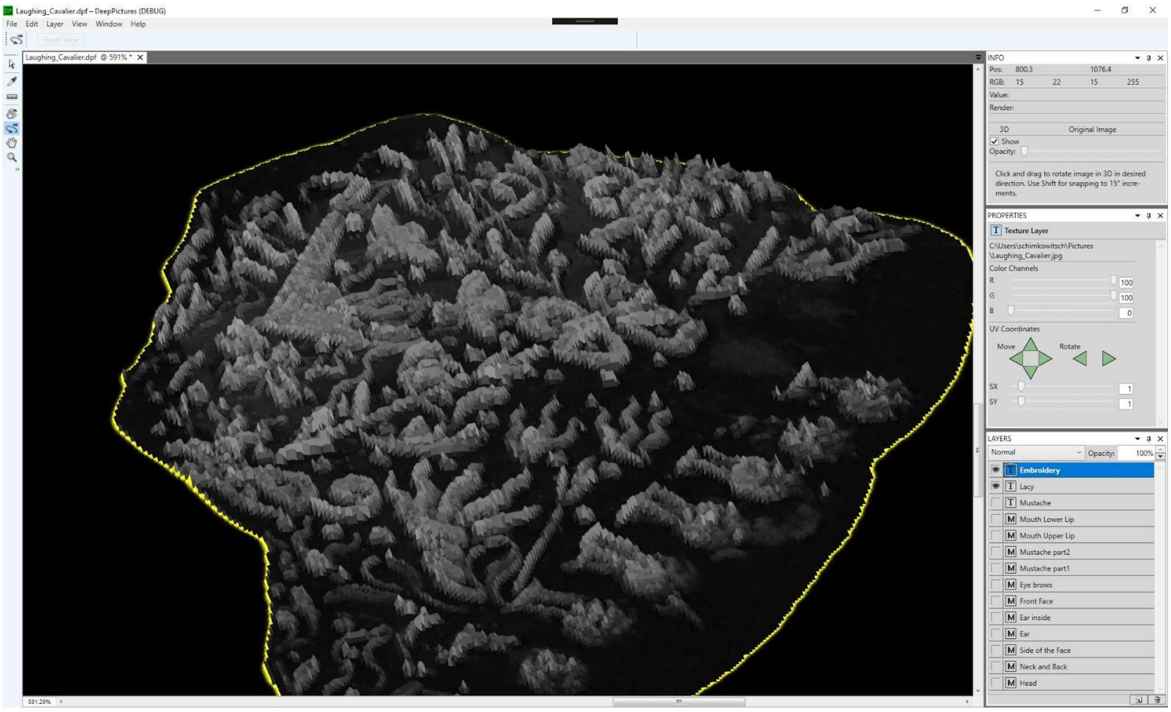

Figure 30: Texture Layer of the embroidery, seen in a rotated view. The resulting height values are over exaggeratedly displayed to show the effect. The texture information is directly translated into height values. It appears quite noisy yet, which will be improved by applying filters in future versions of this tool.

2.2.3 Summary

A large part of the general framework is finished, as well as several layer types have already been implemented. In the second phase, our efforts will be on the implementation of more layer types, as well as on improvements on the framework and already started layer types.

3 Task 5.3 “Context-sensitive tactile audio guide”

As already described in the Introduction, this task is concerned with the creation of an interactive audio-guide that can be controlled with finger and hand gestures directly on a tactile relief, based on optical hand tracking technology. We could build on a simple prototype created in a previous project, which is being improved and extended with new functionality. In order to better describe the implemented improvements, we briefly describe the previous prototype. A more thorough discussion can be found in [7]:

The original prototype consists of a depth camera (an Intel RealSense F200) connected to a computer and rigidly mounted above a tactile relief, which it observes (see Figure 33). In contrast to conventional colour cameras, a depth camera returns a depth value, i.e., how far an object at this pixel is away from the camera. First, the system is initialised with only the relief present. The system stores the acquired depth image, the so-called background image. Whatever is now put on top of the relief creates depth measurements that are nearer to the camera, hence can be easily detected by comparing the current depth image and the background image. As any objects may be added, the foreground is carefully searched for hands, and whether these hands form certain input gestures (see Figure 31b). Finally, depending on the gestures, real-time audio (and potentially other) feedback is given to the user. We differ between two types of gestures (see Figure 32): On-relief gestures for location specific information directly on the relief, and off-relief gestures for general information and other commands. The currently only *on-relief* gesture is the single-finger pointing gesture. With this the user can select one of several location specific audio comments. For this, a label-map (see Figure 31a) is prepared that indicates which positions belong to each comment. Currently two off-relief gestures exist. The user can stop any audio by showing a fist more than 20cm above the relief. And up to five general texts can be released, by showing one to five (1-5) fingers of a single hand in the same height range as the fist.

The work in ARCHES is divided in two subtasks.

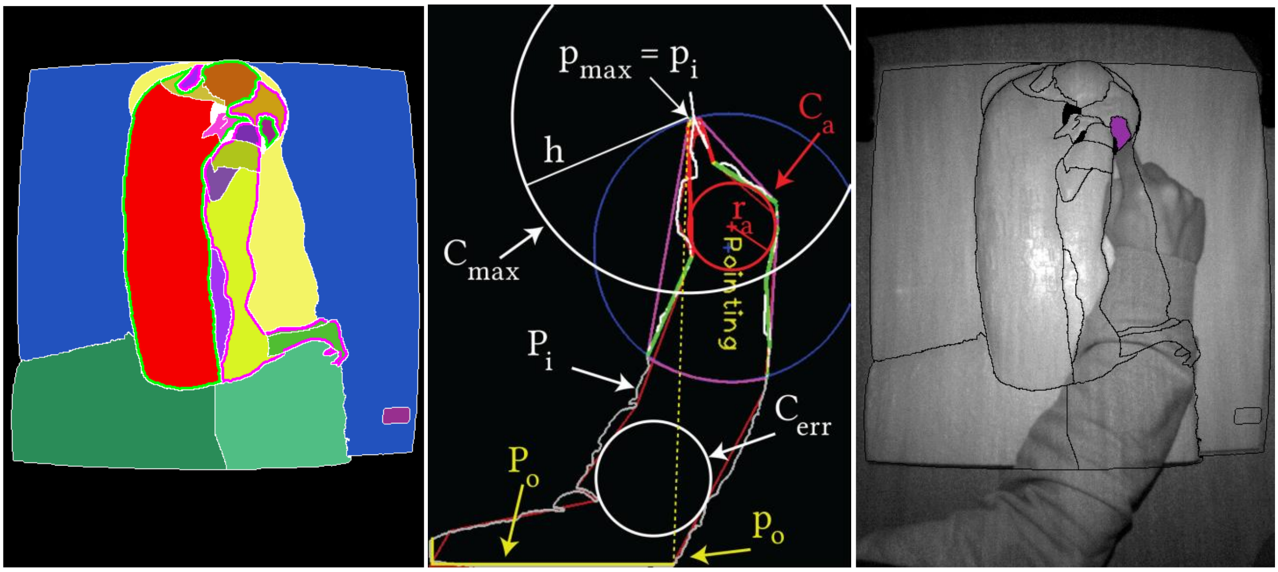

Figure 31: Gesture recognition. From left to right: (a) Hand-drawn label image, warped to camera space. Light-green and purple outlines indicate the merged base labels of the two figures; (b) Hand detection output and palm detection diagram. (c) Infrared image with superimposed label borders as output from our automatic relief calibration algorithm and touched label (purple).

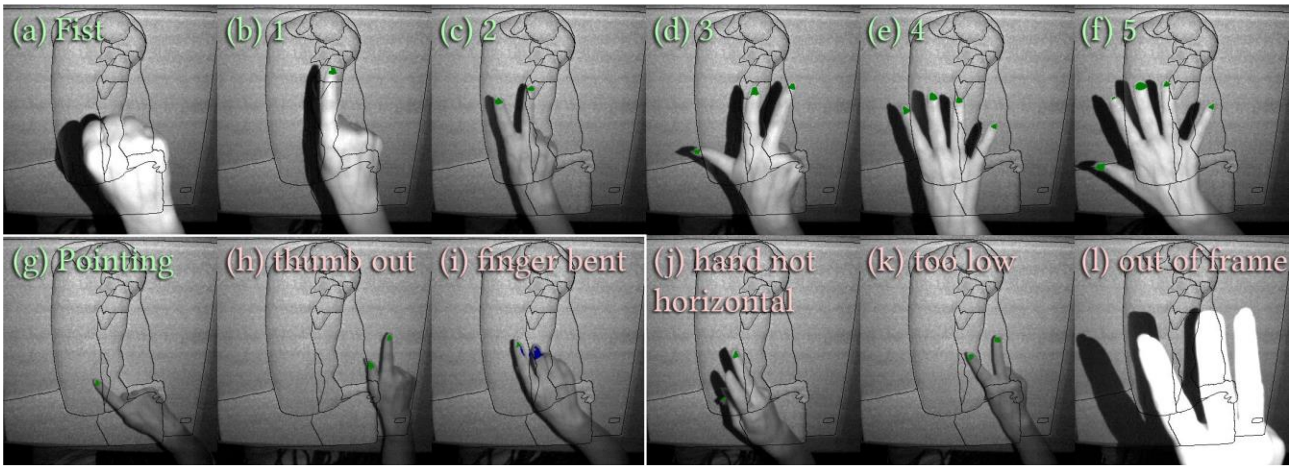

Figure 32: Gestures as performed by a user, with detected fingertips overlaid on infra-red images. (a)-(g) shows the full set of currently implemented gestures correctly detected by the system. Note that the 1-5 finger gestures can be performed with any fingers of the hand, thus this list is not exhaustive. (h)-(l) depict common hand posture errors which are not correctly detected.

3.1 Subtask 5.3.1 "Improvement of touch and gesture recognition"

This task was originally thought to just improve touch and gesture recognition over the previous implementation. The reasons for selecting the original sensor, the RealSense F200, have been discussed in [6]. As proposed in the Grant Agreement, we did a market analysis, whether improved depth sensors became available. Unfortunately, no new sensors entered the marked which are suitable for our project. The only additions were:

- an updated Intel RealSense called SR300 which did not improve over its predecessor in the near range which we require, and
- the ORBBEC Astra Series, which we expected to have lower performance in the very near range we require, as it is focused on for range sensing.

As the performance of the RealSense F200 was very satisfactory, we decided to keep working with this model. However, during first tests in London which had to be setup by non-technicians, it turned out that there are more severe problems with the old setup:

- Our first prototype setup had the RealSense camera mounted on a tripod (see Figure 33 left). This turned out to be not very usable, as people frequently bumped their head into it.
- This often accidentally moved the camera and voided the calibration and degraded the system's performance, until it was recalibrated.
- In general, the setup was deemed to be not ready for use in a museum, where a more robust setup is required.
- The setup requires careful adjustment, as it is required to exactly reproduce the same relief to camera pose if the initial setup, in which the label map was drawn. All six degrees of freedom of camera position and orientation have to match. This is achievable for an experienced operator, when the exact same tripod is used without moving its joints between setups, but is otherwise highly impractical.

The following sections highlight our efforts to overcome these limitations.

3.1.1 New hardware setup

We discussed several possibilities to overcome aforementioned limitations, and considered building a customised mount for the camera and relief to achieve a more rigid setup. Fortunately, we found a more elegant, full-featured and even off-the-shelf solution.

Figure 33: Comparison of old and new setup, © ARCHES project team. From left to right: (a) previous prototype setup using a tripod mounted RealSense F200 depth camera. (b) User testing the new setup with HP Sprout, making a pointing gesture. Depth camera and projector are integrated in the top part of the device. The relief is simply placed on the table.

Our new setup uses the HP Sprout workstation, an all-in-one computer, built specifically for innovative desktop 3D interaction (Figure 33 right). The HP Sprout is nicely designed, and fits much better in a museum space. In addition, it features a large touch-screen which could be used to display additional interactive content in the future. There is also a webcam installed above the screen, which observes the user and could also be used in interesting ways.

The HP Sprout also has a built-in projector, normally used to project an additional screen on a detachable touch sensitive mat on the desk. This opportunity inspired us to use it to project images directly on the reliefs (see Section 3.2.5 for more details). Further, the touch mat could be used as additional touch input around the relief, which we keep in mind if this requirement arises.

The HP Sprout has a RealSense F200 sensor directly integrated in a beam extending from the top of the computer screen, mounted in about 60 cm height, pointing down to the desk, where we place the relief, directly in front of its monitor. This beam also features an additional high-resolution camera pointing down and a desk lamp, which could be used to light the relief in dark places.

This new setup tremendously helps during setup, as the camera on the HP Sprout always has the same height and tilt relative to a flat table it is placed on, leaving only the three degrees of freedom of placing and rotating the relief on that table for manual alignment.

3.1.2 Automatic relief calibration

Despite the advantages of the rigid setup with the HP Sprout workstation, alignment still turned out to be a tedious task and needs to be repeated each time the HP Sprout or relief was moved. Therefore, we decided to implement an automatic calibration process.

The content author draws the label-map now no longer over the infrared image, but over the depth map that was used to manufacture the relief. With the known size and height, a virtual 3D model is created, textured with the label-map. During start-up, the calibration automatically detects the relief in the depth-sensor’s point cloud, and recovers the relative Euclidean transformation that transforms the 3D model to the detected location in the point cloud. As the point cloud is given in the co-ordinate system of the depth camera, the recovered transformation can be used to render the textured relief model as seen from the depth camera, which produces exactly the same label image as required for our system.

The question was how to reliably detect the relief and find the transformation. One option we found in the literature would be to use fiducial markers which we did not want to apply on our reliefs as they disturb the tactile feeling. Another possibility would be to let the user identify special features like the corners (i.e. point on them) and use that as a starting point, but this is not practical in a museum environment, where a fully automatic approach is more desirable. Therefore, we were looking for an automatic approach that could constantly calibrate the relief during inactive periods, or detect the change of a relief and quickly calibrate on that one. A future, faster implementation could even enable real-time tracking and allow the user to move the object during exploration. We explored image-based object recognition methods, but found them difficult to use in our case, as these require distinctive patterns on objects, and our reliefs are typically single coloured. Based on these considerations we looked for a purely geometric, surface-based method, which is more suitable in our case.

We currently use our own implementation of a point-pair feature-matcher based on [5], followed by an ICP optimization based on [4]. Since the method is quite sensitive to surface variations of the models, we had to down-sample and blur the relief model depth map significantly to resemble the relatively low quality received from the depth camera. Further, the camera’s point-cloud is created from an average of the five most recent depth maps to reduce noise, and normals are computed using a plane-fitting algorithm over a 3mm radius to make them more robust to noise. To optimise the results, we carefully calibrated the infrared camera using OpenCV’s calibration methods [1] based on calibration images taken with a checkerboard pattern (see section 3.2.6 for more details).

Currently, object recognition takes around 5 seconds on average, which is very good in our static setup. An acceptable pose is recovered in about 4 out of 5 cases, and is almost always recovered after a second try. The rendered label map still has a few pixels deviation (see Figure 31c), but only the border seems to be shifted to the left, while the contours around the figures fits tightly. We attribute this either to distortions in the plaster relief mould, resulting in an actual physical deviation. Or it could be caused by non-linear distortions in the depth values, received from the depth camera.

All in all, this new method could reduce setup times to a few seconds. Simply place the relief under the HP Sprout, press the calibrate button, and after a few seconds everything is ready. This can now be done, even by non-technically supervised people. If the relief is accidentally moved, the system is quickly recalibrated. This also prepares the system for a future use with multiple reliefs, which can be quickly replaced under the same HP Sprout.

3.1.3 Adaption and improvement of the algorithms

Compared to our previous setup, with the HP Sprout the sensor does no longer look straight down mounted over the centre of the relief, but is mounted higher and further back over the far end of the relief. This requires a higher setting of the RealSense’s motion-range trade-off, dropping the frame rate slightly to around 16.6 fps, which is still more than sufficient for a fluent interaction. In addition, higher depth-map filtering settings are necessary, which smooth away some detail of the depth map, and the effective resolution at the relief drops from 10.7 down to 7.3–8.2 pixel/cm8, making finger detection and localization more difficult.

Unfortunately, the RealSense is mounted in landscape orientation, and tilted so that the upper 20% of the sensor is worthless as it films the Sprout’s monitor. In the remaining vertical view, the used 42 cm high relief just fits in, with very little space left in front of the relief. Thus, the user’s hand is no longer detected when touching the lowest parts, as it is then already largely out of the camera view. Thus, reliefs need to be created a bit smaller to work smoothly with the new setup.

After some adjustments of finger and fingertip detection parameters, and, most importantly, making them varying with camera distance and current frame rate, the detection could even be improved than before. In addition, we profiled the core algorithms, and found several bottlenecks that slowed down the program in the previous implementation. The new implementation now runs significantly faster, and can now use the full received frame rate from the camera, which improves response times.

