User-Appropriate Viewer for High Resolution Interactive Engagement with 3D Digital Cultural Artefacts

David Gillespie
Centre for Digital Entertainment
Bournemouth University

Doctorate of Engineering
22nd of November 2017

Academic Supervisor: Dr Hongchuan Yu, Prof Kate Welham and Dr Ian Stephenson
Industry Supervisor: Danny Boardman

Signature of Author . . . .

David Gillespie

Abstract

The core mission of museums and cultural institutions is the preservation, study and presentation of cultural heritage content. In this technological age, the creation of digital datasets and archives has been widely adopted as one way of seeking to achieve some or all of these goals. However, there are many challenges with the use of these data, and in particular the large numbers of 3D digital artefacts that have been produced using methods such as non-contact laser scanning. As public expectation for more open access to information and innovative digital media increases, there are many issues that need to be rapidly addressed. The novel nature of 3D datasets and their visualisation presenting unique issues that impede use and dissemination. Key questions include the legal issues associated with 3D datasets created from cultural artefacts; the complex needs of users who are interacting with them; a lack of knowledge to texture and assess the visual quality of the datasets; and how the visual quality of the presented dataset relates to the perceptual experience of the user. This engineering doctorate, based on an industrial partnership with the National Museums of Liverpool and Conservation Technologies, investigates these questions and offers new ways of working with 3D cultural heritage datasets. The research outcomes in the thesis provide an improved understanding of the complexity of intellectual property law in relation to 3D cultural heritage datasets and how this impacts dissemination of these types of data. It also provides tools and techniques that can be used to understand the needs of a user when interacting with 3D cultural content. Additionally, the results demonstrate the importance of the relationship between texture and polygonal resolution and how this can affect the perceived visual experience of a visitor. It finds that there is an acceptable cost to texture and polygonal resolution to offer the best perceptual experience with 3D digital cultural heritage. The results also demonstrate that a non-textured mesh may be as highly received as a high resolution textured mesh.

The research presented provides methodologies and guidelines to improve upon the dissemination and visualisation of 3D cultural content; enhancing and communicating the significance of their 3D collections to their physical and virtual visitors. Future opportunities and challenges for disseminating and visualising 3D cultural content are also discussed.

Publications

La Pensée, A., Cooper, M., Gillespie, D., *2012 Things that we thought were straightforward when we started 3D scanning cultural heritage; copyright, data archiving, internet security, and access for all.* Computer and the History of Art (CHArt) 28th annual conference, Consume – Digital Engagement with Art, Association of Art Historians, London

Gillespie, D., La Pensée, A., Cooper, M., 2013. *User appropriate viewer for high resolution interactive engagement with 3D digital cultural artefacts*. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XL-5/W2, 271–276. doi:10.5194/isprsarchives-XL-5-W2-271-2013 – Contributes to chapter 3

Gillespie, D., La Pensée, A., Cooper, M., 2014*. 3D Cultural Heritage Online; In Search of a User Friendly Interactive Viewer.* International Journal of Heritage in the Digital Era 3, 51–68. doi:10.1260/2047-4970.3.1.51 – Contributes to chapter 3

Gillespie, D., 2016. *Copyright and Its Implications for 3D Created Datasets for Cultural Heritage Institutions.* International Journal of Culture and History, Vol. 1, No 2. December 2015 – Contributes to chapter 2

Acknowledgements

I would like to give a big thank you to my previous industry supervisors, Annemarie La Pensée and Martin Cooper for their support and guidance. Thank you both for taking the time to share your knowledge and give me the support I needed during my initial years in my EngD. I would also like to thank Hongchuan Yu, Danny Boardman, Kate Welham and Ian Stephenson for their support and guidance through the final years of my EngD.

I am also grateful to the Centre for Digital Entertainment at Bournemouth University and the EPRSC for supporting this research. I would especially like to thank the Centre of Digital Entertainments not just for the funding but for the on-going support from their Staff. I would especially like to thank Daniel Cox, Katja Haferburg and Michael Board for the support and the opportunities they have provided for me.

I am also grateful to the National Museums Liverpool, to allow me to continue my research that I had begun with Conservation Technologies and allowing me to access the collections to support my research. I am also particularly thankful to Phil Phillips who was a great advocate for the use of 3D cultural artefacts within National Museums Liverpool's galleries, and helping to set up various experiments.

I would like to give a special thanks to my friends and family who supported me during the highs and lows of my EngD. I would never have reached this point without the support from my friends Marjon Kruize, Jamie Pugh, Catrin Jones, Adam Scott, David Jones, and to all of the wonderful people I met in the Liverpool Drama Society. I would also especially like to thank my previous partners, Mary-Ann Gilbert, Michalina Piekarska and Joanna Thomas, who continued to push me forward. Without them, I doubt I would have made it this far.

A final big thank you to my father, for just being there for when I needed support or someone to talk to. He was a big supporter of my time as a student and thought I would be an eternal student. In His eyes I will be. I would like to dedicate my EngD to the memory of my father John Brough Gillespie, who passed away shortly before the completion of this Doctoral thesis.

Table of Contents

List of Figures

List of Tables

List of Acronyms

3D: Three dimension

2D: Two dimension

ABF: Angle Based Flattening

API: Application Programming Interface

BGLAMS: Basic Grey Level Aura Matrices

CDF: Cumulative Distribution Function

CDPA: Copyright, Design and Patents Act 1988

DCP: Discrete conformal parameterisation

FPS: Frames per second

FTT: Fast Fourier Transform

HCI: Human Computer Interaction

HLSCM: Hierarchical least square conformal mapping

HSD: Honestly Significant Difference

LCSM: Least square conformal mapping

LDR: Low Dynamic Range images

MDS: Multi-Dimensional Scaling

MIPS: Most Isometric Parameterisations

MOS: Mean opinion score

MSDM: Mesh structural distortion measure

RBF: Radial Basis Function

SfM: Structure from Motion

SPM: Shape Preserving Method

VDP: Visible Difference Predictor

Chapter 1 Introduction and Background

1.1 Introduction

The roles of museums and cultural institutions are changing; they are no longer just the gatekeepers of our cultural heritage. They are moving away from closed systems, evolving slowly to share and increase access to their collections for research, education, and inspiration, empowering global and local communities. Technological advances have allowed institutions to innovate, granting more access to relevant information with the click of a button. Data has become more accessible and less restrictive for the general public, with the increased digitisation of paper archives, physical collections and catalogues of knowledge. What was once only available in a physical format now can be searched in a database which may include additional photographs of the object in question. Examples of large scale digitisation programmes and collaborations include Google's Art project [1] and the digitisation of the Endangered Archives programme at the British Library [2]. These publicly available collections have allowed researchers to cooperate and study these objects, increasing the knowledge gain for cultural institutions. This in turn has facilitated more data being available for the general public to learn about the collections; within the museums, their archives, and objects that are either too fragile or damaged to be exhibited [3]. The creation of these databases presenting information in the form of text and images has been improved upon in the past decades, and is a valuable approach for the dissemination and sharing of 3D documentation of cultural heritage objects.