Overall, the new setup seems to have positive effects on the ease of use. The lower resolution and higher filtering of the depth map made the single-finger gesture detection more tolerant, as not perfectly hidden fingers are less likely to being detected. And there is a lot more room above the relief. People don’t accidentally bump into the camera, and off-relief gestures are easier to perform and to detect.

3.2 Subtask 5.3.2 “Improvement, extension of user interaction methods and content authoring tool”

As outlined in the Introduction, this phase was mainly concerned with improvements of the setup and the tracking system. Nevertheless, several other improvements could already been implemented. Most of them were a direct response to user feedback during tests with the participatory research group in London (see section 4.3).

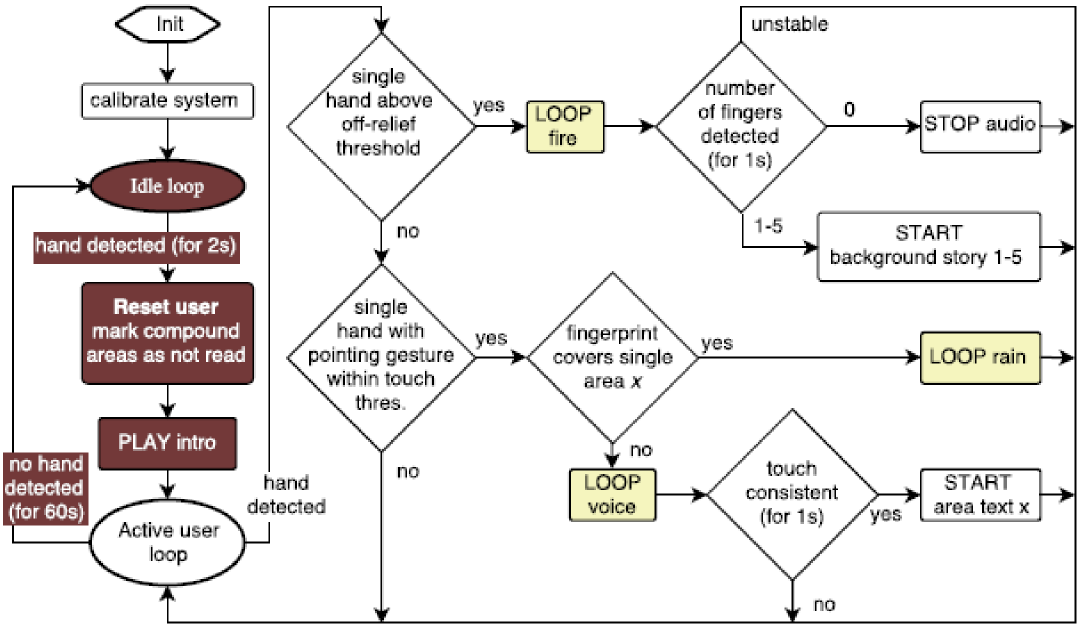

(a) Mapping of input to audio commands.

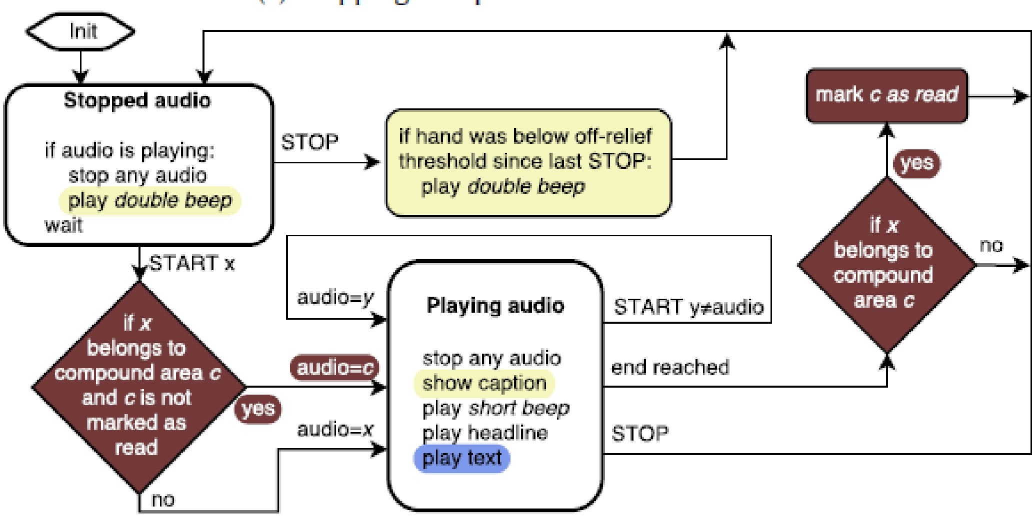

(b) State machine of the audio player.

Figure 34: Program flow and state changes of the user interface. Colour indicates differences between modes. Dark red: introduction and hierarchical exploration only in evaluation study 1. Yellow: additional sound design and captions only in study 2. Blue: no text in trainings-mode.

3.2.1 Additional sound design

The sound design of the original implementation was kept very minimal. For on-relief interactions, each triggering of a new area was accompanied with a *confirmation click sound*, followed by a short name of the region, and after a short pause, a detail description follows. Similarly, the triggering of off-relief gestures was also accompanied with the same *confirmation click sound*. In case of the stop gesture, that was all. For the number-gestures, the detected number was said, followed by the headline of the text, and then by the actual text.

Based on the feedback of the first evaluation in London, we implemented additional sounds to aid participants in making correct gestures:

- A *fire-crackling sound* indicates that the hand is over the minimum height, required for the off-relief gestures.
- Two further sounds indicate that the single-finger touch gesture is correctly detected:
- An *ethereal voice* is used on regions where descriptions are available, while
- A *rain sound* indicates that the finger is on the border between two or more regions, and the algorithm cannot determine which region to select.

Before these additions, touching the border between regions would not give any feedback, and users frequently assumed that the finger gesture was not detected. Now, when hearing the *rain sound* they can move the finger more towards the desired region until the touch is detected.

Alternatively, a user can also use the sounds to scan the bounds of a region with their fingers. An interesting idea for future investigations would be to try distinct sounds for each region: It could help orienting, but could also be too confusing.

We took care to select subtle ambient-like sounds, so that people who need them can clearly hear them, while it does not render the spoken text incomprehensible, nor disturb people who do not need them. Still these sounds double as a reminder also for those who do not need them: A reminder to relax their hands into a non-command pose in order to not accidentally trigger another sound. As further cue, the sounds get louder the clearer the gestures are detected, implemented as percentage of the frames in which the gesture was detected over the last 0.75 seconds. Incidentally, this also produces a nice fading effect. Furthermore, the sounds pans from left to right depending on where the gesture was detected in the camera frame. This might help some users keeping the hand centred under the camera.

Finally, we changed the *confirmation click sound* to a *short beep* which can be easier recognized, and implemented a distinct stop sound as a *double beep*, which is now always played when a sound is stopped: either by the fist-gesture or at the end of each description. This might help to distinguish short text pauses from the end.

3.2.2 Screen design

During previous tests, participants with low vision unexpectedly observed the debug views on the screen, which we had for the purpose of technical support. These were showing the output of the hand detection system on the left and on the right the output of the touch detection output superimposed on the infrared camera image. Some participants seemed to enjoy watching these views. They could see where the interactive regions are, and whether the system detected the touch event, as the interaction region then becomes coloured.

As the target group of ARCHES now explicitly also includes participants who can see, we added a visualisation of the detected fingertip pixels on the touch detection output, so people could directly see which fingertips are detected, and where exactly and how large the detected fingerprint actually is (see Figure 32). In addition, we added simple subtitles for all spoken texts in the lower half of the screen: One static text per description, font size 30, black text on white background. As the new setup with the HP Sprout has an integrated projector, a future implementation could directly project such visualization on the relief and fingers.

3.2.3 Training mode

The current prototype was designed as an installation in a museum, for people who are not familiar with the system. Therefore, the first interaction is to simply put the hands on the relief, which triggers a short introduction explaining the interface. After the system is not used for a given amount of time (currently 2 minutes), the system was reset and waited for the next user.

As this mode did not work well in the initial tests in London (see section 4.3), we implemented a training mode as an aid for the examiner. This is basically the same as the normal mode, but without description texts in order to minimise audio output, and help the participants to better concentrate on the examiner’s words and on practicing the gestures. In this mode, the on-relief gestures only trigger the confirmation beep sound, and the name of the touched parts, and the off-relief interaction repeatedly tells the number of detected fingers (0–5).

3.2.4 Authoring tool for audio guide

As discussed in the Introduction, the Tactile Audio Guide evolved from a previous project. This prototype implementation was only targeted at a single relief, for demonstration purpose, and all interaction and content was hard-coded in the program code. In ARCHES we build on this prototype, but intend to open it for many reliefs. Therefore it is necessary to remove the hard-coded content for the one relief from the program, and let the program load it from multiple places on demand.

We developed a content description file format and a folder structure for the content, and started implementing a first version of an authoring tool, that allows creating the description file and folder structure for the reliefs.

Currently it is implemented as an internal tool, but our intention is, that by the end of the project, this authoring tool will allow museum staff to easily edit the textual and auditory data for their reliefs, as well as to change the default settings.

As mentioned above the current implementation of the Tactile Audio Guide offers up to five general descriptions, and one description each of an arbitrary number of regions on the relief. Similarly, the current state of the Authoring Tool focuses on exactly this content, and already tries to create a hierarchy of regions, that will later allow to divide larger regions (e.g., a whole figure) into smaller regions (e.g., head, body, arms, legs,...). As the development started, before the reliefs for ARCHES were created, all examples are based on the relief for Gustav Klimt’s “Der Kuss” created in the previous project AMBAVis.

There is a single resource folder for the application in which all relevant content is stored (see Figure 35).

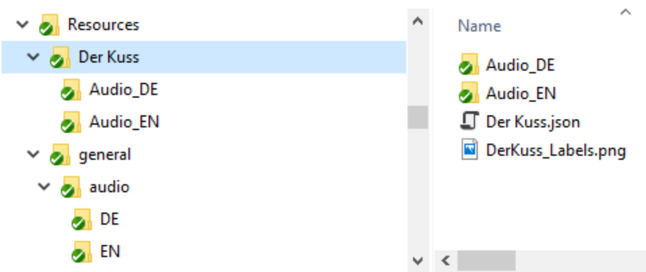

Figure 35: Resources folder.

All general content that is required for each relief is located under the “general” folder. It currently consists of the audio files for the sound effects as well as some text that is used for each relief, e.g., the spoken numbers 0 to 5. There are different sub-folders for different languages.

Each painting has its own folder inside the resource folder, containing all content specific to the painting:

- Audio files (.wav), which are stored in the corresponding subfolder of each language (e.g.

“Audio_de”, “Audio_en”, …),

- Painting label picture (.png), which specifies the interactive regions on the painting
- a JSON file (e.g. “Der Kuss.json”), as the root specification of the content.

The JSON file (see snippet in Figure 36) is loaded first. It contains all textual information, and links to all the other information.


```
"Name": "Der Kuss",
"SectionGroups": [
  {
    "Name": "Intro",
    "TextualAudioDescriptions": [
      {
        "Language": "en",
        "title": "Welcome to the interactive touch relief of Gustav Klimt's \"The Kiss\".",
        "Body": "Please explore the relief with your hands, as you like. If you want to know more about a
        "TitleAudiofile": "_Intro.wav",
        "BodyAudiofile": "",
        "CanBeInterrupted": true,
        "IsStoppable": true
      },
      {
        "Language": "de",
        "title": "Herzlich Willkommen zum Interaktiven Tast-Relief zu Gustav Klimts „Der Kuss“.",
        "Body": "Bitte ertasten Sie das Relief nach belieben mit beiden Händen. Um mehr über einen Bereich
        "TitleAudiofile": "_Intro.wav",
        "BodyAudiofile": "",
        "CanBeInterrupted": true,
        "IsStoppable": true
      }
    ],
    "Sections": [
      {
        "Name": "Brief description of the painting",
        "Color": "#FFFFFFFF",
        "TextualAudioDescriptions": [
          {
            "Language": "de",
            "title": "1. Der Kuss - Kurzbeschreibung des Gemäldes",
            "Body": "Auf dem Gemälde ist ein eng umschlungenes, junges Liebespaar zu sehen. Das Paar befin
            "TitleAudiofile": "_brief description of the painting_.wav",
            "BodyAudiofile": "_brief description of the painting.wav",
            "CanBeInterrupted": true,
            "IsStoppable": true
          },
          {
            "Language": "en",
            "title": "1. Brief Description of the Painting",
            "Body": "The painting shows two young lovers locked in an embrace. The couple is on a stretch
            "TitleAudiofile": "_brief description of the painting_.wav",
            "BodyAudiofile": "_brief description of the painting.wav",
            "CanBeInterrupted": true,
            "IsStoppable": true
          }
        ]
      },
      {
        "Name": "The artist Gustav Klimts",
```


Figure 36: Part of an example JSON-File describing the content for the interactive audio guide.

The structure of the painting is also described in the JSON file. As a preparation for the aforementioned hierarchical labels, a painting consists of sections that are assigned to section groups. Section groups and sections have textual and audio data per language, as seen in the file “Der Kuss.json”. At the start the program loads all JSON files of the paintings from the resource folder.

The painting of the kiss consists of the following sections, as can be seen in Figure 37: background, halo, signature, tendrils, meadow, male figure head, male figure cloak, etc.

In order to make it easier to handle, they are subordinate to section groups. As can be seen in the table in Figure 38, the section group “Malefigure” consists of the sections: head, coat, back and neck, left hand and the right hand of the man and the left hand of the woman. The Figure shows that the painting “Der Kuss” with Language “de” (short for German) and the male figure is selected. Now it is possible to edit the title and detail description of this section group. The audio file of each textual description can be played or can be set again.

If you click on a section (e.g. male figure cloak) on the label picture the tab automatically changed to “Section” and the correct section is selected (see Figure 39).

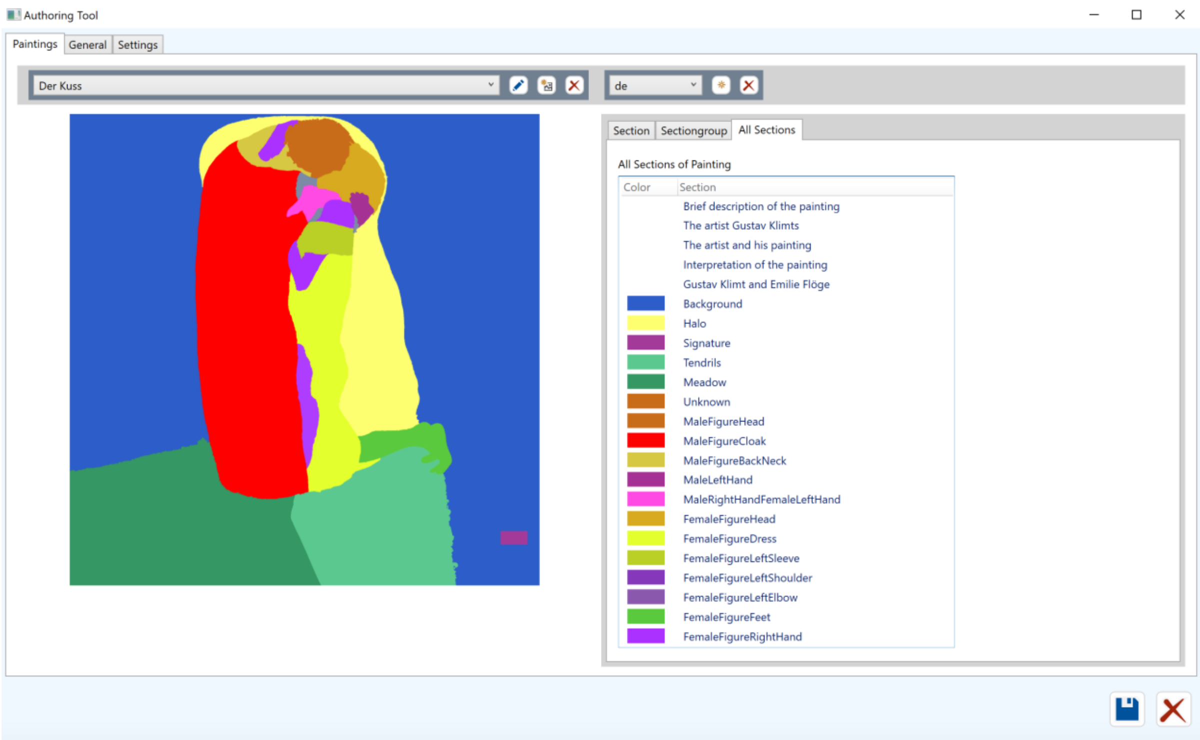

Figure 37: Sections of a picture.

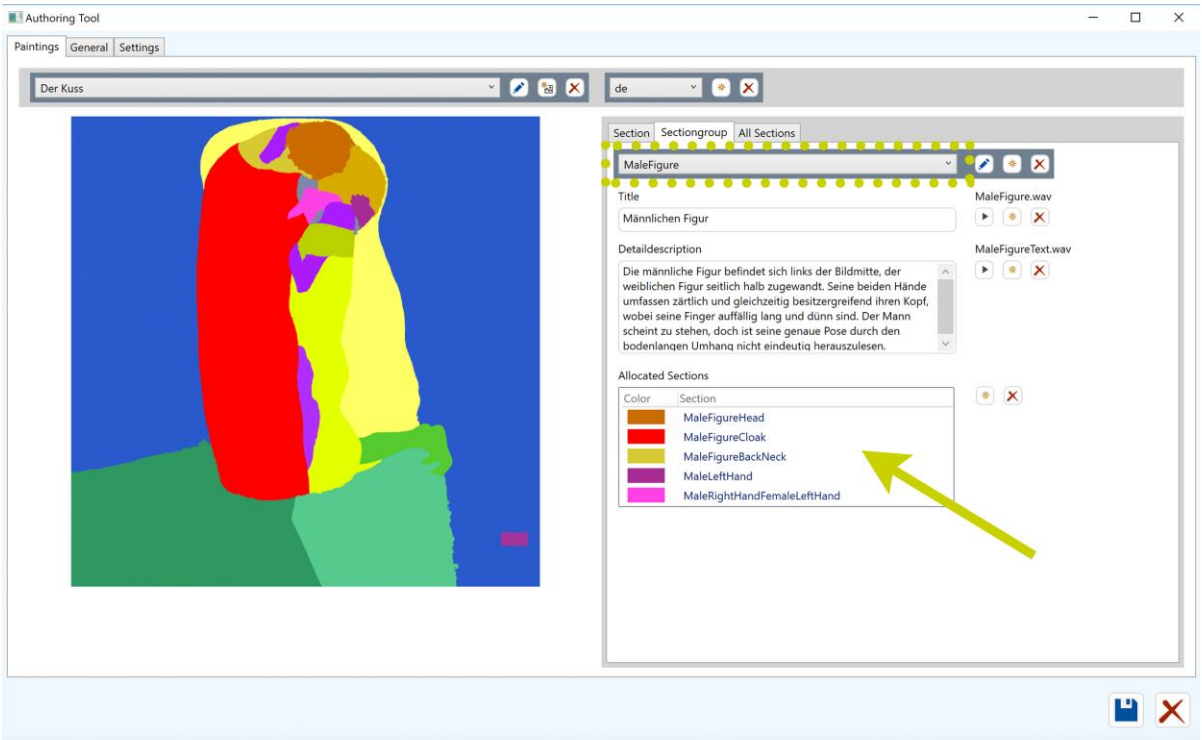

Figure 38: Subdivision of section group.

The selected section is highlighted with the same colour as the section in the label picture. Also the correct section group is displayed. The textual and audio description can be edited in the section tab.

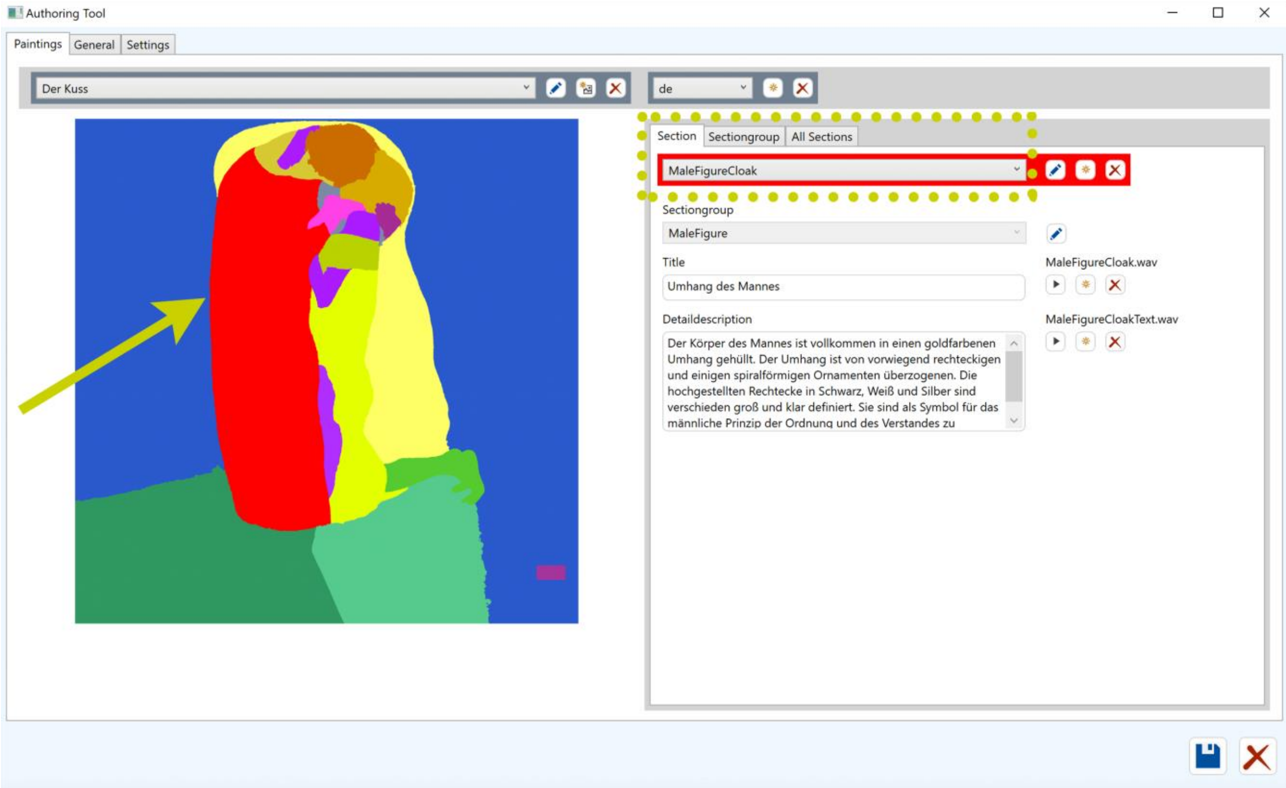

Figure 39: Section “cloak” selected.

Changing the general structure is facilitated by adding, moving or deleting section groups or sections. This structure change is edited for all languages at once. After the paintings have been edited, they will be saved by clicking on the storage icon on the right bottom.

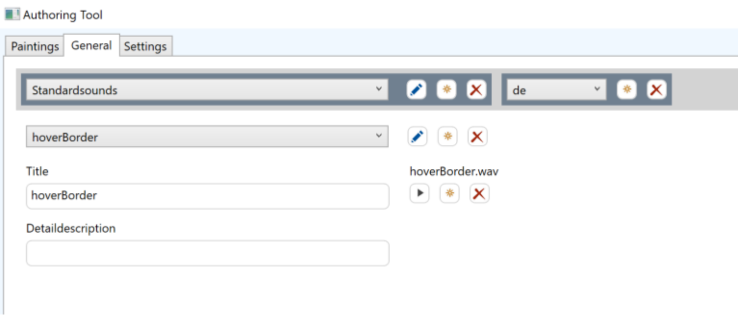

Figure 40: General sound settings.

Similarly, general texts and sounds across all reliefs can be set under the “General” tab, as well as other settings under the “Settings” tab (see Figure 40).

The authoring tool already allows switching between the different paintings and will be expanded in the future in order to allow authoring the new interaction modes yet to be developed.

3.2.5 Projection on the reliefs

The HP Sprout has an in-built projector that is normally used to project on an optional touch mat, in order to create a second touch screen on the desk. As mentioned earlier, this inspired us to investigate, whether it is possible to use this projector to project unto the tactile relief. Normally, the projector is only activated when the touch mat is attached, and automatically turns off, when it is detached. We however managed to circumvent this mechanism, and now can independently turn on and off the projector, set its brightness, and change its calibration settings. The latter was important, since the HP Sprout automatically calibrates on the approximately 30×40 cm large projection space of the touch mat. Having control of the calibration settings, we can maximise the projected space, and could enlarge it at least by 1cm near the monitor, and several centimetres near the user.

This projector can now be used to project onto smaller reliefs (see Figure 41) and even on the user's hands, e.g., to indicate the interacting fingertip. In contrast to reliefs already produced in colour (e.g. [10]), projecting onto a single-coloured relief gives much more freedom. This enables us to not only project the original colour of the painting, but also alternative versions, like high contrast or simplified versions, but also interactive content. This will be investigated in the second phase.

Another challenge was to exactly align the projection with the relief. This could be achieved in a similar way as the automatic relief calibration, but instead of rendering the label image into the depth camera view, the projected image is rendered into the projector optics. For this, however, the projector has to be calibrated, intrinsically, but also the exact offsets to the depth camera need to be known. This whole setup calibration that we developed is described in the next section.

Figure 41: Projection of the original coloured image onto a tactile relief. Note that the projection quite exactly matches the relief, thanks to our new calibration method.

3.2.6 Setup calibration

To make the automatic relief-calibration work, we needed to carefully calibrate the whole setup: The RealSense's colour and depth camera, and the projector. We implemented our own calibration method, based on methods offered by OpenCV [1].

The calibration of the RealSense is based on calibration images taken from a checkerboard pattern, printed and glued on a planar board. For the calibration of the projector, we found a demonstration video [2] and an example implementation [3] that uses a special, second calibration board, with a small calibration pattern on one side in order to compute the board's location, and a white space where the projector projects another calibration pattern on it, in order to find its calibration parameters. In order to simplify this process, we managed to combine the whole process into a single calibration process with a single calibration board. The advantage of our method is that we can use the whole board to print a large calibration pattern, and simultaneously use the same space to project the full calibration pattern of the projector. Using a larger space increases the quality and speeds up the process. Figure 42 shows the calibration setup.

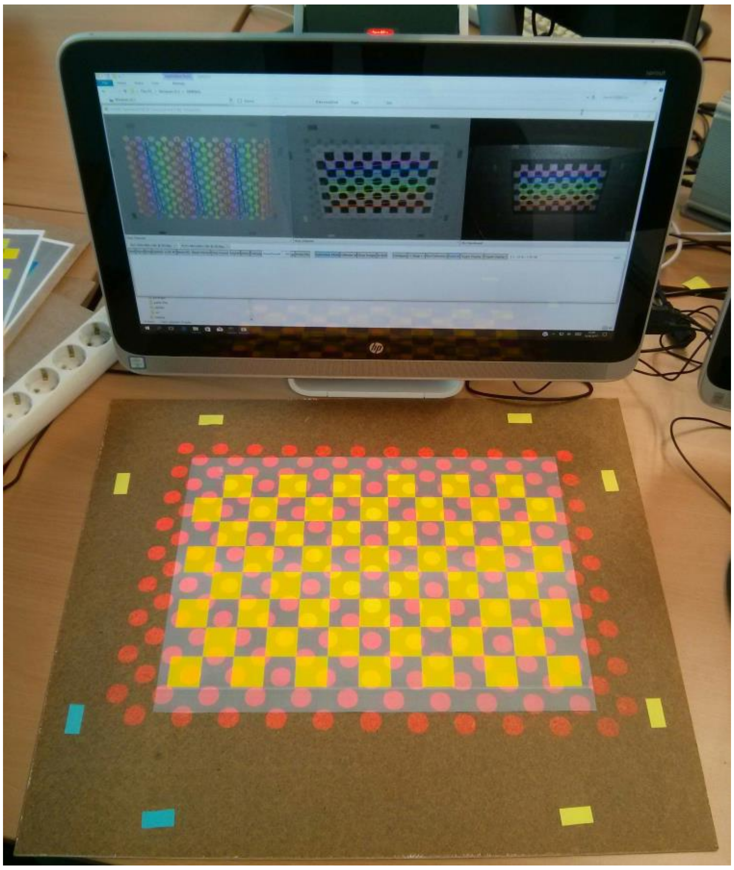

Figure 42: Calibration setup of the HP Sprout using a printed checkerboard and a projected circle pattern.

To make this work, we used the different colour channels of the colour camera and a coloured calibration pattern (see Figure 42): The checkerboard calibration pattern is printed in the colours gray and yellow (we used a colour laser printer), and the projector uses a different pattern made from circles in a deep red colour. Using those colours, the two patterns can be easily distinguished when individual colour channels are used for the calibration (see Figure 43b). In the blue camera channel (left images in Figure 43) the checkerboard pattern of the board can be detected, as the yellow of the calibration board is dark (the yellow pigment filters out any blue colours), while the gray parts still contains all colours, and appears light. The red light of the projector is reflected by both, the yellow and the gray parts, and can therefore be detected in the red camera channel (middle images in Figure 43). We took care, that both parts appear with the same level in the red camera channel, by varying the percentage of black pigment in the gray areas, so that the printed pattern does not disturb the detection of the circle pattern. In addition, the black pigment used for the gray parts absorbs infrared light as used in the depth camera, while the yellow pigments fully reflect infrared light. Thus the checkerboard pattern can also be detected in the infrared channel of the depth camera (right images in Figure 43), and the combined calibration can be performed. We also took care, that the gray part is not too dark in the infrared channel, so that depth measurement can still take place, and the recorded depth values can be further used for a calibration of the depth values. It required some experimentation with different colours and different intensities. Also we had problems with colour compression performed in the colour camera, and had to adapt the colour segmentation to the camera's native YUY2 colour space. With the current implementation and colour selection, calibration works very well.

a)

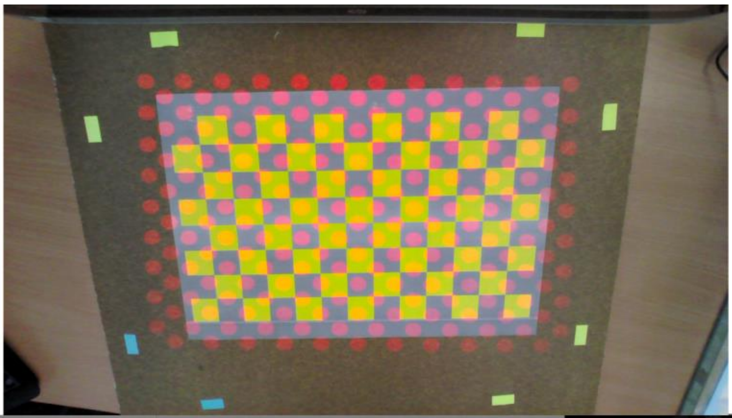

b)

c)

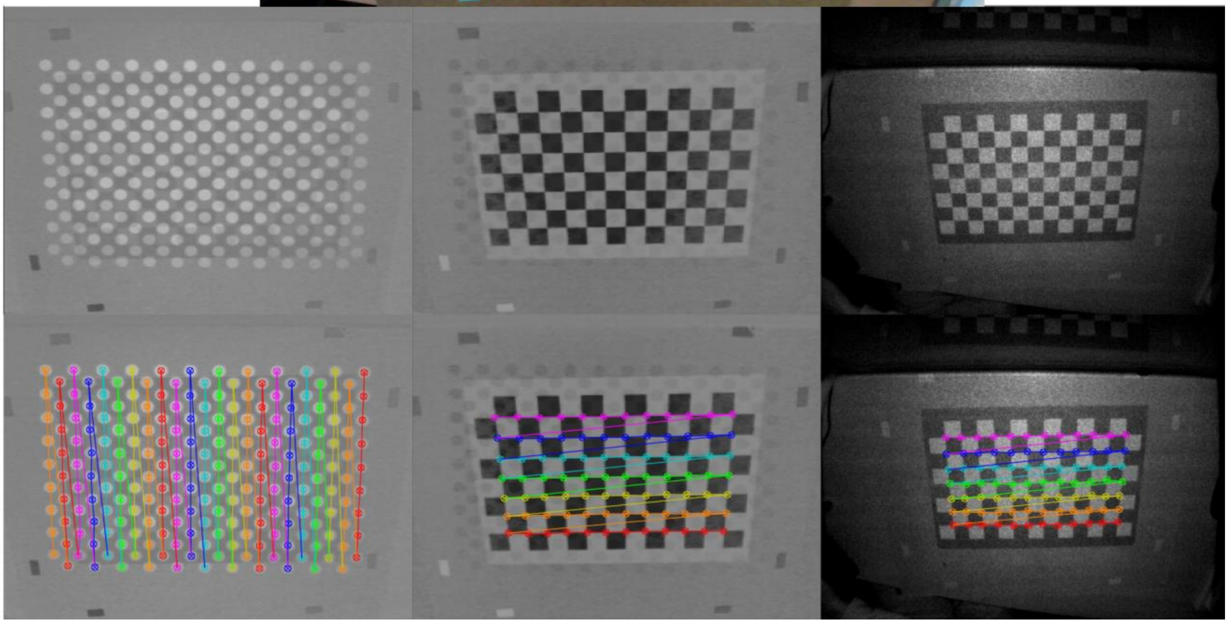

Figure 43: What the camera sees during the calibration process. (a) The RGB image that the colour camera receives. Rows (b) and (c) show the results of the individual colour channels: Left - red channel; Middle – blue channel; Right – infrared channel of the depth camera. Row (c) shows additionally the detected features that are used in the calibration algorithms.