1.2 3D Digital Models: Opportunities and Challenges

The last decade has seen the rapid adoption of technologies that can accurately measure the physical world and create high resolution digital models of cultural heritage artefacts or sites. One of these documentation technologies is non-contact laser scanning, which was adopted from the aeronautical industries and can record the surface of an object with sub-millimeter accuracy [4]. This technology allows for the creation of three dimensional (3D) digital models that enable the general public or researchers to explore and interact with these objects or sites remotely. The dataset can be used for a range of applications including: documentation, archiving, surface monitoring, gallery interactives, educational sessions, conservation, and visualisation. It can also be used in conjunction with rapid prototyping technologies such as computer numerically controlled (CNC) machining or 3D printers to create replicas of the original artefact in a range of different materials [4]. The applications of 3D datasets are also increasingly more important with regards to the current threat to cultural heritage objects and artefacts that face destruction via natural disasters or other forms of destruction. UNESCO recently launched Unite for Heritage in an effort to safeguard and protect heritage under attack [5]. It highlighted the importance of heritage in its empowerment of people, [5] and the importance of digitisation and documentation to safe guard cultural heritage for future generations [6, 7].

The documentation of cultural heritage artefacts via laser scanning is non-invasive, causing no damage to the surface of the object during the recording process. The created 3D dataset allows for the surface of the original object to be measured and recorded without the need for any physical interaction with that original object. It also allows for an object to be measured over time, where the original artefact may degrade due to environmental changes.

The 3D replicas created through laser scanning are also often used as a surrogate replica for the original artefact, either as a 3D digital representation of the original artefact or as a physical stand in. These surrogate replicas can be used for a variety of purposes depending on the needs of the cultural heritage institutions, researchers, and educators, or within other industries. The created 3D digital replicas can be used for immersive video games, augmented or virtual reality and 3D printing and visualisation purposes [8]. There is also developing technology such as the Oculus Rift [9], VIVE [10] and 3D televisions that help to facilitate new immersive experiences and interaction with 3D content. However, the novel nature of a 3D dataset is presenting a rather unique challenge with respect to its dissemination and display. These issues include: the legal and intellectual property (IP) rights of 3D files created from cultural heritage objects, the interaction and accessibility of engaging with 3D cultural heritage objects and lastly providing a fit for purpose 3D digital heritage object that can provide a digital surrogate for the original artefact, yet be manageable for both dissemination and visualisation.

As 3D content is becoming more readily available, there is a need to address the issue of presenting high resolution 3D cultural artefacts, providing knowledge, safeguards and the best possible experience for users interacting with them.

1.3 Context

The dissemination and sharing of 3D cultural heritage datasets can offer many benefits. Despite this, access to the bulk of this 3D content by user groups including the public, researchers and conservators, locally and remotely is still rather limited. To offer the best perceptual and visual experience for users interacting with 3D digital culture, there are three major areas that cause issues for the adoption of 3D content and its dissemination, which will be introduced below.

1.3.1 Intellectual Property and Legal Issues

3D digital datasets of cultural heritage artefacts provide an extremely accurate recording of the original artefact's surface. This digital file can provide an unlimited number of digital surrogates of the original artefact but also can be used to create physical surrogates of the original artefact in a range of different materials if used with 3D printing. This has led to cultural institutions raising concerns over the ownership of the digital file and if it is possible to license these objects to other institutions, similar to how they license two dimensional (2D) photographs of works within their institutions [11]. There are other major concerns, especially if the 3D dataset was illegally infringed and shared online without appropriate permissions. There are worries this could lead to the loss of control of the digital file, loss of potential revenue streams, and the degradation of the digital cultural heritage artefacts.

1.3.2 Dissemination and Interaction

Text and 2D images are good for the dissemination and sharing of cultural content, as they have small file sizes, can inform and provide information to visitors and are widely adopted within cultural heritage. However, for research purposes or allowing a user to fully understand the shape and nature of the artefact, they are limited. The use of 3D datasets, allows for researchers and the general public to interact with cultural heritage in new and exciting ways and provide new research opportunities. However, unlike text and 2D images it is difficult to disseminate and access these files due to their file size and the lack of an appropriate way to interact with these datasets. This is vitally important especially in an environment where the user chooses when and how to interact with information [12] and the novel nature in interacting with a 3D dataset within a webpage. There is also a concern over accessibility, as 3D content is normally rendered within specialised software limiting the ability to disseminate and access this 3D content. To access as large and diverse an audience as possible a key proponent would be the investigation into key technology that would be platform agnostic. The technology would need to broach all operating systems, and consider mobile devices as well.

1.3.3 Accurate depiction of the original artefact

Cultural heritage institutions have raised concerns about the visualisation and depiction of the 3D datasets, and whether it is an accurate representation of the original artefact. With the adoption of new technologies such as 3D technologies, guidelines and practices are needed. One of the most important and relevant document is the London Charter [13]. The London Charter [13] provides internationally recognised government practices in the field of 3D visualisation in research and the communication of cultural heritage. It was created by a panel of experts in 3D visualisation technologies within cultural heritage to provide best practices that would be aimed to provide intellectual transparency and documentation to allow visitors to understand the nature and “claims” being made by the computer visualisation [13]. It also provides guidelines and best practices for the display of 3D digital replica of the artefact, ensuring it is an accurate representation, conveying to the user the material the artefact is made from, and whether the visual representation is based on scientific evidence or speculation [13].

However, it does not provide a definition of accuracy with regards to the actual created dataset. It states that the presented information and visualisation should accurately inform the users of the presented knowledge, making a distinction between evidence hypothesis and probability of it being true [13].

Accuracy will be referred to a lot within this thesis when discussing the 3D datasets and derived stimuli against the original dataset, and is worth discussing now. Accuracy is defined as “The quality or state of being correct or precise” [14] or more technically the definition of accuracy is “The degree to which the result of a measurement, calculation or specification conforms to the correct value or a standard” [14]. This could infer that a 3D dataset or stimuli is accurate, if it can be used as substitute for the original artefact in certain situations. This could include returning the same measurements between two points on the digital and original artefact or offering the same views on the artefact and digital replica from the same viewing angle. However a 3D dataset cannot be accurate with regards to all of the properties of a physical artefact. A 3D dataset cannot be picked up and weighed with the hand or guessed with the eye, it is also not possible to ascertain the physical nature and material through touch. Even the most accurate 3D dataset of an artefact or object will be limited in some capacity with regards to the original artefact.

This raises concerns, especially if the recorded 3D dataset did not capture the surface material information. Also with the adoption of hand scanners, the recorded datasets may be a rough representation of the original artefact without their texture information. This is extremely important when cultural institutions wish to disseminate their 3D datasets; they may apply decimation or simplification techniques to the 3D dataset to reduce the file size and quality of the file. Though they may reduce the datasets file size, they may reduce the visual quality of the 3D dataset too much, which discourages users interacting with 3D datasets due to the low quality [15].