To our knowledge this method is new and was never published. The method is currently summarised as the Bachelor Thesis of Laura Luidolt, who helped with the implementation, and we consider a follow-up publication of this very useful by-product of our research.

4 Task 5.4 “Organisation of multisensory activities and testing”.

As mentioned in the Introduction, this task is concerned with the actual design and production of a test setup that will be installed at each participating museum to be tested and improved. And, as also explained there, only material for the London-based museums – the Victoria and Albert Museum and the Wallace Collection – are being prepared until this deliverable, as the consortium agreed that the choice of object has to be done by the respective participative research groups. Since these have yet to be formed at the other locations, we cannot currently prepare these materials. Production will begin, once the participative research groups have been formed, and the selection process is complete. We report on the performed work, the preliminary results and the results of first testing workshops. Results of Subtask 5.4.2 “Production of relief mediums” are discussed in the simultaneously published Deliverable D5.1.

4.1 Subtask 5.4.1 “Design and production of reliefs for evaluation”

The goal of this subtask is, that each museum has one of their paintings or exhibited objects be translated into a tactile relief, which will then also be used with the Context Sensitive Tactile Audio Guide (see section 3) and will be part of the testing with the participative research groups. We will comment on the progress so far, for each museum separately.

The process is the same for each artwork.

- First the participative research group selects a number of objects they are especially interested in.
- We then write a short statement for each object, how suitable the object may be for relief creation.
- After a discussion, we select one of the objects.
- We write a design document, in which we describe in detail how we see the object and plan to create a tactile relief.
- This document is discussed with the consortium.
- When the design document is approved, we create a first version of our design.
- Screenshots and digital models are discussed with the consortium, and corrections are done until everyone is satisfied.
- We agree with the consortium on the size and material for production.
- The resulting model is sent to a model shop for production.
- The produced relief is then tested with the participative research group.
- In ARCHES we have the budget to create a second version, of the design, in which feedback by the participative research group will be considered.

4.1.1 General considerations

In general, we are limited in size for the HP Sprout, especially in the height of the relief. One limitation is the coverage of the depth camera. In order to facilitate hand-tracking also at the bottom of the relief, it must not be higher than 40 cm. The main limitation is however the projector (see Figure 44), with approximately 40 cm x 30 cm coverage in landscape orientation. Thus, we have to produce at least one relief of each art work that fits into that size. But as it is a really nice feature to have the image or other content projected onto the relief, we believe that we should not miss this opportunity. Alternatively, if the reliefs are really too small with this limitation, we could just build it larger and only have a part of it projected on. Or we could make two relief version, a smaller one for the interactive content, and an additional larger one.

Concerning the material, it would be optimal to have the reliefs made plain white, which would look best for the projections. A very good material we already worked with is Corian®, which is very durable, smooth and very well cleanable. However we will consult the museums about the best material choice. If multiple copies are needed (e.g. to test the same reliefs at multiple locations) plaster moulds could be made, which are not very expensive, plain white but not as robust.

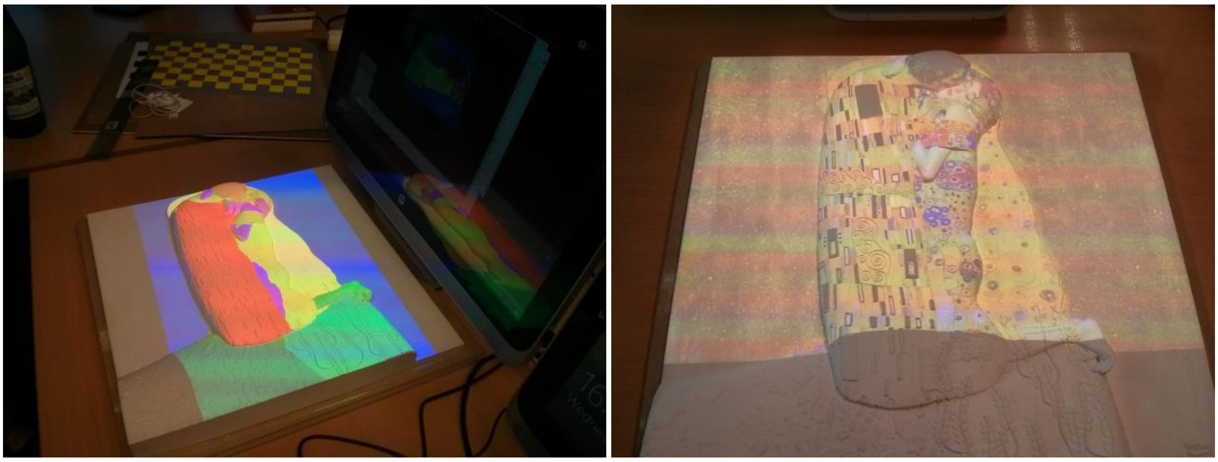

Figure 44: Projecting on the tactile relief. On the left hand side the semantic segments of the painting are projected on the tactile relief of “The Kiss”. On the right hand side the original image is projected on the tactile relief (the relief is unfortunately too big for the limited projection area, as can be clearly seen).

4.1.2 Victoria and Albert Museum

4.1.2.1 Selection of artwork

Five objects from the Europe 1600-1815 Galleries have been chosen by the participants as their favourite objects (see Figure 45). This list contained largely objects, and only a single painting. As our proposed workflow is mainly targeted at paintings, we initially considered the only painting as the object of our choice. However, after visiting the V&A we realised that the V&A is to a large part a museum of decorative objects. Therefore we decided to make an exception and use a three-dimensional object. However, the result should still be in relief form, in order to make it usable with the Interactive Audio Guide (IAG).

Therefore we needed to make a 3D scan of the chosen object. For this purpose we performed first test scans during our first visit in London (2017-03-16). We used a photogrammetric scanning method that could even be used inside the museum during normal opening hours, as it only requires taking a large number of photos from different angles. A software solution (we currently use Agisoft Photoscan Standard Edition) is then used to compute the 3D model from the photos. The selection of the object was therefore largely dependent on whether or not we could actually scan the object as it is exhibited in the museum.

The bed by George Jacob (1780-1785) is placed in a closed niche, fully covered with a thick glass window. All photos we created had lots of reflections, and we could only take photos from a very limited number of angles. Therefore scanning was not successful.

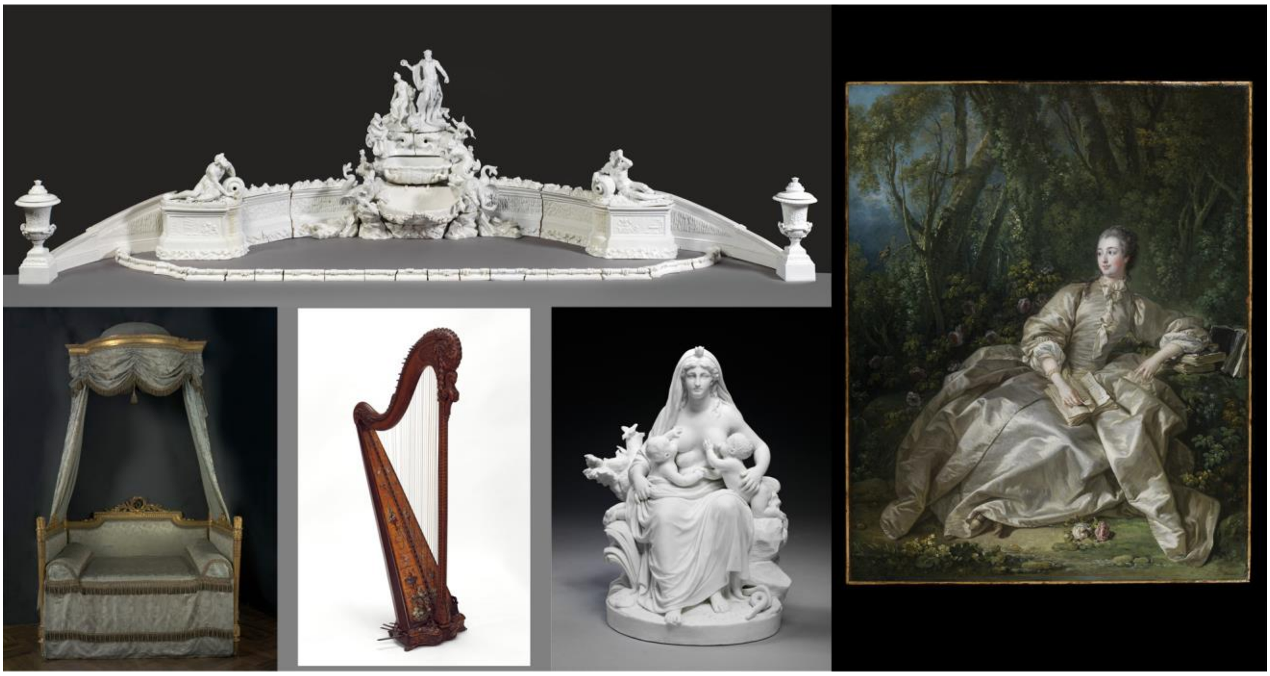

Figure 45: Objects selected by the participative research group from the Victoria and Albert Museum.

The harp by Jean-Henri Naderman (1787) has very beautiful carvings and sculptures which would have been interesting as a tactile model. However, it was placed inside a very dimly lit room with limited access. Therefore scanning was also not successful.

The table fountain “The Triumph of Amphitrite” made by the Meissen porcelain factory (1745-1746) is quite large and consists of many individual porcelain pieces that are arranged on a large and perfectly lit table. We took 631 photos with a handheld Sony RX 100 III. Reconstruction with Agisoft Photoscan was largely successful (see Figure 47), with some parts missing or artificially filled.

The statue “Nature” designed by Louis-Simon Boizot (1794) is a rather small statue, placed in a cuboid glass showcase and perfectly lit from above. We could take 134 good photos from all around the object. The reconstruction with Agisoft Photoscan was very successful, and could yield a quite detailed 3D scan (see Figure 46). The only drawback was that we took the photos with fixed exposure settings, and parts on the back were underexposed, resulting in some large holes. All in all, this was the most successful scan of the first round.

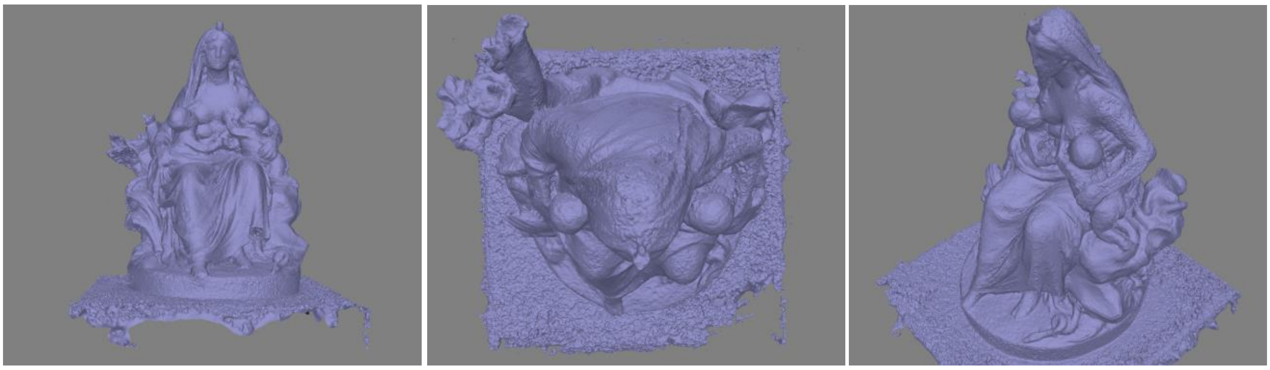

Figure 46: 3D geometry of “Nature” reconstructed with photogrammetry.

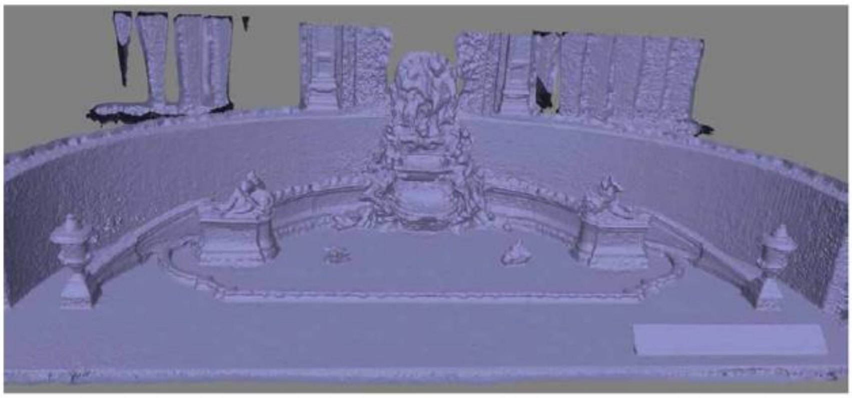

Figure 47: The table fountain reconstructed with photogrammetry.

In a second visit for the plenary meeting (16th May 2017) we took additional photographs with a special permission to shoot outside the opening hours with a tripod and professional equipment with a Canon EOS 5D Mark II to enhance the quality of the data we already had (see Figure 48 and Figure 49). 324 photos of the stature “Nature” and 414 photos of the table found have been taken, each bracketed with three exposures separated by 2 f-stops. These could be merged to HDR photos in order to overcome any over or underexposed parts. In addition another 262 photos of the table fountain have been taken with the handheld camera, in order to cover difficult to reach parts.

Figure 48 Taking photos of Mother Nature.

It was decided to use the table fountain “The Triumph of Amphitrite” as the statue “Nature” loses some of its meaning when not represented in full 3D. As explained above, the tactile audio guide so far only works with reliefs. Therefore there is currently little use for a smaller and touchable 3D version of it. However, as we now have the 3D data, a 3D print could be considered in the future. The table fountain on the other side can be flattened without losing too much information and is generally a very unique object in which people take interest immediately.

Ultimately, we did not have to use the pictures of the last session, as we got access to existing data which was created during the renovation of the fountain and has very good quality. During our research we found a video describing the renovation process. Indeed, many of the parts of the current exhibition have been recreated from 3D scans of existing parts. We contacted the person responsible for the renovation, Reino Liefkes and got permission to use the scans. However, only scans of the individual parts were available to the museum, and the data was too large to be transmitted via internet.

After further research we found a short video clip of the full fountain rendered in 3D, made by Patrick Neubert. We found his address and could contact him. Indeed he was responsible for the video and digitally arranged the individual pieces to resemble the full state of the fountain. He was willing to share the high quality data. Thus we can now design the reliefs based on this dataset.

Figure 49: Taking pictures of "Fountain".

4.1.2.2 Design document

We suggest that the table fountain should be made in some kind of overview relief for the first version. To convert it to 2.5D we would use a slightly angled view from above, so that most features are visible. We would try some skewed projection that allows us to simultaneously show the arch-structure the fountain describes, while still retaining the frontal view of the different statues. In this manner the fountain will be flattened into the limited depth space of the relief, while still retaining the three-dimensional features.

For the second version we can respect the feedback of the participatory research group. We could make a close-up of one of the most important figures, like the central piece, or a collage of multiple close-ups, to allow a more detailed depiction.

4.1.2.3 Current state of the design

Based on the design document, and further input from the consortium, we designed a first version of the relief (see Figure 50 and Figure 51). The Idea was to show the arch-structure of the whole setting while still showing the individual sculptures each from their best recognizable view. We tried to recreate the layout to be largely true to the original arrangement.

The large group in the centre is shown a bit from above, so that people can see/touch into the little basins. The different figures are rendered in 3D, and each of them should be easily separable in depth. We tried to select a view in which the figures do not overlap, so all can be touched.

The two river gods (left and right) are shown slightly from above, mostly from the front, to show the little reliefs on them, and the two male figures lying on it. It should be possible to feel that the base is rounded on the left and right.

The two vases in front are shown completely from the front, so that it is possible to trace their shape without getting irritated by perspective.

The bent walls are flattened as much as possible, to get the figures the most depth available. However, the structure of the walls itself should not be compressed too much.

There are currently two stones missing which are to the left and right in the pond. These were not included in the scans we got. Further, smaller holes are scattered around the surface. These are scanning errors, where the laser scanner did not get the shape, and need to be fixed one by one.