Attempting to define accuracy can be extremely ambiguous with a 3D dataset, especially when attempting to encompass all of the criteria of the real-world object. Therefore for this thesis, accuracy will be defined more specifically. Accuracy will be based on being able to reliably present the "real-world" content in a digital format; where a visitor will be able to recognise the real-world artefact, surface details and/or infer the material on the 3D dataset without the need for direct or physical contact with the original artefact. This specific definition of accuracy and its limited criteria allows for the accuracy of the 3D datasets and their stimuli to be measured within this thesis.

1.4 Problem Statement

This Engineering doctorate was undertaken in partnership between the Centre for Digital Entertainment and Conservation Technologies, to solve issues related to the dissemination and interaction with high resolution 3D cultural heritage artefacts. Conservation Technologies was a financially self-sustained department within National Museums Liverpool (NML), exploring potential technologies to document objects in 3D. This was of particular importance as it allowed conservators and researchers to explore the surface of an object in full for potential tool marks, monitoring surface changes due to corrosion or weathering and removing the need to apply a direct mould to the object. Conservation Technologies used non-contact triangulation based laser scanning to record the objects surface for objects that ranged in different sizes and materials [4]. The accuracy and resolution of the final dataset is entirely dependent on external factors, such as the material and lighting conditions during the scanning process and/or if it can be affected by human error or software limitations during the post processing. Conservation Technology's datasets on average are sub-millimeter accurate. During their tenure with NML, they created a loose archive of roughly 400 individual datasets, which were used for specific applications such as: replication, documentation, archiving, surface monitoring, identification and interpretation, gallery interactives, educational sessions, conservation, and visualization. The research in this thesis was driven by Conservation Technologies' wish to allow the datasets that had previously been created for one purpose to be reutilised for learning and engaging with the general public and being openly accessible online. The main issue that needed to be addressed from Conservations Technologies' point of view was the dissemination and visualisation of the 3D datasets. However, Conservation Technologies went into Administration in June of 2013 and the company had fully closed down on the 31st of August 2013. Research was refocused within NML, still focusing on the dissemination and sharing of their 3D Cultural artefacts but factoring in the legal and IP rights of the 3D objects and providing solutions that would not increase the strain on their already limited budget [16]. This would allow NML to display their 3D content in new and exciting ways within their website and galleries, allowing them to exploit their 3D datasets for the greater good.

The research presented in this thesis is motivated by the use of digital 3D cultural heritage as digital surrogates of the original artefact and its associated information, to provide in-depth knowledge to disseminate and share these with the general public. To achieve this, a new knowledge base and tools are needed to disseminate and share 3D cultural heritage locally and remotely. The following issues have been identified when reviewing on-going research and activities within cultural heritage.

1. The current legal framework and IP rights of scanned cultural heritage artefacts are not truly understood. Being heavily influenced by photography, it is assumed in contracts with employees or external parties that they will assign the property to the cultural institute and not the scanning technician. Photography can document an artefact, taking many pictures to create a full documentation, where artistry is recognised in the use of lighting, and different exposures [17]. Laser scanning makes a document that was created through an entirely mechanical process and is a slavish reproduction that may not warrant any IP rights [18]. There currently is very little information regarding the legal or IP rights of scanned cultural heritage artefacts.

2. Even with increased understanding of the legal and IP rights of cultural heritage artefacts, there is no comprehensible research that may alleviate the fear of sharing and disseminating 3D cultural content online. There is mixed research concerning the sharing of content regarding piracy, lost and gained revenue, yet once again there is very little information regarding the sharing and dissemination of 3D content.

3. Recent affordable hardware and software have led to an increase in the documenting and recording of artefacts in 3D and its associated information. However the access to the 3D dataset and information varies widely between user groups and is addressed case by case. There has been an increase in dissemination of 3D datasets, yet the access to these is still limited especially for the general public.

4. The display and creation of content for an artefact within a webpage or gallery is decided upon by curators and an institution's web team. This content would include text, multiple images at differing angles and possibly video as well. However, when asked to provide information for 3D content, it proved extremely difficult, and they approached it in a similar way to 2D content. They did not have the experience or knowledge to create tools or interactives for 3D content that would be usable and provide a good experience for visitors to their website.

5. Currently the recording of 3D cultural artefacts may lack surface materials. Some devices may take images to create a coloured model of the recorded artefact, yet there are issues such as; lighting and visual errors being recorded and the details being dependent on the resolution of the point cloud. These errors and artefacts can put people off interacting with the 3D cultural artefacts. There is a need for knowledge that would allow a cultural institute to place a surface material on a 3D dataset.

6. While there are guidelines provided by the London charter [13] for the display and visualisation of 3D digital cultural heritage, there are still barriers for the dissemination and display of 3D content. This is due to the large files size and high polygon counts of the 3D datasets, and in some cases the lack of surface material. The file size and polygon count can be reduced by using decimation or simplification techniques, and a UV map can be created for the cultural heritage artefact. These approaches may be applied in an arbitrary fashion, resulting in visual errors and degradation being applied to the 3D dataset itself. There are no guidelines or information available to inform

institutions on the best methods or how to create a perceptually acceptable 3D dataset to display and disseminate.

1.5 Research aims and objectives

The aim of this research is to encourage cultural institutions to disseminate and display their high resolution 3D digital cultural heritage, in a fit for purpose manner offering the best possible experience for end users, while being informed by the research. The research in this thesis aims to create a framework, which informs cultural institutions about their rights concerning their 3D artefacts, solutions for dissemination and how to create a good perceptual experience for users interacting with 3D content. To achieve this, the following objectives are investigated in this thesis:

1. Identify the legal and IP rights of scanned cultural heritage artefacts, which are afflicted with copyright, or lie within the public domain.
2. Examine how IP has impacted the dissemination of other media forms, and their solutions for dissemination and sharing their works.
3. Undertake an in depth literature review of the current state of dissemination and sharing of 3D cultural artefacts, to understand both technical and user needs.
4. Conduct a user study to understand a user's needs when designing a Graphical User Interface (GUI) and how they would wish to interact with 3D models.
5. Undertake a literature review of possible ways to apply a material to a 3D model, and state of the art research on how to assess visual quality of 3D meshes.
6. Conduct a quantitative user study, to identify and evaluate the relationship between texture and polygonal resolutions and how it may affect the perceived quality of 3D cultural heritage artefacts.
7. Examine the effectiveness of textured 3D models versus non textured 3D models to assess how high level of detail geometry compares to the surface material.
8. Assess whether the interaction with 3D cultural content alone is effective enough to encourage further interaction and exploration of cultural institutions collections.

1.6 Contributions to knowledge and structure of the thesis

Finally, this thesis makes contributions and expands the current knowledge for cultural heritage:

- **Chapter 2** provides in depth knowledge regarding the legal and IP rights of files created from laser scanning of cultural heritage objects. This knowledge includes both the legal state of the created file that is derived from objects that are utilitarian in purpose, within the public domain and derivatives created from the 3D file. Cultural institution's rights are also discussed with regards to the created file and how they may handle freedom of information requests regarding the created 3D files. Lastly information is provided regarding how IP issues may impact sharing and dissemination with possible solutions. This is beneficial for the cultural heritage sector, as it allows them to understand their rights and how it may impact on their plans to disseminate and share their 3D objects. This chapter will investigate and answer number 1 and 2 of the objectives list.