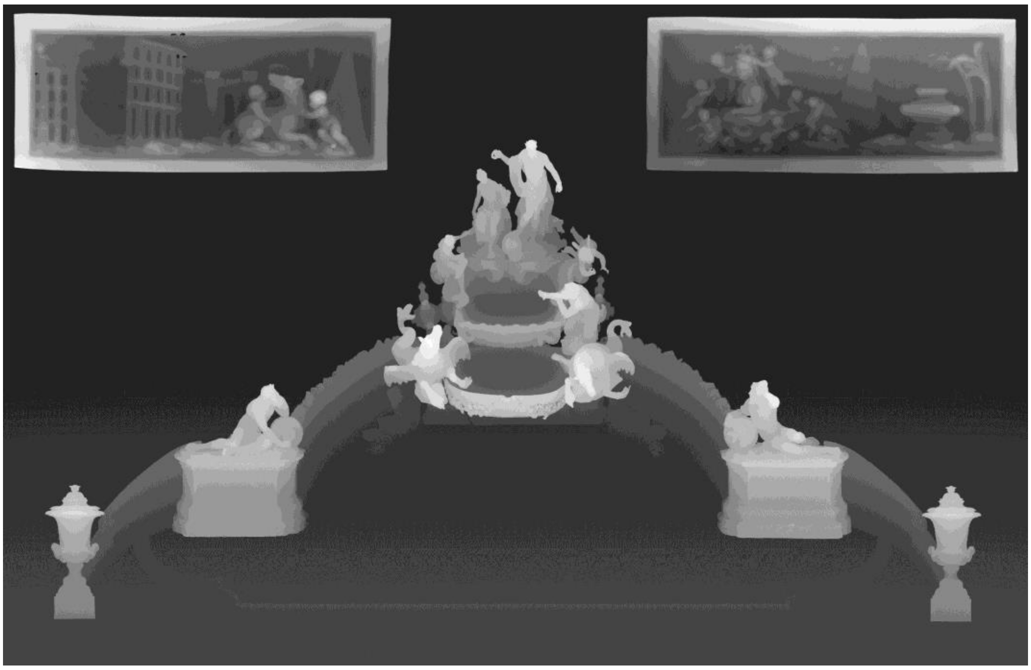

Figure 50: Height field data.

Figure 51: 2.5D relief.

After a first review of the design, we agreed that the reliefs below the river-gods are important and should be made tactile as well. We therefore included a close-up if each in the respective corners above. We tried to match the size of the figures with the full fountain, so that users can get the same amount of detail. The target size of the relief is now approximately 400 mm x 260 mm, and uses the full depth range of 25 mm, so it should be quite plastic.

4.1.3 The Wallace Collection

4.1.3.1 Selection of artworks

The participative research group pre-selected five objects; primarily paintings in this case (see Figure 52). The following paragraphs summarise the arguments during discussions that lead to the final choice:

'Gothic' equestrian armour (1480-1500): This is a full-scale exhibit of a knight on his war horse, both with beautiful crafted armoury. A full 3D replica would be the best choice to get a faithful object representation. A relief is probably not optimal. However there are very thin features, which are hard to fabricate and which would break easily. And, much worse, the reflective material makes it very hard or impossible for a 3D scan, except we could somehow matt it (e.g. by covering it in chalk dust) which was rejected out of conservatory reasons. Therefore, this object was clearly rejected.

The Swing, by Jean-Honoré Fragonard (1767): This painting has a very nice story, a happy mood, with the funny guy who falls just to catch a glimpse under her dress. There are many small details to find and explore and talk about: Stone angels, relief on pillar, the flying shoe, water, dog, rake, ropes, and buildings in the background. However, we felt, that it is simply too much, with all the vegetation and the extreme depth and number of layers: from front tree, girl, sitting servant, stone angels/fountain and background, all in close proximity. The thin ropes in the foreground and the flying shoe are not optimal for relief-depictions. We would need to strongly simplify the painting, which was against our ambition to make it close to the real painting. As much as we liked the painting, we ruled it out, because the relief could get too complex with too many small structures.

With both objects rejected, we were left with three portraits, and it was hard to find a decision.

Figure 52: Objects selected by the participative research group at The Wallace Collection. From left to right: Equestrian armour; Lady with a Fan; Rembrandt Self-Portrait; The Laughing Cavalier; The Swing.

Rembrandt's Self-Portrait in a Black Cap (1637): This painting has the advantage that it is painted by a very famous artist, which is always a plus. But it is very dark in the lower left, which makes it difficult to extract faithful textures in that region. Most depicted surfaces are soft: fur, leather, hair, and the velvet-cap. This is in strong contrast to the hard relief-depiction. Also the contrast to the necklaces could be hard to feel, as there is no distinct change in geometry but more in material properties & reflectivity. Therefore we found it not optimal.

Lady with a Fan, by Diego Velazquez (1640): This is a very dark and sad painting, which is good for setting a mood. Her sad expression and the sad eyes might be a challenge, especially the strangely enlarged left eye. The style is a bit blurry, which makes it harder for us to find distinct edges. The different cloths and accessories make it interesting. Some institutions we previously worked with had issues with depicting breasts on a touch relief, which however was no problem with the WC. Her headscarf and the body shape around her left hand are a bit dark, which makes it hard to see how it was meant to be, which could be hard to recreate. Overall we found it a suitable candidate, with some difficulties.

The Laughing Cavalier, by Frans Hals (1624): It depicts a very realistic face, an interesting pose and clothing, and a simple background. It is probably quite a challenge to get the clothing plausible and accurate. There are so many details. Especially the lace-decorated collar and sleeves might be difficult. The embroidery and buttons are probably not too hard, maybe a bit small, but definitely interesting. Also the distinctive hat is definitely interesting. Overall, we found this painting also a suitable candidate, with some difficulties.

When we had to decide between The Laughing Cavalier and Lady with a Fan, we chose The Laughing Cavalier, not only because it is the most iconic and important painting for the WC, but also because it has a lot of interesting features that will make for a nice tactile image. Additionally, there is a good story to tell surrounding the painting and so there will also be interesting content for the audio guide.

4.1.3.2 Design document

The composition is quite simple, with just a single figure, which is probably easy to recognise, and then there are many details to be discovered.

We agreed to convert the whole picture into relief, as it is hung in the gallery, no close-up. The scan that was made with removed frame will be cropped to exactly the visible part inside the frame.

The painting is in portrait orientation. With the height limit posed by the HP Sprout setup of 30 cm, this results in a total size of approximately 25 cm x 30 cm. With previous reliefs we used the rule of thumb to fit the relief inside Din A3, which would be 29.7 cm x 36.1cm. Thus the realizable size is only 83% of that optimal size. Alternatively, we could cut off unimportant parts at the top and bottom, to make the relief a bit larger within the same height, or just produce it larger and have only the upper portion projected on. The decision will be made when the final design is accepted.

The background will be the bottom plane. As it is not important, it should be quite flat, maybe with the shadow added in a bit lower, so that people can feel that there is a shadow. But it should not distract from the important parts. The text in the top right corner will be made tactile, only slightly raised. It is however quite small with a font height of only 3.5 mm, too small to be readable. However, we still try to edit it in, to have a tactile reference for a possible audio-guide description, which could then read the text when touched and explain its meaning.

The face is an important feature, so we try to recreate it in 3D. Nose, eyes, mouth, ear, hair, eye brows, moustache, beard should be made tactile.

The figure's hat is ‘cavalier hat’ made from beaver skin, which was very popular in the 17th century. We will model the hat as realistic as possible, using pictures and descriptions of similar cavalier hats as reference. It should have a soft and smooth fur-like texture.

With the clothing we will have to be careful so that the different parts can be easily distinguished. Probably we need to leave out the folds for the sake of readability.

There are several different layers of clothing:

- The white undergarment that can be seen through the doublet at some points is linen. As the visible parts are quite narrow, these areas cannot be felt directly as the fingers are too wide. Nevertheless, we try to slightly render them to show that the doublet is cut at these parts.
- The white fabric on his chest is probably the undergarment showing through.
- Over it is the embroidered doublet. The embroidery is important to be raised and tactile. It shows lots of symbols that give the viewer clues as to who the Cavalier might be and the purpose of the portrait. For example, many symbols represent the pleasures and pain of love, so most people believe that this is a marriage portrait. It's important that the different symbols are represented in the relief, but some of the superfluous squiggles and dots could be left out. We have to find a way to make sure the result is not too cluttered.
- The rows of buttons on the chest and on the sleeve will be made as small half-spheres.
- There is some leather-patch over his under-arm. This might be slightly raised, and much smoother than the embroidered doublet.
- The collar and cuffs are lace, which are an important part of the costume as well. The collar will be made a bit higher than the rest, following the shape of the chest and shoulder. The cuffs are probably pulled over the doublet, as it is clearly draped over the leathery part. The lace looks rather smooth on the top, with maybe slight waves, and the texture of the painting made tactile. It looks like it is made from several layers, and quite rough on the ends. We will try to render some of the layers in slightly different heights.
- The black scarf is on top of the collar. We will slightly model the visible folds.
- The black sash he wears (around his waist and over this shoulder and down his back) is smooth and shiny and should feel very different from the rougher, raised embroidery. We will try to slightly hint at the folds, but they should not distract the viewer.

The right arm is largely covered by the rest of his body. We just slightly model it in the lowest levels so you can feel that there is something.

The left arm is coming out to the viewer at the elbow, so we need to model that. The folds of the upper garment at the upper arm seem interesting and important.

The hilt of his rapier (sword) is quite important to include as it signifies that this man is a gentleman. It consists of the round part in his elbow, and the wirework showing through below his underarm. This will be modelled in 3D.

4.1.3.3 Current state of the design

The design of the relief is at the time of this writing in progress. We have already segmented the painting into its relevant parts, and have traced important details. This is the most important step, as it ensures that the relief will stay true to the original painting, and already lays out the edges of the relief. As there is a large emphasis on the small details, like the lace and the embroidery with its symbolic meaning, most elements have been traced by hand. This will ensure that these details will have a good tactile quality in the final relief. Tracing and Segmentation was especially difficult in this painting, as several boundaries were not explicitly drawn, or were overly dark in shadowed regions, and a huge amount of detail had to be traced. The current state is depicted in Figure 53.

Figure 53: Segmentation of the painting “The Laughing Cavalier”. Semantic labels shown in false colour.

As the second step, we did a design research, in order to get a better intuition about the depicted items. We searched for photos of similar objects, and clothing from the respective epoch (see Figure 54 to Figure 57 for some examples). This will help to better infer the 3D shape as the painting is in some cases somewhat ambiguous. This clarified, e.g., how the hat bends in 3D, that the doublet has cuts through which the undergarment can be seen, and how the handle of the rapier might be formed.

Figure 54: Cavalier’s hat made from leather.

Figure 55: Lace ruff and cuffs, as worn in the early 16th century.

Figure 56: Doublet with decorative slits.

Figure 57: Rapier.

Further we designed a first draft of a relief that depicts the embroidered details on the sleeve (see Figure 58). This could be manufactured as an additional relief that allows exploring the details. The relief of the full painting is currently work in progress.

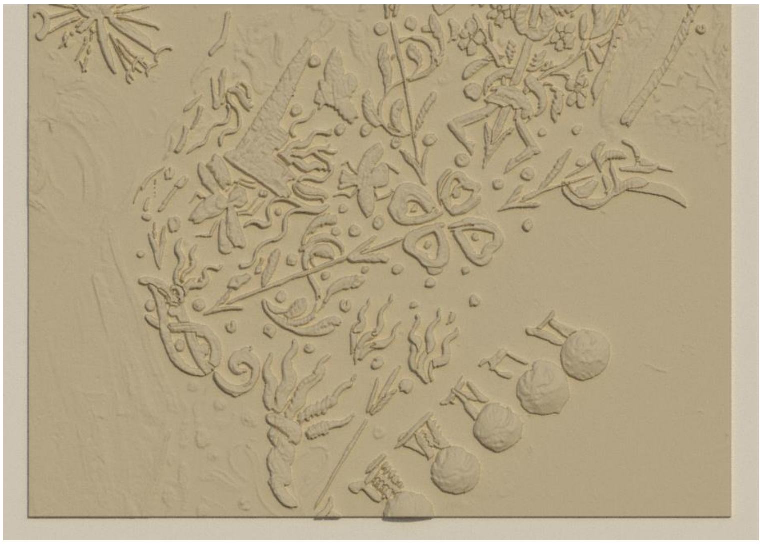

Figure 58: Relief showing the embroidered details on the sleeve of the "Laughing Cavalier".

4.2 Subtask 5.4.3 "Creation of tactile audio guide for produced reliefs"

4.2.1 Hardware

Before we decided to use the HP Sprout for the test setups in the project, we could use a unit on loan. After successfully porting our software to the new platform, we were convinced that it is the currently best option, and looked into ways to acquire the required number of devices (one for each museum, and one for VRVis, according to the Grant Agreement). Unfortunately, at that time, neither the original HP Sprout nor the updated version HP Sprout Pro was available for purchase any more. We went into negotiation with HP and found out, that a new Version Sprout Pro G2 is planned, but no delivery date could be given. For the time, we could continue our work on the loan unit. Once we were asked for testing sessions in London, we were lucky to find a used HP Sprout model at a UK-based online shop, and could deliver it directly to the V&A museum. For the time being, we decided that a single model is sufficient for the two London museums, as they share the same participative research group.

During two stays in London (one for testing purposes and one for the project meeting), the device was set-up, calibrated, and the local team was introduced to the technical basics to set-up the device. We already successfully demonstrated the possibility for remote maintenance and support during test sessions (via a Remote Desktop connection). To facilitate testing even before specific content was created, we were allowed to produce a plaster copy of an existing tactile relief and could use the corresponding audio content.

We continued our negotiations with HP and were offered a preview version of the new HP Sprout Pro G2. However, the price was significantly higher than for the original version, and was way beyond the granted budget. We also could get approximate specifications, and found out that

- the depth sensor was changed to an ORBBEC Astra device, with unknown quality in the near range, and which would have required a software rewrite of our tracking solution,
- the projector was exchanged for a full HD model, but the projection space was narrowed and widened to approximately a 16:9 aspect ratio, which is less optimal for portrait reliefs, and
- generally the effective working volume was narrowed down, which would have negative effects on tracking performance, especially towards the user.

These drawbacks convinced us to stay with the original model. After the announcement of the new version by HP, sales offers for the original version re-appeared, and we quickly acquired the remaining 6 devices, to a very competitive price, and even got the updated HP Sprout Pro versions.

All devices have been tested for production defects. In one device, a defective projection unit needed replacement, as the desk light did not work. Further we received replacements for two touch mats as these had scratch marks on their surface.

Currently all models are setup with our software, readily calibrated and will be stored at VRVis for easier software maintenance until needed at their final destinations (i.e. when the museums have set up their participative research groups). This aims to reduce unnecessary travel costs, as the setup can be tested with the produced reliefs in house, until everything works perfectly.

4.2.2 Guidelines for audio guide content

As the content will not be authored at VRVis, we wrote guidelines for the texts and also included some ideas for future content. The guidelines are as follows:

We can offer up to five general texts, which are for the whole painting and not for a specific part.

One text should be a short description of the painting or object as a whole, for blind or visually impaired people as a first orientation. It should contain basic information: how large is it, what is in it, composition, colours, style, is there anything very specific for the object? This text serves as a start to get the visitors exploring. More details can then be offered as separate texts, see below. You can of course define your own stories, whatever you think fits to the piece. Just keep them moderate in length, around two minutes maximum. Please give a headline, and then the text.

In the future, we might also combine them with multimedia, sign language, projections, and so on. If you already have some ideas, please let us know, so we can plan future features. But for this version, only text is required: audio and subtitles.

In addition, we can offer texts for different parts of the object that can be triggered when being touched. There can be as many parts as we like, but they should not be too small, as they would be difficult to find and touch. It would help us, if you could scribble over an image or the painting or a photo of the object where these locations are. VRVis will then draw the exact map, once the final reliefs are designed. For each of these areas we need:

- A header or the name of this part in just a few precise words. This is used for orienting in the painting, when quickly scanning over it.
- A description of that region, up to 30 seconds in length. Talk about things specific to that area, what is it, how is it shaped, are there interesting colours or textures? It also helps to reference nearby places, e.g. at the upper arm, say that the elbow is to the lower right, or that it goes under a part of the dress. Things like that to help a blind person orient themselves inside the painting.

Finally, this is a list of some further options, we could implement in a future version. When you think about the texts, maybe you find something that would be interesting to describe in one of these ways. Please keep notes already that might inspire us to different interactions. We can discuss them later for a future version.

- Hierarchical descriptions. At our previous relief we had one description for the whole male and female figures, describing their arrangement. And then several sub-parts of them (head, hand, arm, dress, feet...) which are accessible once the person listened to the full part.
- We could create special modes to describe the colour and texture, so there are different layers of descriptions.
- We could create multiple shorter texts which cycle through each time you touch the same part.
- Multi-touch would also be an option. If there are interesting relations between some parts, a separate description could be triggered when both parts are touched with both hands.
- Each part or the painting as a whole could have some background sound or music played. Currently we have a fixed sound playing, when correctly touching a region. We could use different sounds, so people instantly hear when they touch the same part. This would need to be a loop-able sound.
- The audio comments could be enriched with music and sound effects, especially the longer off-relief descriptions.

We can also have different languages. It is also possible to have a simplified version as a separate "language".

4.2.3 Audio guide content

Based on the recommendations in the previous section, the following audio guide texts have been written.

4.2.3.1 The Fountain (V&A)

Starting Point: If you touch anywhere on the relief, this text will start.

- Table Fountain made from porcelain, 1745-1747
- Please choose between normal English, and simplified English. You can also choose to hear an overall description of the work, and audio descriptions for the main figures, Neptune and Amphitrite.
- This explanation will start with a short introduction. Then you can hear other texts by touching different areas of the relief.

Figure 59: Sections of the fountain relief represented in the audio guide.

4.2.3.1.1 Normal English

Introduction

About two hundred and seventy years ago, Heinrich von Brühl, a German Count, very powerful and rich, bought a palace on the edge of his hometown of Dresden. He reformed and enlarged the building, and in the garden he put a gigantic fountain, showing ancient gods and fabulous creatures. Later, the Count decided that he also wanted a smaller version of the fountain. He wanted it to be out of porcelain so it could be installed on a table, to surprise his guests at an important dinner party.

The table fountain was made by Johann Joachim Kaendler, the "chief modeller", and the most famous of the sculptors of the Meissen porcelain manufacture, with the help of several assistants. It was completed in early 1747, just in time for a very special occasion, the celebration of the royal marriage between the crown prince of France and a princess from Dresden. The porcelain fountain was the central piece of the table decorations during the dessert course, and was much admired by the international guests. Not only for its beauty, but also for the fact that it was actually running with scented rose water.

Overall Description of the Fountain [optional]

Just like the stone fountain in the Count's garden, the table fountain has a central water basin, which is contained by a low brim, towards us, and a high wall, behind. All the figures are mounted on the rear wall. The central group of these figures is about half a meter high, and shows Neptune, the god of the sea, standing at the very top. Beside him, seated, is his wife Amphitrite, the queen of the sea. At their feet are two water basins and other figures including two sea horses. Following the rear wall to the left and to the right, we will find two massive stone platforms, with a lying male figure on top of each platform. The one on the left represents the river Tiber, the one on the right, the river Nile. On the platform below them, the rivers are also represented in a relief, one for Rome and one for Egypt. From here on, the rear wall becomes lower and lower until we reach the end of the fountain, which on each side is marked by a vase.

Parts

- **Neptune:** [Description – optional] Neptune stands on an elevated platform, a shell-chariot. He has one foot resting on a dolphin's head. His left arm is stretched forward, and he is pointing downwards, while his right arm is lifted, holding a laurel wreath. He wears a crown and a beard and, as an ancient god, he is only dressed with a piece of cloth falling from his shoulder down to his hip. [End of Description] Neptune was the god of freshwater and the sea in Roman religion. He is the counterpart of the Greek god called Poseidon. Neptune arrives on a chariot made of a large shell, drawn by two sea horses. He is accompanied by his bride Amphitrite. At each side, we see the river gods representing the Nile and Tiber. It is thought that Neptune is about to present his laurel wreath to honour Prince-Elector Augustus III of Saxony. In this way, Count von Brühl intended to invoke a comparison between his ruler and the first Roman emperor, also called Augustus, who lived around the birth of Christ. Augustus reigned over the Roman Empire, and also over Egypt, represented by the rivers Tiber and Nile. He initiated an era of peace known as the 'Pax Romana', the Roman Peace. With this reference to the great emperor of the same name, Count von Brühl wanted his people to see Augustus III of Saxony as a great ruler as well.
- **Amphitrite:** [Description – optional] Neptune's wife Amphitrite is sitting next to him. She, too, is barely dressed, so we can see her breasts and one leg fully exposed. She is young and beautiful, and looking slightly downwards. [End of Description] Amphitrite is the queen of the sea. When the porcelain fountain was used at the royal wedding feast, for the future king of France and a Saxon Princess, the guests must have interpreted that Amphitrite represented the bride. This way, the meaning of the whole scene changed. In fact, the piece was at that time referred to as 'the Triumph of Amphitrite', as people saw a comparison between Neptune's bride and Princess Maria Josepha of Saxony, about to get married to the crown prince Louis, the son of king Louis XV of France.
- **Sea nymph:** Below Amphitrite, we find another sitting female figure looking downward, a sea nymph. She holds a shell in one hand and a piece of coral in the other. One of her feet has broken off and has not been restored.
- **Putto:** Below Neptune, we find a small figure with wings. It is probably a putto, which is an Italian name for a figure depicted as a chubby male child. In the Baroque period, the putto came to represent the omnipresence of God and is usually shown naked and sometimes winged.
- **Triton:** Below Neptune and the small putto figure, there is Triton, the son of Neptune and Amphitrite. Triton is a merman. His upper body is that of a human and the lower body that of a fish. He is blowing into a twisted conch shell, which he uses to calm or raise the waves.
- **Upper water basin:** The upper water basin looks like a shell. It is rough on the outside and has a

curvy brim. The inside has a smooth surface, as if it was made from mother-of-pearl.

- **Left wheel:** On each side of the upper water basin is a wheel of the shell-chariot. It is not a normal wheel, but rather a paddle wheel, so it can drive the chariot on and under the water. This is the paddle wheel on the left side, the other can be found behind the figure of Triton, on the other side of the water basin.
- **Right wheel:** Behind the figure of Triton there is a water-wheel of the shell-wagon. As is the case with the wheel on the left side of the water basin, just under the two female figures, this is not a normal wheel, but a paddle wheel that can drive the chariot on the water.
- **Left sea-horse:** Just below the paddle wheel, we can see a seahorse, a creature that is half horse, half fish. It looks up towards the group of Neptune and Amphitrite. It also has a pair of wings on the shoulders, and its front hoofs look like webbed feet. The lower part of the body is covered with fish scales and a big fin.
- **Right sea-horse:** On the right side of the lower water basin, we find one of two strange creatures. They are seahorses, half horse, half fish, who pull the chariot over (and under) the water. The head of the winged sea horse looks down, yet its tail fin is up in the air. Just like the other sea-horse, on the left of the chariot, it also has a pair of wings on the shoulders, and its front hoofs look like webbed feet.
- **Lower water basin:** The lower of the two water basins has a rocky outside, decorated with two swags of sea shells, and a smooth surface on the inside. It is flanked on each side by a sea horse.
- **Rear wall:** Just like the Neptune fountain in the garden of the Palace, the table fountain has a high rear wall where the figures are mounted. The artist, Johann Joachim Kaendler, had modelled all the figures during the first months of the project, but then left the work to his assistants. He again turned his attention to the centrepiece a year later, in November and December 1746, when he modelled the rear wall with icicles.
- **Left river god:** The two lying river gods Tiber and Nile look very similar. On this side, we have Tiber, who wears a laurel wreath.
- **Left stone base:** Below the river god Tiber, on the massive stone plinth, there is a relief of Rome. You can find it in an enlarged version on the upper left corner of our tactile relief.
- **Left relief (reproduction):** This is an enlarged version of the relief of Rome, which can be found on the stone base under Tiber. The scene of Ancient Rome shows the famous Coliseum on the left, and a wolf on the right. This is the mythological she-wolf who breast-fed the little boys Romolus and Remus, the founders of Rome.
- **Right river god:** This river god represents the Nile. His head is partly covered by a cloth, which indicates that the source of the river Nile was unknown at the time. His torso is naked and he is leaning on an amphora that pours water into the central basin.
- **Right stone base:** Below the river god Nile, on the massive stone plinth, there is a relief of Egypt. You can find it in an enlarged version on the upper right corner of our tactile relief.
- **Right relief:** This is an enlarged version of the relief of Ancient Egypt, which can be found on the stone base under the river god Nile. This relief shows the Sphinx on the left, crowned by a large number of children figures who resemble small angel figures from baroque altars, called Putti. To

the right we see a pyramid, a tomb and a palm tree.

- **Left vase:** In December, with the wedding date approaching, work on the vases ran a little late. The main artist, Johann Joachim Kaendler corrected one of the vases modelled with a relief of Bacchus, the god of wine. In the same month, one of the assistants, Johann Gottlieb Ehder finished the relief of one of them, while another assistant, Peter Reinicke, who had not been involved so far, modelled the relief for another.

- **Right vase:** After Count von Brühl's death, the table fountain entered in the possession of Count Marcolini, who had taken over many of Brühl's official functions, including that of director of the Meissen porcelain factory. Marcolini restored the porcelain fountain, replacing all the damaged parts with new ones made at the Meissen factory, using the original moulds. The two vases were among these restored parts. After the restoration, Count Marcolini installed the fountain in his summer palace as a water feature, which included a goldfish basin.

4.2.3.1.2 Simplified English

Introduction

This is a model of a fountain. This model belonged to a man called Heinrich von Brühl. Heinrich von Brühl had a summer palace. Heinrich’s summer palace was near a town called Dresden. Heinrich’s summer palace was a small palace with a very big garden.

(FX: BIRDSONG AND WIND IN TREES, WITH A FOUNTAIN RUNNING)

In Heinrich’s big garden was a real fountain made of stone. Heinrich had this model made of the fountain in his garden. Of course, the fountain in the garden was much bigger than the model.

Heinrich had this model of a fountain put on the table during a very special dinner party. Heinrich also had it filled with pink coloured Rose water. Imagine seeing this beautiful white model on the table when you are having dinner. It would have looked amazing.

(FX: GLASSES AND VOICES, PERHAPS MUSIC...AS AT A PARTY)

This model was made more than 250 years ago. The model is made of porcelain. The model was made by a man Johann Joachim Kaendler. Johann was famous for making Meissen pottery. The fountain made by Johann is full of figures.

Parts

- **Neptune**: In the middle, you can see Neptune. In many stories, Neptune is the god of the sea. (FX: SOUND OF THE WAVES)

- **Amphitrite**: Beside Neptune is his wife, Amphitrite. She is young and beautiful. Amphitrite is the Greek goddess of the sea. (FX: SOUND OF THE WAVES)

- **Sea nymph**: This sitting figure is a sea nymph. She lives under the sea.

- **Putto**: Look at this small boy with wings. He is called a putto. This is an Italian name for a naked little angel with wings.

- **Triton**: This is Triton, the son of Neptune and Amphitrite. Triton is a merman. His upper body is that of a man. His lower body looks like a fish. He is blowing into a conch shell. (FX: SOUND OF SOMEBODY BLOWING A CONCH)

- **Upper water basin**: The upper water basin looks like a shell. It is rough on the outside and smooth on the inside.
- **Left wheel**: On each side of the upper water basin is a wheel. It is not a normal wheel, but rather a paddle wheel, so it can drive the chariot on the water. (FX: SOUND OF a paddle wheel)
- **Right wheel**: Behind the figure of Triton there is a water-wheel. This wheel can drive Neptune’s chariot on the water. (FX: SOUND OF a paddle wheel)
- **Left sea-horse**: Look at this creature. It is a seahorse. The seahorse is half horse, half fish. It also has wings on the shoulders. (FX: SOUND OF a horse, and water)
- **Right sea-horse**: On the both sides of the lower water basin, we find strange creatures. They are seahorses. Seahorses look like horses, but have a fish tail. They also have wings. (FX: SOUND OF a horse, and water)
- **Lower water basin**: This is the lower water basin. It has a rocky outside and a smooth surface on the inside.
- **Rear wall**: This is the back wall of the fountain. It is covered with icicles.
- **Left river god**: This model is meant to make you think about rivers as well. Can you see two men lying down on pillars? The two men lying down on the pillars are to make you think about the River Nile and the River Tiber. On this side, we have Tiber, who wears a laurel wreath.
- **Left stone base**: Below the river god Tiber, on the pillar, there is a relief showing the city of Rome. You can find it in a larger version of it in the upper left corner of our tactile relief.
- **Left relief (reproduction)**: This is the large version of the relief of Rome. You can find it on the pillar below the river god Tiber. It shows the city of Rome, with its famous Coliseum on the left, and a wolf on the right. This wolf is actually a she-wolf. The story goes that she breast-fed two little boys. And that the boys later founded the city of Rome.
- **Right river God**: This lying figure is a river god. He represents the river Nile, in Egypt.
- **Right stone base**: Below the river god Nile, on the pillar, there is a relief of Egypt. You can find it in a large version on the upper right corner of our tactile relief.
- **Right relief**: This is a large version of the relief of Ancient Egypt. It can be found on the stone pillar under the river god Nile. This relief shows the statue of a woman. She is called the Sphinx. She is surrounded by a large number of children figures. To the right we see a pyramid, a tomb and a palm tree.
- **Left vase**: On the left end of the fountain, we find a big vase.
- **Right vase**: On the right end of the fountain, we find a big vase. Can you see the other vase, just on the opposite side?

Extra chapters

- **The wedding party:** In early 1747, Maria Josepha of Saxony was married by contract to the future king of France. The preparation of this important union was maybe Count Von Brühl’s greatest diplomatic achievement. The groom, Prince Louis, was the son of the French King, Louis XV, was not present at the wedding party in Dresden, in January. He was still grieving for his first wife who had died half a year earlier, in Paris. According to the tradition, the 17-year old was married to his two years younger spouse in absence, by contract. The real wedding party took place in Paris, a month later. The first wedding ceremony, a splendid supper, took part in von Brühl’s palace in Dresden and featured the table fountain as the centre piece of the main table.

- **The Meissen manufacture:** Among his many titles, Count von Brühl was also superintendent of the famous Meissen porcelain workshop. In fact, he used his position to order many fine works from the Meissen manufacture, the most famous one a service of tableware with more than a thousand pieces. Porcelain can be formed by hand or moulded into forms, and is then fired at very high temperature, usually around 1400 degrees Celsius. It was first made in China around the 7th or 8th century. Meissen porcelain or Meissen china is the first European hard-paste porcelain. The production of porcelain at Meissen, near Dresden, started in 1710 and attracted artists and artisans to establish one of the most famous porcelain manufacturers, still in business today. Its signature logo, two crossed swords, was introduced in 1720 to protect its production. The mark of the crossed swords is one of the oldest trademarks in existence.

- **The Restoration:** By the time the fountain was acquired by the Museum, it was incomplete and much damaged. Only the main figures were on display while the remainder, including many fragmentary pieces, remained in store. Until 2011 when the V&A Museum started a research and restoration programme where many of the damaged parts were carefully restored by the Museum's conservators. They have also used scanning technology to create virtual, 3D models of the missing parts of the fountain. These scans are based on a second, later version of the fountain, which was identified in the stores of a museum in Dresden. The conservators used cutting-edge 3D printing and machine-tooling technology to replace the lost parts. For the tactile model that you have in your hands, we used very similar techniques. For this model, we simplified some details to make them easier to touch and to understand.

4.2.3.2 The Laughing Cavalier (WC)

Starting Point: If you touch anywhere on the relief, this text will start:

- The Laughing Cavalier, by Frans Hals (1582/3 - 1666)
- Netherlands, 1624
- Painting , Oil on canvas
- Size: 112.5 cm x 98 cm x 9 cm
- Inscription: 'AETA SUAE 26 / A° 1624'

4.2.3.2.1 Normal English

Introduction: This portrait shows a young man at the age of 26, and was painted by the Dutch artist Frans Hals, nearly 400 years ago. It has been called The Laughing Cavalier, even though the man is neither a noble man, nor is he laughing. The painting is about 85 high and 70 cm wide, , and was painted on canvas. We can see the man, who looks directly at us from his hip upward to his head. He wears a hat and very lavish clothes, in the style of a rich Dutch merchant of his time. This work is probably the artist’s most famous portrait and demonstrates his virtuosity in the lively characterisation of his sitters.

Inscription: In the upper right corner of the painting, we find an inscription. It says 'AETA SUAE 26, and below A 1624'. This is Latin and means that the man was painted when he was 26 years of age, in the year

1624. The inscription is painted in black, on a grey background. We find similar inscriptions in other portraits by Frans Hals.

Figure 60: Sections of "The Laughing Cavalier" represented in the audio guide.

Hat: In the process of painting, the artist seems to have sat close to his subject, depicting him from below. This is why we see the underside of the black felt hat, which makes the brim look huge. Furthermore, the low viewpoint gives the man an elevated position. He seems to be in control, looking down on us.

Face: The man's face and head are executed from fine, blended brushstrokes. The young man wears an upturned moustache and a goatee. His eyebrows are bushy and his hair is a little curly. He has twinkling, green-brown eyes that look down on us in a rather arrogant pose. His lips are closed as he smiles but does not laugh, despite the title The Laughing Cavalier. The identity of the sitter has not yet been firmly established. The painting has inspired a novel with the same title by Baroness Orczy in 1913, and a musical by Arkell and Byrne, in 1937.

Ruff: We cannot see the man's throat or neck, as he wears a big ruff. This kind of a large round white collar was very fashionable at his time, but it must have been rather uncomfortable to wear. The frill is made from layers of starched textiles. Just under the chin, a small piece of black cloth is mounted on the ruff, as if it was tie. This black tie combines with the black hat and a black sash around his shoulder and waist.