- **Chapter 3** provides results from a small Human Computer Interaction (HCI) study, to help understand user's needs and how users would expect to interact with small to medium 3D cultural heritage artefacts within a webpage or mobile device. This is very important in new and novel environments, as if the navigation is too complicated, it could put users off interacting with the 3D content and change their perceptions of the presented 3D content. The results themselves were surprising as it showed that users preferred to interact with 3D content with similar controls that are found within 3D software packages. This knowledge allows other cultural institutions to implement their own interactive viewers without the need to explore different navigation and interaction styles. Objectives 3 and 4 are covered in this chapter.

- **Chapter 4** provides an in depth literature review regarding the ability for cultural heritage institutions to be able to apply materials and assess the visual quality of the 3D cultural artefact meshes. The knowledge includes state of the art techniques for surface parameterisation and solid texturing to apply a material to a 3D dataset. Surface parameterisation allows for the automatic creation of a UV map to apply 2D textures. Solid texturing allows for a texture to be synthesised in 3D and applied to the 3D dataset, to make it appear as though it was carved from the material. Lastly information is provided that allows cultural institutions to assess the visual quality of 3D meshes. This is important and beneficial for cultural heritage, as it allows them to manage and reduce the models they wish to disseminate without losing the visual appeal of the full 3D dataset. How visitors perceive a model can be just as important in the dissemination of 3D cultural content the accuracy and tools provided by the cultural institution. This chapter provides the knowledge to satisfy objective 5.

- **Chapter 5** provides research and knowledge to indicate an acceptable cost to the polygonal and texture resolution of the 3D cultural heritage artefact that offers the best perceptual experience through a large quality assessment study. The study was performed with naïve museum visitors investigating how differing levels of texture and polygonal resolutions affected their perceptual experience of 3D cultural heritage artefacts. Key components of this study included the first study, to the best of my knowledge, that compared non textured 3D models against textured models, a comparison study comparing differing levels of polygonal and texture resolutions, the creation of a 3D stimuli through the use of a 2D image metric and the comparison of this stimuli versus the original artefact or a replica. The research presented in this thesis will be useful to cultural heritage institutions addressing the issues they may face when they wish to both disseminate and visualise their 3D content without reducing the overall perceptual quality. Number 6, 7 and 8 of the aims and objectives are investigated and researched within this chapter, providing the knowledge to satisfy these aims.

- **Chapter 6** presents the conclusions and future work. It reflects on the contribution and impact of the research conducted in this thesis and the opportunities for future work. This is followed by the references and the attached appendices.

Chapter 2 Legal, Ethical and Intellectual Property

2.1 Introduction

Cultural institutions such as museums are moving towards a paradigm of open content and access [19], where they wish to enable education through openly sharing information with the general public. This information would be freely available within the institutions or on an online platform; allowing visitors to access their collections, exhibits, photographs, archive and other works. This adoption of open content has seen some success within the UK [1, 2, 20], yet there are still some problems with regards to the dissemination and use of this content, the main concern being ownership of the IP (IP) under UK law. There are many licensing options available for openly using and sharing information such as the creative commons licenses [21], but a museum itself is fraught with IP perils. A museum is not considered an IP intensive industry but it is in the odd position of being both a licensor and licensee of IP [11]. There is a conflict between these different approaches which seem to be mutually exclusive, with little to no way of them coexisting together [19]. This is further complicated when you consider 3D datasets of cultural heritage objects.

The last decade in cultural heritage has seen a large adoption of technologies that can create high resolution 3D digital models of cultural heritage artefacts, which can be used to create digital or physical replicas of the original artefact [4]. With regards to the newly created 3D datasets, which include public domain works; should cultural institutions be allowed to obtain IP protection for these datasets, even if they are of public domain objects? What of objects that are utilitarian in purpose, would they be protected under IP laws? If so, what implications may this have? How will it impact the dissemination and sharing of information?

Access to digital cultural resources for education, inspiration for artists and visitors, and sharing culture is a key mission for many museums within the UK and internationally. Their mission empowers their communities, letting them learn about past civilisations, but this can be hampered by legal, ethical and IP issues. This chapter addresses these issues, which are currently a major concern for cultural institutions, wishing to open their collections to open access. This chapter focuses on the application of IP laws within the UK to 3D digital cultural heritage artefacts. The application of copyright or other IP rights for cultural heritage artefacts internationally lie outside the scope of this thesis. The chapter will then discuss if any new works derived from the created 3D dataset would warrant IP protection under UK law. The last section discusses the impact of IP law and the effect it can have on the dissemination of 3D datasets.

2.2 Background

2.2.1 Legal and Ethical Issues

Works and objects that are donated to a cultural heritage institution collection will most likely relinquish their rights to that work and assign it to the institution; which can include IP rights. This gives the cultural institutions the right to:

- Secure contracts and govern how future photographs taken of the work, may be used by third parties.

- Restrict the access to this object for the purposes of photography or laser scanning by the general public
- License the images to third parties [22].

This can lead to legal and ethical issues surrounding the scanning of these works within the museum, to create a 3D work, for both the cultural institution and its use. This is due in part to the novelty of being able to create 3D works, where most institutions may treat it similarly to the digitisation to 2D images with regards to legal or ethical issues. This is not the case as a 3D dataset can used for a range of different applications [19], and a range of different objects can be scanned that have different purposes sculpture, artwork, traditional expression, utilitarian purpose.

The assigning of copyright for the created 3D dataset from laser scanning or other means is a common issue for cultural institutions. Due to the high costs of laser scanning equipment it is likely scanning is commissioned by a third party, which has an impact on the ownership of the copyright of the digital dataset. As described in the Copyright, Design and Patents Act 1988 (CDPA), section 11, the Author of the work would be the copyright owner not the commissioning party [23]. The museum would need to secure all rights to the new files created via a contract assigning all rights to the museum or acquiring a license to use the 3D file. So this in a sense would allow the museum to protect its rights to the digital 3D dataset but it would also mean members of the general public could scan a work and claim copyright on the digital file but not for the original object.

Another issue would be the creation of a 3D work from a traditional cultural expression, which would be an object or item held with cultural institutions collections, which a culture may consider sacred to them, such as religious artefacts [24]. It could also be the case that these items and objects may have originally been taken from a culture without the consent of the community and this may lead to complex issues. Due to the complex nature of traditional expressions, their legal status is unclear under the law, especially under IP laws, where the expression may be within the public domain, and usable by all, but could be considered sacrilegious within the cultures community [24]. More information can be found within [24] by the World Intellectual Property Office.

2.2.2 Intellectual Property

In the UK, IP is an umbrella term that covers primarily Copyright, Patents, Design and Trademark, which confer to the author of the work certain exclusive rights regarding how the work may be used and distributed for a set amount of time. These rights also allow the author to dictate how others may use their work, be attributed for their work and it is also possible to transfer these rights to other individuals or corporations with the exception of moral rights within copyright [25].

The use of Patent law or Trademark is unlikely to affect the dissemination and sharing of 3D digital cultural heritage artefacts, or be infringed upon. Patent law is not applicable, as it covers an invention that is “new, involves an inventive step and is capable of industrial application” [26]. Trademark IP laws are also just as unlikely to apply, as they cover the broad scope of signs, design or graphical representations that identifies a product or service from a particular source from others of the same type [27]. However, these may not be the case for more modern works, which will be covered by copyright law [28], but may also have incorporated a Trademark within the artistic work, or the creator may have IP protection for a design right or a patent. Other issues that may arise include; unauthorised use of the trademark for commercial use, and if it is used in a manner that falsely gives the impression that there is a connection between the seller and the trademark holder [29]. The only issue that could possibly arise would be the use of a trademark being embedded within a 3D dataset and sold without permission of the trademark holder. Design Law covers the protection of the shape or configuration of a marketable product [30]. This is unlikely to apply, due to cultural artefacts being made marketable but possibly could apply to 3D printing. However this is an extremely complex area of law and lies out of the scope of this chapter, but is discussed in great depth by Dinusha Mendis [31] with regards to 3D printing. However, it is likely that Copyright will have the most impact on the dissemination and sharing of 3D cultural heritage artefacts.

2.3 Copyright

Copyright within the UK is a right conferred to an author to protect the expression of their idea, if it is original and has needed substantial degree of skill to create. This covers artistic, literary, dramatic, musical, films, sound recordings, broadcasts and published editions (typographical works). From the moment of its creation, copyright subsists in the new work [28] and will do so for the life of the author and 70 years after their death. The author may also seek damages if an individual or corporation has used the work without permission or a license from the author. There are exceptions within UK law for fair use, which include: private research or personal study, education, criticism, reviewing or reporting [32] and more recently parody [33].

2.3.1 Copyright and Cultural Heritage Objects

Copyright is not supposed to exist in a work that is within the public domain or for objects that are considered useful, yet this is not often the case, such as the Museums licensing photographs of public domain paintings [34]. This raises concerns for 3D datasets created from objects within a cultural institution's collections, especially if a piece is within the public domain, a utilitarian object or a traditional cultural expression. Works that are within the public domain can be exploited by anyone for free without permission or license. However, this is a contentious issue currently as institutions may generate revenue though worldwide licensing agreements for the use of photographs of paintings and other artworks that clearly are in the public domain [35, 36].

This is further complicated when you consider a 3D dataset, which is created first through the scanning of an object by a technician, but then post processed to create a digital file, which could be used to recreate an accurate replica of the physical object. Under current UK copyright law [37], a work warrants copyright if it is deemed to be original but has also been created through a degree of labour and have both skill and judgement applied during the creation of the expression of the work [37]. Copyright does not exist in a piece that was created through an entirely mechanical process or is a slavish reproduction [18]. This would apply then that any images taken of a work that was within the public domain, to be a "slavish" reproduction and not warrant copyright protection.

The USA court case of Bridgeman Art library, Ltd. V. Corel Corp [18] addressed the use of 2D photographs of works within the public domain. The case held the decision that the digitisation of works within the public domain did not warrant new copyright as the work was not original and a slavish reproduction of the original public domain art. Even if a substantial amount of skill and effort was used in the reproduction of the piece to a new medium, as noted in a previous case "Sweat of the Brow" alone is not the "creative spark" which is the sine qua non of originality" [38]. The court did argue that even under UK Copyright law; copyright would not subsist in the new work. However, the outcome of the case is not binding in the UK, but the significance of it has led to many debates regarding the decision. Within the UK, the Museums Copyright Group commissioned a report on the outcomes of the case [39] and a seminar at the Queen Mary University of London where in attendance art professionals, IP lawyers and other creative individuals decided that the decision should be reversed [40]. However, this may have been more recently contested within the UK with cases such as the Temple Island V. New English Teas and Nicholas John Houghton, where the originality for copyright in photography is defined where "taking of the photograph leaves ample room for an individual arrangement" [41], where images of public domain paintings and art could be considered slavish reproductions with no room for individual arrangements and warrant no copyright protection. This has been further strengthened by the UK Intellectual Property office, stating that copyright will not be applied to a replication of an older image that is within the public domain [42]. They state even if additions are added to the public domain image, to remove stains and improve upon the image, it still will not warrant copyright protection [42].

However, in the UK though the CDPA states that copyright will be granted if the work is original and considerable skill and judgment was used in the creation of the expression, courts have not adopted a literal reading of the law. UK courts have adopted a stance that if a reasonable amount of skill and judgement in the creation of an expression is acceptable to warrant copyright protection [43]. Yet the copyright implications of scanning 3D cultural heritage objects that lie within the public domain are ambiguous. Using the reasoning of Temple Island vs. New English Teas [41], it should not be possible to get copyright protection in the design document that is created. The idea of granting copyright protection to recreated artistic works due to the skill and judgment that take place to create the file is supported by Ong [44]. He justifies that it could be in the best interest of the public for replicas to be recreated of works that are of cultural importance or are not readily available to the public [39]. This is further supported by cases of Antiquesportfolio [17] and Painer [45]. The case of Antiquesportfolio held that there can be copyright in the protection of photographs of 3D antiques, due to the angle of the camera, lighting and camera focus which led to exhibiting features of the antique such as colour and details on the item. Due to the skill being shown in these photographs, it was concluded that it did warrant copyright protection even though these skills were at a basic level [17]. A similar conclusion was reached in the Painer case, which focused on the copyright protection for works based on reality such as portrait photos [45]. Also in an ironic twist, the UK Intellectual Property Office, which stated digitised works do not warrant copyright protection, applauds the licensing practice that the National Portrait Gallery of public domain works [34]. This blog post on the UK Intellectual Property Office website has muddied the water in regards to licensing public domain pieces. However, they have left a comment on the page illustrating that it was written by a third party and it does not represent their views [46].

The decision of the Bridgeman Art library, Ltd. V. Corel Corp has been applied to the 3D recreation of a physical object in the US court case of Meshwerks, Inc. V. Toyota Motor Sales [47], where it was decided that a replication would not warrant copyright [47]. However, as pointed out by Michael Weinberg [48], that this may change in USA law if it grows to recognise the artistry in the artistic process for the capture and the creation of a 3D dataset via postprocessing, which is recognised in photography [49]. During the capture, post processing of the data and creation of the digital design file from scanning, it could be argued that the skill and creativity shown, in selecting certain views, lighting conditions and processing the data, and in its creation, would make it an "intellectual creation of the author reflecting his personality and expressing his free and creative choices" [45]. Taking all of this into account, the copyright protection of a created dataset is still rather ambiguous and it will need to be settled with a future court case within the UK or the EU.