Costume: The man wears a dark doublet, a close-fitting padded jacket, with buttons and embroideries in vivid colours. By studying emblem books of the time, researchers have been able to identify some of the motifs as symbols for love. These symbols for the pleasures and pains of love include bees, as they make sweet honey but they also sting. We can also see love arrows, and lovers' knots, as well as flaming cornucopias, which are goats' horns overflowing with flowers, fruit, and corn. As allusions to gallantry and courtship, they may indicate that the work was painted as a betrothal or wedding portrait. In this case it would be normal to find a companion piece of the bride. However, such a painting has not been identified yet.

Sash: The man wears an exquisite black sash around his shoulder and waist. It is painted with broad, loose brushstrokes. This black sash goes well with the black hat and the rest of his costume. Hals mixes white into the black to give a sense of the folds and patterns in the material.

Arm and elbow: The young man poses with his arm akimbo, with hands on the hips and his elbows turned outwards. This allows us to study the heavily embroidered sleeves of his dazzling costume. The stitching in white, gold and red thread may offer some important clues to the identity of the man, and why he had his portrait painted this way. We see symbols of love, and also a special motif called a caduceus. This is a herald's wand, with two serpents twined round it, and two wings at the top. It is normally carried by the Greek messenger god Hermes. Hermes is the god of trade and commerce, which might indicate that the man on the portrait is a wealthy textile merchant.

Pommel: At the crook of his elbow, clamped between his arm and his body, we see the rounded knob, a pommel, which indicates the end of the handle of a small sword, a rapier.

Lace Cuffs: The cuffs are made of white lace. It shows complicated geometric forms and is so oversized that we cannot see the hand beneath it.

4.2.3.2.2 Simplified English

The Laughing Cavalier, painted by Frans Hals. Frans Hals lived in the Netherlands. He was born in 1582, or maybe a year later. We don't know it for sure. He became 84 years old and died in 1666.

This painting was made in 1624. At that time, Frans was 42.

This is an oil painting on canvas.

The size with the wooden frame is 112.5 cm high, by 98 cm wide, by 9 cm deep

The picture has in Latin inscription that says: 'AETA SUAE 26 / A 1624'. We will find out what that means soon.

Introduction: This artist Frans Hals painted this picture nearly 400 years ago. It is a portrait of a young man. We do not know the man's name. So the famous painting has no title. However, somebody once called it The Laughing Cavalier. This is strange, because the man is not a noble man. And he is not laughing. But still, this is what the painting is called today. The painting is about 85 high and 70 cm wide. It was painted on canvas. It has a dark wooden frame around it. The young man in the picture looks directly at us. He is shown from his hip upward to his head. He wears a hat and very expensive clothes. This is probably the most famous portrait by the painter Frans Hals. It shows that he is a great artist.

Inscription: In the upper right corner of the painting, we find an inscription. It says 'AETA SUAE 26, and below A 1624'. This is Latin and means: This man was painted when he was 26 years old. This was in the year 1624. The inscription is painted in black, on a grey background.

Hat: When the artist painted the young man, he was sitting very close to him. Maybe the young man was sitting on a high chair, and the artist on a low chair. This is why we see the man from below. He looks down on us. This makes him seem a bit arrogant, doesn't it? And we see the underside of his hat. It looks very big this way. What is the hat made of? Probably of black felt.

Face: The man's face and head are painted with a small paintbrush. It must have taken some time to get the face right. It looks very real. We can see that the man wears a moustache. The moustache is turned upwards. And you see that he also wears a small beard on the chin? This kind of beard is called a goatee. His eyebrows are bushy and his hair is a little curly. The young man has twinkling, green and brown eyes. He looks down on us in a rather arrogant pose. His lips are closed. He smiles, but he does not laugh.

Ruff: We cannot see the man's throat or neck. He wears a big white collar, called a ruff. To wear a ruff was very fashionable at this time. It does not look very comfortable to wear, though. It is made from layers of stiff textiles. Just under the man's chin, we see a small piece of black cloth. It is mounted on the ruff, as if it was tie. The black tie goes well with the black hat and the rest of his clothes.

Costume: The young man wears a dark jacket. It has many buttons and embroideries in vivid colours. The embroideries have been stitched on the jacket. Some of the things we see stand for love. Love can be sweet like honey, but it can also hurt. This is why we see bees as symbols for love. They make sweet honey but they also sting. We can also see love arrows, and lovers' knots. Finally, we see a goat's horn full of flowers, fruit, and corn. It has little flames, too. This is called a cornucopia. It stands for gallantry and courtship. That is, for the desire of a man to win the heart of a woman. Maybe this portrait shows a man who is in love? A man who is going to get married, maybe? In this case it would be normal to find a companion piece of the bride. However, such a painting has not been identified yet.

Sash: The man wears a black sash around his shoulder and waist. A sash is like a ribbon, or belt made from a piece of cloth. The sash is painted with broad, loose brushstrokes. This black sash goes well with the black hat and the rest of his costume. The painter, Frans Hals mixes white into the black colour. This way he can give a sense of the folds and patterns in the material.

Arm and elbow: The young man poses with his arm akimbo. This means that he has his hands on the hips. His elbows are turned outwards. This pose of the young man allows us to look closely at his sleeve. The sleeve of his dazzling costume has a lot of embroideries on it. The stitching is made with white, gold and red thread. Even if we do not know who the man is, these embroideries might help us to find out. First, we know that he must be rich. A poor man could not buy such an expensive jacket. Second, we see symbols that stand for love. For example, arrows and knots. So it is possible that the man was in love, and that he was going to get married. Then there is a very special motif, called a caduceus. This is a herald's wand, a kind of stick. The stick has two serpents twined round it, and two wings at the top. This stick is normally carried by a Greek god called Hermes. Hermes is the god of trade and commerce. What does this mean? Maybe the young man was a wealthy textile merchant?

Pommel: Under his elbow, clamped between the arm and the body of the young man, we see a rounded knob, a pommel. This means that he is carrying small sword. This kind of sword is called a rapier.

Lace Cuff: The cuff is made of white lace. It shows complicated geometric forms. The cuff is so oversized that we cannot see the hand beneath it.

4.3 Evaluation study

As a starting point of the improvements made to the Interactive Audio Guide, we decided to perform a first user study with the London exploration group very early on. We got the opportunity to get copies of an earlier relief, Gustav Klimt's "The Kiss", as a test-case until the new reliefs for the museums are ready. These test sessions were interleaved with improvements on the programs as outlined in section 3.2, so that feedback could be immediately acted upon and the changes tested in the following session. The participant-led method provided a greater depth of feedback as well as a more critical approach which benefited our analysis.

We also asked participants to identify their access needs, rather than their type of disability, which provided the opportunity for a complimentary, overlapping but differently-orientated dataset to compare against the first study. The different group, spread of participants, research method and approach resulted in significant progress both in the technical innovation of the IAG and analysis of its wider applicability.

We started with an informal 5 hour session with 25 people across a wide range of sensory, learning and cognitive impairments. After this first feedback we implemented a structured evaluation which took place over the course of 2 days with 14 people (10 female, 4 male, aged 18–75, average 45, see Figure 61).

Sex	Male	Female
	4	10

Age (years)	0-9	10-17	18-29	30-39	40-49	50-59	60-69	70-79
			3	5		2	3	1

Why do you visit museums? (You may tick more than one)	Don't visit any	For educational purposes	For enjoyment purposes
	1	8	11

How often do you visit a museum per year?	Never	< 1x year	≥ 1x a year	≥ 1x a month	≥ 1x a week
	2	2	3	5	2

How do you visit the museum? (You may tick more than one)	Friends	Project	Alone	Family	Support	Other	no visit	School
	7	6	4	3	2	1	1	

What technology do you use? (You may tick more than one)											
iPad	iPhone	W. Laptop	A. Phone	A. Tablet	Kindle	MacBook	W. Tablet	W. Phone	Dictaphone	Other	None
8	7	7	4	3	2	1	1	1	1	2	2

Have you used tactile reliefs before?	No	Yes	I don't remember
	9	3	1

A ... Android W ... Windows

Figure 61: Demographic data.

The approach taken in this study was based on the concept that people cannot be neatly allocated in disability categories: they instead have access needs that may relate to one or more categories of disability or impairment. For example a participant who would be typically classified as visually impaired may prefer the sort of one-to-one support typically associated with those who have learning difficulties. Likewise, visual impairment often accompanies learning difficulties and vice versa [8]. Asking participants about their preferred access needs and not their disability not only enables catering for those needs but also furthers the creation of a universal tool that can be enjoyed by everyone regardless of category or need.

The key needs that emerged from the participants in this second study were: one-to-one support (7); audio description (4); captioning (3); simplified information (3); tactile books (3); Braille (1); and British Sign Language (1); see Figure 62 for further details. Though only one participant in each case required Braille and BSL these were the main way in which they processed information and their preferred medium. Most participants were interested in museums, 7 going at least twice a year, and 3 at least once a year, however 4 rarely went to museums, if at all (see Figure 61).

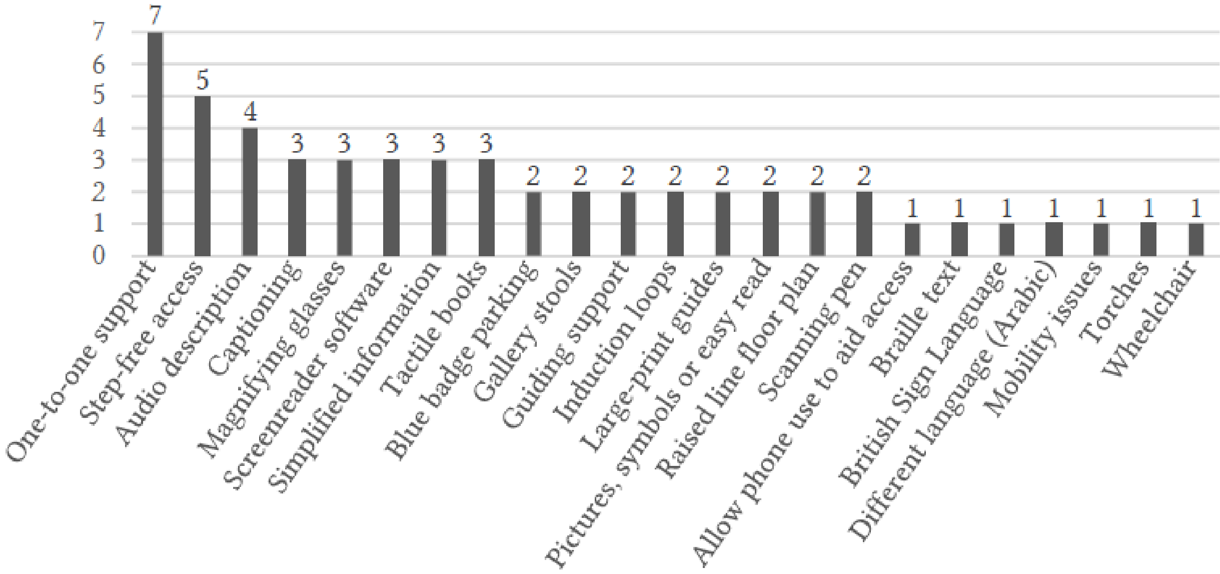

Figure 62: "Please indicate your access needs (You may tick more than one)." (13 responses)

The same basic testing and evaluation method was used as in the first study, with one relief being tested; however, an important difference was the use of a participant-led research method.

34 questions were asked, most being on a 10 point Likert-scale, 1 being the most negative, 10 the most positive ranking. These are summarized in Table 1 to Table 3.

4.3.1 General impression

In general, the response was positive: it was "fun", the multi-modal approach "made the information easier to process" and participants felt they "connected" to the painting and got to "explore it more deeply". Several people said that they liked having the background information about the painting as this helped put it in context. The tactile element helped them pick out details they would otherwise have missed. After the first session one participant requested a clear training mode. The training mode that was created (see section 3.2.3) was well-received but it only allowed participants to explore the gestures – it didn't give them the step-by-step guide that was needed. Also, it became clear that audio or descriptive instructions do not work for everyone: "With learning disabled people you can't just tell them what to do; you need to show them what to do". This comment from a support assistant shows that there is a need to create a more social training mode, with videos and images. In particular people with learning difficulties took more time understanding the process. They had not used tactile elements previously and were therefore not used to it.

Overall, it seemed that those with the most severe level of visual impairment appreciated the system most, as people with sight do not have equivalent dependence on touch and audio. Of course, it is important to note that those with visual impairments are often more used to using touch as a means of learning and orientation, so this is not necessarily a direct or fair comparison. Some participants found it easier than others to get used to navigating by touch and audio.

1	2	3	4	5	6	7	8	9	10	avg.	sum
1. How did you find using the IAG? (1 = not good, 10 = very good)											
				3		2	2	3	4	8.0	14
2. How understandable did you find the instructions? (1 = not, 10 = very)											
		3			3	2	2	2	2	6.8	14
3. How simple were the gestures for you? (1 = not, 10 = very)											
	1	1		2	1	2	2	3	2	7.0	14
4. How important is it, that the audio guide is triggered only by certain gestures and does not talk all the time? (1 = not, 10 = very)											
1					1	1	3	2	5	7.6	14
5. How easy was it to select certain parts? (1 = not, 10 = very)											
	1	1		2		3	3	3	1	6.9	14
6. How important is the tactile element for you? (1 = not, 10 = very)											
1				3		2	2	1	5	7.5	14

Table 1: Questions concerning the technology of the interactive audio guide.

4.3.2 Interface

As before, the system was designed to more closely resemble a kiosk in a museum; therefore an introductory text should be sufficient for use. Our original intention was to use kiosk mode, with initial mandatory instructions as used in a previous study. However, after the first session it was clear that participants needed (and were given) one-to-one support. They consistently fed back that a training mode, taking them through the gestures, would be helpful. As a result, a training mode was created for the second session (see section 3.2.3). This training mode allowed the instructor to layer the information necessary to use the IAG. The addition of this mode made a significant difference, especially to the accuracy of the participants' gestures and how quickly they were able to pick them up. However, though they appreciated this addition, most fed back that instructions accompanied by a video would give them the best chance of navigating the system in true kiosk mode. Further developments to this end and testing in a museum environment are required before the IAG can truly work in kiosk mode.

Having received one-to-one support and testing the setup, nine people rated the question "How did you find using the IAG?" above 8, with all but three giving it a ranking of 7 and above (Table 1, Question 1). Participants were giving feedback such as that the IAG "allowed me to precisely connect what I was feeling to the description of the painting" and that "the description provides extra detail that is engaging and interesting". Some liked it because they "could actually feel the painting" (2×), they liked "the voice and liked exploring the portrait by touch whilst the voice told you what you were feeling", it "gives you an all round experience regardless of your access needs" and it "Helps to explore the painting more deeply at an interactive level". For others it "took time to adapt to it and getting used to the directions" but had a "very positive experience once up and running".

One user found that there are "Not very good tactile markers to find your place", and another would like to see "more visual content especially an image of the picture to make a comparison". Over half found the instructions easy to understand (ranking above 7) and all but three gave a ranking of 6 and above for this (Table 1, Q2). Some found that the starting instructions are "not clear enough" and that they needed "step by step instructions on how to use it". Among those that gave the lower rankings on both questions were people who had not listened to the instructions as well as (and partly in combination with) having learning or cognitive impairments that make retention of information more difficult. These participants needed and received one to one support from the instructor and, in a museum setting, would frequently be accompanied by a support assistant.

Half the participants mastered the interface through the training mode and could reproduce all gestures without any intervention from the researcher. As with the previous study, others needed guidance or slight manual corrections of their hands. Two participants had difficulties using the gestures and indicated that they would prefer a “keyboard” or “button” alternative. A support assistant also pointed out that some disabled people might have additional conditions like Parkinson’s or arthritis. However, generally, when asked how easy it was to perform the gestures, nine participants rated 7 or higher (Table 1, Q3). Comments included, “it’s really instinctive”, “cool!” and “I like the fist!”. The design goal, to only have the system play audio when it is explicitly requested by the user, was supported by the responses: 7 participants gave a ranking of 9 or 10, and all but two gave a ranking of 7 and above (Table 1, Q4). The repeated feedback was that people wanted to explore the relief and the information at their own pace and not be bombarded by the constant audio information.

4.3.3 Off-object gestures

In the first session, the off-object gestures were the ones participants struggled with most. Getting the correct height, keeping the hand flat and being able to spread the fingers without distorting the hand shape – these were all common problems. Some conditions affect the physiology of the hands – e.g. shorter, thicker fingers – which meant that the gestures were harder to make for those participants and that if was harder for the algorithms to detect the gestures. Those with the most visual impairment sometimes found it hard to know when they were at the correct height, making the correct shape and if their hand was flat, as they had no visual or tactile cue for checking and correcting this.

Taking these difficulties into account, we created a training mode for the second session (see Section 3.2.3), and added a *fire-crackling sound* to guide participants to the correct height to trigger the audio (see Section 3.2.1). This allowed participants to practice making the gestures and gave audio feedback when these were performed correctly. This gave the participants a standard to refer back to when they were navigating the IAG. Both additions significantly improved the overall experience and meant less guidance and physical adjustments were needed from the instructor.