However, it is clear that if a 3D dataset were created from an artistic work that is still protected by copyright, it would be considered a derivative work of the original and subsequently would infringe copyright [37]. However, Bradshaw [50] identified 2 possible cases where copyright may be eroded for a physical copy: a) the definition of the object produced, and b) the digital file as a design file.

2.3.2 Definition of an Object

Copyright may subsist in the digital file, and its subsequent derivative works including 3D printed models. However, there is an issue with the classification of the 3D printed object in itself. It could be assumed that as the dataset is based on an artwork, it could be defined as a sculpture as defined in section 4 of the CDPA 1988 [37]. Yet within court cases the definition of a sculpture, has always been a somewhat difficult area to define, with various rulings attempting to clarify what a sculpture is [51]. This led to a judgement in Lucasfilm Ltd. & Ors V. Ainsworth Anor [52] that a work would qualify as a sculpture if the object had an "intrinsic quality of being intended to be enjoyed as a visual thing" even if it had other purposes [37]. LucasFilms ltd. tried to assert that the original clay model of the Storm Trooper helmet would be classified as an artistic work under the CDPA 1988 [37]. However, the courts upheld that the clay model had a utilitarian purpose, as it is a helmet first and foremost, therefore there could be no copyright infringement under section 51 and 52 of the CDPA [37]. While artistic works will be protected, this ruling lends itself to 3D printing of utilitarian objects as section 51 of CDPA states:

"(1) It is not an infringement of any copyright in a design document or model recording or embodying a design for anything other than an artistic work or a typeface to make an article to the design or to copy an article made to the design.

(2) Nor is it an infringement of the copyright to issue to the public, or include in a film, anything the making of which was, by virtue of subsection (1), not an infringement of that copyright. [37]"

In Section 51 of the CDPA 1988 a design document is specified as:

"Design document" means any record of a design, whether in the form of a drawing, a written description, a photograph, data stored in a computer or otherwise [37].

Where in section 51, an item produced via a design document will not result in copyright infringement [37]. This was the case in Lucasfilm Ltd. & Ors V. Ainsworth Anor, where the original artwork for the helmets was designated a design document for the creation of the 3D helmets. This would raise many implications for IP and is answered by Bradshaw et al. [50], but section 51 expresses that artistic works are exempt from this protection. However, in section 52 of the CDPA, originally if artistic works were used for mass production for the generation of revenue, the copyright on the work would be severely reduced to 25 years [37] but this has recently been changed to reflect copyright in other sectors, and retain the copyright for the life of the creator plus 70 years [53, 54].

2.3.3 Other Protections

While a museum may be able to have ownership of the 3D dataset, it is still expected to use the digital file to share the knowledge and use of for educational purposes in a public forum. However, as explained by Koller et al. [55] a museum may be afraid too, as they would lose control over how the object would be represented and there is a chance of the model being pirated if the model was disseminated over the internet [55]. To combat unauthorised distribution there are tools available which are codified in the US Digital Millennium Copyright Act 1998 [56]. Michael Weinberg explains briefly how this is an acceptable tool to stop the online distribution of infringing content [48], but it does not allow a museum to maintain control over its IP fully.

Also of note is the ownership of copyright of the 3D dataset. As described in the CDPA section 11 [57], the author/creator of the work would be the copyright owner not the commissioning party [57]. The institutions would need to secure all rights to the new files created via a contract assigning all rights to the cultural institutions or acquiring a license to use the 3D file.

A final important fact to note is that many museums and cultural institutes within the UK can be subject to a Freedom of Information Act request [58]. This allows a member of public to make a request for and to use certain information held by a public institution, body or department. If a member of public makes a request to the department, they are entitled to:

"(a) To be informed in writing by the public authority whether it holds information of the description specified in the request, and,

(b) If that is the case to have that information communicated to him. [58]"

I can be assumed that the 3D dataset would be available under a Freedom of Information request, but this is not the case. There are exemptions and clauses that protect the public body to decline the request for information, two clauses that may be used to protect the datasets are:

Commercial interests:

"Information is exempt information if its disclosure under this Act would, or would be likely to; prejudice the commercial interests of any person (including the public authority holding it). [58]"

Law Enforcement:

"The purpose of protecting the property of public authorities from loss or misapplication [58]"

This same protection could be applied to stop the general public who wish to attempt to create a 3D dataset from a sculpture or object. The museum is not required to grant that person access to that work for documentation, and could ask the person to leave. This can also be applied to both public and private institutions.

2.4 Impact on Dissemination

While the legal clarity surrounding IP protection for these datasets is ambiguous, it opens up new areas of innovation, dissemination and potential revenue for cultural institutions. There is an ever growing market for this type of 3D content especially in 3D printing, which was valued at $2.2 billion in 2013 by Wohlers Associates [59] and this is figure is still expected to grow in the coming years. Yet, due to the high value these datasets possibly command, it is worth considering how these datasets can be protected and possibly monetized. Especially when considering sharing this information on the internet, where duplicating the dataset and sharing it with other users is extremely easy without informing the original owner. There may also be concerns regarding piracy and loss of control of data, yet it should be considered that the gains of openly sharing these collections will always outweigh the perceived loss [35].

The possibilities and opportunities that a 3D data set offer, will invariably attract users who wish may to download or replicate one of the datasets protected by IP law. As such there are methods that can be used for the protection of the datasets, wither this is in partnership with a 3rd party platform for 3D printing or through the use of a licensing for the actual dataset file. There are systems that have been developed for interaction and visualisation of cultural heritage models, without the user downloading the actual dataset.

Koller et al. [60, 61] created a visualisation system called ScanView to enable users to interact with 3D models created during the Digital Michelangelo project [62] while protecting the 3D dataset. They created a portal that would allow users to interact with a low resolution model, where they used a combination of technique such as remote rendering of the full resolution 3D dataset with subtle distortions to the model. This closed designed system meant no datasets where downloaded and it was not possible to recreate a dataset from extracted images. Yet while the system was acceptable, the use of the distortion and noise distracted the end user from the experience of interacting with a dataset [63] and hindered in some instances of interaction. Similar solutions have been launched within other industries but it is now possible to use the internet to disseminate and engage with 3D models directly [64, 65]. The internet also allows for a subtle form of protection commonly referred to an always "always on" system that requires users to connect to a server to authenticate the use of the work [66, 67]. While this approach is not as restrictive as the likes of ScanView, it presents unique problems especially when content is only available when the server is online [68, 69].

There is an alternative to these types of systems, which is the adoption of a traditional licensing, which is commonly used for photographs of cultural content. This approach allows a cultural institution the possibility of generating revenue through licensing agreements or allowing open access through a Creative Commons license [21]. There have been successful attempts at the sharing and printing of 3D content, both from video games such as Second Life [70] and even more successfully by FigurePrints [71]. Where they obtained a license from Blizzard, which allowed them to print gamers in game characters from World of Warcraft in full colour and have them delivered to the customer [71]. There has even been some success with cultural heritage institutions using online platforms to disseminate 3D content such as The

British Museum using SketchFab [72] and The Metropolitan Museum of Art [73]; allowing users to download low resolution datasets, and print these objects from either museum's collection under a creative commons license.

The use of licensing has advantages over of a closed loop system, which in some cases has been shown to actually lead consumers to piracy, as they feel victimised by restrictions and the lack of freedom to use the content for other purposes [74]. There is other research that points to the ineffectiveness of these systems to deter piracy and can actually drive users to restriction free content [75]. A greater surprise is that companies that have removed restrictions on their content have not seen a loss of revenue due to piracy [76] and in some cases actually increased profit [77].

While many institutions may worry about piracy and the loss of control over their works or 3D datasets, there have been many debates regarding monetary damage from illegally sharing files. In 2011 Motion Pictures Association of America (MPAA) claimed that the US economy lost $58 billion due to piracy [78], but the figure itself was extremely exaggerated due to issues such as counting every illegal download as a lost sale [79] and assuming adults would buy an additional 200 DVDs a year [80]. This topic is highly contentious both in the estimations of damage and how to effectively combat piracy but it appears to have little effect on 3D content.

The Pirate Bay is a platform that has been blocked in the UK and other countries due to allowing users to share illegal copyrighted files [81]. In January of 2012 the Pirate Bay introduced "Physibles", a section to share 3D datasets for 3D printing [82]. Yet the section contains very few files in comparison to the large amount of torrents available to download [81]. With such a low number, it can only be assumed there is not a lot of interest in either uploading or downloading of 3D content from these types of sites.

It is more likely for 3D content to be hosted on 3D online platforms such as Sketchpad or Thingiverse, which provide an interesting insight in to their users, the most popular categories for 3D models and the impact of sharing 3D content. These sites are legal file sharing platforms that offer the ability to interact with the 3D content in the browser and encourage the sharing of 3D content under a creative commons license as well [83]. These sites host a vast range of 3D content from spare parts, games characters, cultural heritage items and an assortment of other categories. As well as these platforms, there are others that allow users to buy 3D datasets to print at home that have been created by designers, companies or amateurs. One of the largest sites is Shapeways, which allows designers and companies to create shops on their platform to sell their designs, which are then printed by Shapeways in a range of different materials. This secures the access to the digital file, as it is only shared with the platform and also protects users who may get a defective purchase where blame would be laid with Shapeways. These sites also take on the burden of possible copyright infringement away from the designer or company. If Shapeways was to print a piece that was infringing copyright, it is likely that Shapeways would be the infringer for creating a derivative work from the original. This has forced Shapeways to check for IP infringement before proceeding to print any work, and if any infringement is suspected they will not print the piece [84]. This policy is being used by other similar platforms, although it can be subject to human error it is reducing the amount of infringement that is taking place on these sites.

2.5 Discussion

A review of the current law throughout this chapter highlights the complex nature of IP issues, in regards to cultural heritage and recreating either images or 3D models of objects within their collections. While using a 3D model for a range of purposes may not infringe copyright for a number of reasons [32], it could infringe under one of the other branches of IP law [31].

Though what is apparent is that if a 3D model were created of an object that still has copyright applied to it without permission, it would be a copyright infringement. Any derivative work from this created 3D file would also be an infringement [57], if it is an artistic work which may warrant copyright protection. This protection may not extend to all objects within a cultural heritage collections, it should be noted that objects that are utilitarian in purpose are not protected [37].

The complex nature of IP laws leads to an ambiguous answer for the question "Is it possible for a cultural institution to hold the copyright for a 3D file of an object that is within the public domain?" The answer itself could be either yes or no depending on how the object was created. Was the 3D model created from a work that is to be enjoyed visually and created with artistry similar to photography to warrant copyright protection [37]? This could be the case but it is not clear cut, especially in the UK where the Intellectual Property Office, has said that it is not possible to warrant copyright protection on public domain objects which may have been modified slightly [42] and the supporting evidence of Bridgeman Art library, Ltd. V. Corel Corp [18] and Temple Island V. New English Teas and Nicholas John Houghton [41]. This will be a highly contested area of debate for cultural institutions, which will need to be resolved in a court of law. The implications of such a case may help to provide a decisive answer. However, this may have further implications within cultural institutions that may rely on money being generated by images taken of their objects that are within the public domain.

There is also an important question that would need to be raised in regards to how IP law would be applied to a famous work of painting. Depending on how the 3D model is created, it will either be without color and only the surface details of the frame and the painting, or have the color information as well. These two things are very separable from each other, with a separate 3D model containing just the frame and canvas surface details and a separate texture file. The 3D frame and canvas would possibly be identified as being utilitarian objects, where the texture file IP status is ambiguous and would need to be addressed in a court of law too. Though assuming that the texture file did get IP protection, what would happen if someone obtained the 3D file and used it for derivative works? According to Lucasfilm Ltd. & Ors V. Ainsworth Anor [52], this 3D file could be used for derivative works without infringing, but could not be used with the texture file without infringing copyright of the texture file. This may not be the case, but it is one of the ambiguities IP laws that would need to be addressed with further research or within a court of law.

For the dissemination of these 3D works via the internet or 3D printing, there are many approaches that could be used, including licensing [21], DRM [67], open access [1, 2, 20]. These allow for a cultural institution to disseminate, be attributed, monetize [75, 76] and control their 3D work. These can offer flexible options, regarding the amount of control that a cultural institution may exert over their 3D objects, but there is always a chance that depending on the option used, it could increase and push people towards piracy, as they feel they are being restricted and punished. In the case of files being pirated, there are tools such as the DMCA [56] that can be used to remove pirated files that may be found on websites. In many cases it is shown by lessening the control and allows the user to use the platform in a way that they want to use it, can lessen piracy and increase revenue. These business models seem to have evolved for new and innovative business strategies that have changed focus from stopping piracy to actually competing with piracy with an easy to access platform.

Bringing this all together, it shows that IP issues are complex and are likely will be an issue regarding cultural heritage as well as other industries and uses of 3D models from use for the promotion of collections to 3D printing. Yet for an institution that wishes to create its own open platform for sharing and educating people within their galleries or websites, this chapter has addressed some of the issues that may plague them. While it is ambiguous in if it is possible for IP to reside in an 3D model file created from a public domain object, some issues that may cause concern has been addressed in some form if they wish to openly share the 3D file and derivative works.

2.6 Conclusion

The creation of extremely accurate data sets of cultural heritage objects has been very beneficial revealing hidden information and allowing users to interact with these objects in new and exciting ways. However, with these benefits issues have been raised concerning the IP issues while increasing the access to these 3D datasets. This is becoming increasingly more important as the technology to scan and replicate objects is heading towards the domestic market.

This chapter has attempted to explain how copyright under UK law may be applied to a cultural heritage object, even if it lies within the public domain. As discussed in this chapter, it should be possible to acquire copyright with scanning, but a test case would be needed to test this hypothesis within the EU or UK. If copyright subsisted in the digital files, any derivative works that would be created would also warrant copyright protection unless it was of a utilitarian object or exempt for fair use.