Participants with joint difficulties and smaller, thicker hand shapes pointed out how they would prefer a button mechanism at the bottom of the relief and/or an alternative set of gestures that worked on the movement of the fist rather than specific hand shapes (i.e. one shake for item 1, two for item 2 etc). One participant even had to support the gesture-making hand with her other arm as she got easily tired of waiting until the camera read her hand gesture. These alternatives would also be useful for those with learning or cognitive impairments as there is not so much to learn and doesn’t require remembering as long a sequence of instructions.

4.3.4 On-object gestures

Using the pointing gesture with the index finger became easier for the participants after clear instructions and corrections. The main issue was that participants didn’t close the fist completely and left the thumb extended. The participants also struggled with the correct angle. The tendency was to have the finger vertically which was uncomfortable for them. In addition the participants tended to press on the relief harder rather than gliding over it.

Participants relaxed after being reassured that they were doing the right gesture. A more social element would give the participant the confidence needed to explore the relief.

Another limitation detected after the first session was that BVI participants were unsure where individual sections or the relief itself stopped. As a result, we implemented a *rain sound* that is triggered as the hand approaches a new section or boundary of the relief. This addition helped BVI participants significantly but also some of those with cognitive and learning impairments.

One of the aspects that need to be thought through further is the automatic necessity of the additional sound design (see Section 3.4.4). Participants with learning difficulties were calmly guided through the process of testing and largely ignored the sounds but for independent testing these sounds might become distressing. Being able to switch the function on/off should therefore become an option for the user.

4.3.5 Content

When asked whether the IAG helped gaining a sense of the whole painting, nine gave a rating above 6 (Table 2, Q2). This increased to twelve participants when asking about the details of the painting (Table 2, Q3). All but one of the participants felt the level of detail was just right; one participant wanted more detail as they were particularly interested in exploring the painting at a deep level and were also highly competent at navigating the system (Table 2, Q7). The physical replication of the details such as “the way that the different shapes represented the pattern of the clothes” were particularly popular and received high praise from most of the participants with an average ranking of 6.7 (Table 2, Q1). There was an overall desire for colour on the relief, regardless of whether there was a copy adjacent to the station.

Overall, 11 of the 14 were happy with the amount described (Table 2, Q4). One participant went through the parts but didn’t listen to the descriptions as they were short of time and wanted to have tried each of the aspects before they left. One of the challenges with the content was a) testing it on people with hearing loss and b) testing it with participants with learning difficulties. During the first session we presented subtitles in a text editor in an ad-hoc attempt to make it accessible for one participant with hearing loss, although in a very small font. For the second session proper subtitles were implemented. This enabled us to test the IAG with D/deaf participants.

People with learning difficulties used this as well as it helped them “focus” more on the content. Further comments from participants suggest that inverted colours would ease the reading and “make the text pop out more”. This was a suggestion from a support assistant with dyslexia and is also typically found with people who have a visual impairment.

Nevertheless, several participants still found the language difficult and too “academic” and “formal”. This further supports the above-mentioned statistic that 23.1% of our participants (3 out of 14, Figure 61) have an access need of simplified information. Here, the differences in the level of education attained were noticeable in how accessible the participants found the information: broadly speaking, the lower the level of education, the more difficult the participant found the information.

1	2	3	4	5	6	7	8	9	10	avg.	sum
1. What was your impression of the relief? (1 = not good, 10 = very good)											
1				4		4	2	1	2	6.7	14
2. How well could you get a sense of the whole painting? (1 = not, 10 = very)											
1			3	1	3	3	1	1	1	6.0	14
3. How well could you get a feeling for details? (1 = not, 10 = very)											
1			1		4		2	4	2	7.2	14
4. How happy were you with the number of parts described? (1 = not, 10 = very)											
		1	2		1	2	3	3	2	7.3	14
5. Would you find it useful in a museum? (1 = not, 10 = very)											
1						2	1	1	8	8.0	13
6. Would you rather go to a museum that offers such audio guides? (1 = does not matter, 10 = much rather)											
1				1	2		3		6	7.3	13
7. How would you rate the amount of detail in the relief?											
too little		OK		too much							
1		12									13
8. Would you prefer other types of description? (e.g. music; poetry; storytelling)											
No		Maybe		Yes		I don't know		I don't mind			
3				10				1			14

Table 2: Questions concerning the content, i.e., relief and audio description of "The Kiss" by Gustav Klimt and the application of this technology.

Overall, the language needs to be simplified but there may also be a case for providing a setting where you can choose your level of description with an option consciously aimed at those who require simplified content. This could also include symbols and pictures to support the content which would be a valuable addition for those with cognitive and learning impairments. In contrast to the first study, this study was conducted in the United Kingdom and the language used, English.

One participant had severe difficulties understanding the speaker and translated text "maybe it's his heavy accent... I don't know". Therefore, if this application is to be used beyond Austria, text creation and audio recordings by native speakers would need to be considered.

The majority (71%, 10 out of 14, Table 2, Q8) said that they would enjoy an alternative, more creative, form of description such as music, poetry or storytelling. Only 3 participants didn't like or see the need for this. The most consistent suggestion was a storytelling approach that included a first person narrative. Such a narrative would also help those who need simplified language as they can find it easier to relate to narrative content rather than the purely factual. As one participant said, it needs to be "more informal and creative rather than just giving the facts in quite a scientific way.

At the moment a bit robotic, make it a bit more human". Having music as part of the descriptive process was also popular "as it creates an atmosphere". Finally, a strong case was made for having audio content related by BSL with a signer in a pop up video: this is especially important for those who have BSL as their first language. Ideally, content would be created in BSL itself, in collaboration with a signer, rather than necessarily being a literal translation of the text. This would increase the quality of the language for users.

4.3.6 Acceptance and field of application

All but 3 participants had never tried a similar technology before. This number is not necessarily surprising as a participant said "I would have never tried it. I am used to the visual learning and that's it". Compared to similar experiences participants said: "Good because it is live but not enough explanations compared to other reliefs, e.g., Living paintings[1]."; "Great improvement on others."; "Very different experience (never done paintings)". 62% stated that this technology would be very useful within the museum environment (8 out of 13 ranked this 10, Fig. 16, Q5). Only one participant said that it would not be useful for them. It is interesting to note how these responses map onto the users' experience of technology in general: all used a Smartphone, all but one used a tablet of some form and 8 out of 14 had a laptop or Mac Book (see Figure 61). Therefore, the feedback on the IAG specifically and this sort of technology in museums generally, was coming from participants familiar with using technology in everyday life. Therefore, it is important not to assign unfamiliarity with technology in general as a too significant factor, especially in analyzing the more critical comments or noted problems in navigating the IAG. In the future more time and a clearer, layered and social tutorial should be offered to them and would support a confident and competent use of the IAG.

Six participants said that they would strongly (10 points on the Likert-scale) rather go to a museum if this would be available to them, three who rated 8 and three who remained rather neutral between 6 and 5 (Table 2, Q6). Though half the participants visit museums at least once a month, we had 2 who visited museums less than once a year and 2 who indicated they would never visit a museum. It is also important to note the conditions under which participants visit museums: the half who visit infrequently are much less likely to go outside of a project or organized trip and require support in physically getting to the museum as well as one to one support once on site (Figure 61). This highlights that even with high-tech solutions to accessibility such as the IAG, there remains a social element that is vital, especially for those who have the greatest access needs.

However, it still remains the case that innovations such as the IAG open up new potential for these participants in seeing museums as relevant for and accessible to them. We just need to make sure the social support is also in place.

Several participants particularly praised how good this system would be for children and in a school or education setting – they were keen for this aspect to be developed. Given the bulky and heavy nature of the IAG especially with the HP Sprout, a more portable version may need to be developed if this is to be practicable, unless such devices are already available there.

4.3.7 Multi-sensory experience

The IAG has the potential to further enrich the experience for non-BVI users. Therefore, we included another set of questions into our Evaluation regarding different ideas to further enrich multi-sensory experiences (see Table 3). 5 participants (36%) think that a projection would very much add to the value of the relief (i.e., ranked 10, Table 3, Q2). When asking participants about possible projections onto the relief, 13 participants indicated that they wanted the original painting; 4 wanted a more simplified version; 3 wanted high contrast; and 3 asked for have an animated version of the painting (Table 3, Q3). None of the participants said that they would rather have it blank. For future applications these results will have to be taken into consideration and implemented, especially if aiming to widen the application beyond BVI users.

When asked about a haptic response (such as vibration) 6 (43%) indicated that they would like to have it (Table 3, Q4). Vibration could help participants with orientation but also help those who have both hearing and vision loss: such intersectional needs are currently rarely catered for and providing this would set the IAG apart. Participants who currently use technology such as smart phones (and particularly those who use

[1] Living paintings is a charity that provides touch to see books for BVI people, see http://www. livingpaintings.org

the accessibility features for low vision and hearing loss) were already familiar with this technology and keen to see it used in this context.

Regarding how suitable the users from the participatory research group found the material on a scale from 1 to 10 (Table 3, Q1) and if they would prefer a different one, some participants indicated that they wanted something rougher and matte-textured such as clay. This was more common to those with visual impairment but also was mentioned by other participants. Another indicated that they would like the material to be chosen according to the artwork for example rougher material replicating brushstrokes for Van Gogh and smooth, shiny materials in the case of The Kiss to represent the gold and metallic elements. This feedback indicates a preference for the creation of a “collage” tactile effect and almost a move to the relief itself becoming an artistic object in its own right. This proposes an interesting avenue to reflect on.

The majority of participants (11, 79%, yes and maybe) were interested in sonifying the relief – representing colours and materials through sound (Table 3, Q5). 6 (43%) definitely wanted this feature added and were excited by the idea – “it would be the final piece of the puzzle, making it a complete experience”. Such a development has the potential to enable participants to relate to paintings based on personal experience. Even participants with complete sight loss are likely to have had some sight at some stage in their life and therefore have some memory and knowledge of colour [9]. For those who have been blind since birth colours can be translated into temperatures (blue equalling cold, red as hot etc.), so this feature has a wide applicability. Again this might be introduced as an on/off switch option to give the user more control of their experience.

8 participants indicated that they would like to test a 3D model compared to a 2.5D model (Table 3, Q6). Three participants with learning disability had difficulties exploring the faces of the depicted figures and identifying their gender. By exploring it from every angle this might solve the issue. Even though the production of such model is rarer and more costly, depending on the nature of the artwork this might be preferable. However, the use of the IAG would have to be reconsidered and the system redesigned to full 3D model tracking so that these can be freely turned and have accessible touch-sensitive parts all around.

1	2	3	4	5	6	7	8	9	10	avg.	sum
1. How did you find the material? (1 = not good, 10 = very good)											
1		1				1	3	2	6	8.1	14
2. How would you rate the added value of projections on the relief? (1 = not good, 10 = very good)											
			1	1		2	5		5	8.1	14
3. What would you like to project on it? (You may tick more than one option)											
Original Color		Simplification		High Contrast		Animation		Nothing			
13		4		3		3					
4. Would you like to have a haptic response? (i.e. vibration)											
No		Maybe		Yes		I don't know		I don't mind			
1		5		6		1		1			14
5. Would you like color/texture represented by sound as you move your fingers?											
No		Maybe		Yes		I don't know		I don't mind			
2		5		6				1			14
6. Would you prefer a relief or a 3D model?											
I prefer a relief				I prefer a 3D		I don't know		I don't mind			
4				8		1		1			14

Table 3: Questions concerning multi-sensory experience.

4.3.8 Important qualifications to the data

People with learning and cognitive impairments are often keen to please and to give the "right" answer in a question situation as well as not wanting to show that they don't know or understand. For example, one participant had answered the age category in the questionnaire but this was corrected by their support worker as in reality the participant didn't know or couldn't remember their age but didn't want to admit this. Therefore, without careful analysis the data can be unreliable and misleading. To offset this and avoid any potential skewing, we have complemented the answers given by these participants with the data and observations of the examiners and the participants' support assistants. We can therefore have greater confidence that the data used is accurate and a fair representation of participants' experiences of the IAG.

It is common for disabled people in general, and amongst our participants specifically, that people do not fit solely in one category such as BVI or Hard of Hearing (HoH) – they have intersectional impairments and needs, often the combination of a sensory impairment and learning/cognitive impairment. Therefore it is important that this is taken into account when reading the raw data and drawing conclusions on the use of the IAG for specific groups. We have taken care to build this intersectionality into our analysis of the data and our conclusions.

4.3.9 Conclusion

In contrast to our old setup, the HP Sprout platform offers a commercially available and museum-ready option that minimizes setup-time and has a nice look and feel, although sensor-placement is not optimal and the effective workspace and accuracy might suffer a bit. Furthermore, we showed, that through a more diverse user group and the use of a participant-led method of testing, the development of new technologies aimed at better accessibility is not only more inclusive, but also more efficient.

The main findings from our study are: the key importance of multi-sensory nature; the ability to personalize systems for specific needs, e.g., sign language, captions, colour, simplified language; the limitations of needing precise gestures to control the IAG especially for those with mobility issues. Especially the second user study highlighted the need for a simple and socially engaging tutorial to teach users how to use the system. While this finding may perhaps be obvious in hindsight it is a very useful reminder about the importance of compelling training techniques for new access technologies.

Beyond the visually impaired participants, participants with different access needs said that the tactile element "explains the painting by getting people involved and helps them understand more. We all want to touch things!", "It helps with low attention span – taps into curiosity.", and "It gives me a deeper understanding of the piece".

Overall, the work provides an approach that may not only reduce barriers to the accessibility of visual art for people with many different disabilities but may provide an entirely new modality that helps all museum visitors appreciate art in an exciting new way.

This study was a very important first step in using the system for a wider audience, and gives valuable insights. As these results are of high interest, we plan to use them as part of a publication in the near future.

5 Conclusions

In this deliverable we gave a detailed report about the different tools that have been developed since the start of the project.

We introduced the relief design software we developed for the semi-automatic generation of tactile relief files, including some new and enhanced functions that will significantly advance the design process. While the software is still in development, we could already show the basic design and functionalities.

We also introduced the context sensitive tactile audio guide. We improved the setup according to the needs of the participatory research group (concerning calibration, screen design, interaction) and developed an authoring tool for the audio content, currently as an internal tool, but targeted to be used by the museums in the future. The new platform based on the HP Sprout offers a museum-ready setup, and was positively accepted. This is especially evident in the first evaluation study that was conducted within the participative research group in London.

Furthermore, we included some insights into the current state of the design process of the tactile reliefs for two museums. These will be produced in the near future, and are then available for user tests by the participative research groups, as well as the content for the audio guide that was authored so far. Design of the reliefs and audio guide for the other museums will start once their participative research groups are formed and the artwork selection process is complete.

To conclude, work package 5 is right on track and good in the time plan. All proposed features could be implemented. With the HP Sprout’s possibility to project onto the reliefs, we even went beyond what we could imagine while writing the Grand Agreement. We collected a lot of ideas for further developments, which are now being ranked to be implemented in the following stage.

References

[1] Gary Bradski and Adrian Kaehler. 2008. Learning OpenCV: Computer vision with the OpenCV library. "O'Reilly Media, Inc.".

[2] Alvaro Cassinelli, Easy Camera + Projector Calibration, YouTube video. Available on line (29 September 2017): https://www.youtube.com/watch?v=pCq7u2TvlxU

[3] Alvaro Cassinelli, Example_CameraProjectorCalibration, code at github, https://github.com/alvarohub/Example_CameraProjectorCalibration

[4] Yang Chen and Gérard Medioni. 1992. Object modelling by registration of multiple range images. Image and Vision Computing 10, 3 (1992), 145 – 155. DOI: 10.1016/0262-8856(92)90066-C

[5] Bertram Drost, Markus Ulrich, Nassir Navab, and Slobodan Ilic. 2010. Model globally, match locally: Efficient and robust 3D object recognition. In Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on. IEEE, 998–1005.

[6] A. Reichinger, A. Furhmann, S. Maierhofer, and W. Purgathofer, *„A Concept for Re-Usable Interactive Tactile Reliefs"* in *„ICCHP 2016 Part II (LNCS)",* Heidelberg, Deutschland: Springer, 2016, pp. 108-115

[7] A. Reichinger, S. Maierhofer, A. Fuhrmann, and W.Purgathofer, "Gesture-Based Interactive Audio Guide on Tactile Reliefs" in "Proceedings of the 18th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '16)", New York, NY, USA: ACM, 2016, pp. 91–100.

[8] Royal National Institute of Blind People. n.d., *Across the spectrum: learning disability and and sight loss.* (n.d.). Available on line (30-08-2017): http://www.rnib.org.uk/services-we-offer-advice-professionals-nb-magazine-health-professionals-nb-features/across-spectrum

[9] Royal National Institute of Blind People. Key information and statistics. n.d. url: http://www.rnib.org.uk/knowledge-and-research-hub/key-information-and-statistics (visited on 2017-09-01).

[10] J. Stewart and A.R. Quattrociocchi, *"Method and apparatus for three-dimensional digital printing",* US Patent 9,061,521 (23-06-2015)