This chapter has also highlighted the possible ways to disseminate and share the datasets via websites through the use of licenses or a closed system. Where it would be possible to allow the datasets to be shared openly through a Creative Commons licenses and possibly generate revenue through various business options including 3D printing. Yet it should not be under taken lightly as there will be risks involved and possible repercussions for the cultural institution. Though there are risks involved, the benefits of sharing cultural heritage content will always outweigh the risks.

Chapter 3 Dissemination and Interaction of 3D cultural artefacts via the web

3.1 Introduction

The use of non-contact laser scanning first originated in the automotive and aeronautical industries [85, 86], where its use in reverse engineering led to its adoption in cultural heritage for 3D documentation [4, 61, 87, 88]. This 3D documentation technology offers the potential for new and exciting experiences, for visitor and researchers to interact with artefacts that are too large, or that are too damaged to be displayed or handled [4]. Yet the novel nature of the 3D datasets is presenting problems affecting their accessibility and dissemination due to a number of reasons such as their file size, and visualisation. However, with new technological advances, and improvements in internet speed, we have reached the point where it is possible to interact with 3D models directly within websites. There has been an increased use of 3D content within the last couple of years [64, 89, 90, 91], and this is piquing the interest of cultural heritage institutions. This level of engagement is very important for cultural institutions, especially when the internet is becoming ubiquitous in the lives of their visitors [92] and their wish to connect to as wide and varied an audience as possible. Moreover the way in which the world engages with information is changing - we are relying more and more on the internet to provide the information we seek [93, 94]. This ubiquity and shift in demographic engagement also means that cultural heritage institutions, museums in particular are finding that their websites are becoming digital venues, gallery spaces, education forums, and (digital) destinations for visitors in their own right.

However, the display of 3D cultural content via new mediums such as the internet is raising challenges not just due to their file size. There are issues such as designing for an anonymous demographic of website visitors, providing navigation and controls that are suitable for non 3D specialists and providing the appropriate information regarding the cultural heritage artefact. There is a need to pursue research that would allow cultural institutions to fully communicate the significance of their 3D cultural content to their physical or virtual visitors.

The work undertaken by the author of this thesis during this chapter, focused on exploring the creation of an interactive prototype 3D interactive WebGL based viewer to display 3D cultural models within the NML website. The prototype was developed in 2012, however other web based interactive viewers have surpassed it such as those provided by Sketchfab [95]. This chapter examines the processes that were explored to understand visitor's needs, and expectations when interacting with 3D cultural content, focusing on usability and interaction. The chapter presents research methodologies for low-level prototyping, highlighting the benefits of these methodologies for cultural heritage institutions. This chapter also contributes results from a small HCI study, which focused on how users expected to interact with 3D cultural content within the interactive viewer, exploring a range of different navigation styles. However, it should be noted that although the user testing conducted in this chapter is to understand the needs and expectations of visitors interacting with 3D digital cultural heritage; the results themselves may not translate into a successful application. There are many complex social interactions, and complex anomalies that may not be quantifiable by HCI methodologies [96], which may hinder the success of an application, especially with a small user study as the base for the final application.

The chapter discusses recent work on interactive viewers in the wider world and within cultural heritage. It also addresses the importance of creating a relationship with visitors and understanding their needs which can help dissemination of information and open cultural institutions to new audiences. It will then discuss the prototyping of the interactive viewer and how user testing was conducted with a low-level and developed prototype. The final section will discuss the impact of this work undertaken within the NML and how this has influenced the use of 3D datasets within the website and gallery spaces.

3.2 Background and related work on interactive viewers and web based approaches

Interactive viewers to display and interact with 3D models have been popular within computer vision for decades, originally displaying 2D renders of 3D models before progressing to real time interaction. A popular implementation is the use of image based solutions, which takes images from different viewpoints of an object which are then 'assembled' to provide a '3D like' interaction. This approach can offer both realistic and unrealistic visualisations (figure 1) of a 3D dataset depending on what context the model will be used. This approach can also allow for multi-platform access to the content, a current problem for some technologies in current use on the internet [97]. These solutions are often presented in Adobe Flash, Quicktime VR and JavaScript. However, viewing of the model is restricted by the number of images that are being used in the interactive.

There are other image based approaches such as Light Fields [98], which allows users to choose between different viewpoints of a model, but this approach suffers from very large file sizes and is not suitable for use on the internet. A more efficient approach is the use of Surface Light Fields [99], which uses additional scanned geometry of the object alongside images. This approach allows users to interact and navigate around an object in a traditional manner but once again can be limited by large data file sizes.

A more recent approach is the use of Polynomial Texture Maps (PTM) or Reflectance Transformation Images. This provides a fixed viewpoint of the object, and allows the user to change the lighting parameters to fully understand how light interacts with the object. Mudge et al. [100] presented a PTM viewer that combined laser scanned datasets with multiple PTM data, to allow users to explore a site and take full advantage of the PTM data. However, due to the file size of the scanned datasets and the PTM, a large amount of storage is needed for each view point.

3.2.1 Choice of Interactive Technology

However to display 3D content within a web browser, that would work across platforms and within different browsers, two technologies were investigated at the start of this study, to identify a possible technology to base the viewer on. These two technologies were Adobe's Flash Player and HTML5 and WebGL. Both solutions are widely available; Adobe's Flash player is a widely available and distributed plugin that is used for displaying media or interacting with Rich Internet Application within browsers. It is a free download that is supported across all major platforms [101] and in 2009 reached over 100 million installations on different platforms [102] and has a large online community. HTML 5 is a new and emerging technology first proposed in 2004 by the World Wide Web Consortium (W3C) [103] but came into the lime light in 2010 due to Apple choosing it over flash to be used on iPhones [104]. It is now, however, supported in all major browsers with the WebGL API supported on the web kit browsers Google Chrome, Mozilla Firefox, Opera and Apple's Safari [105]. WebGL is a 3D graphics API based on OpenGL ES 2.0 that allows 3D models to be directly embedded within a webpage [97]. Both solutions can be used to target a wide and varied audience but currently Flash Player is supported on all browsers apart from Apple iPhones and iPads, with Adobe has issuing a statement that Flash will no longer be supported on Mobile browsers [106] and more recently even Linux, where the flash player will be supported only by an API used within Google Chrome [107].This decline has become even more apparent with a recent statement from Adobe, stating they will stop the distribution and updates for Flash by 2020 [108].

Both are capable of creating 3D interactive content but they both need a supported graphics card [109, 110], to view 3D content in Flash Player it is dependent on Adobe Flash Player 11 and the Adobe Stage3D API [109]. Stage 3D is capable of rendering advance 2D and 3D content by exploiting the power of the GPU [109]. HTML 5 relies on browsers supporting HTML 5 features especially the canvas element and WebGL API, but does not need a plugin to work. HTML 5 relies solely on WebGL a low-level 3D Graphics API based on OpenGL ES 2.0 [97]. WebGL is currently capable of supporting models with 1,000,000 polygons [111] and is currently pushing this number upwards with a leading JavaScript 3D engine Three.js. For computers that do not have a supported GPU, Adobe Stage 3D can offer a "software mode" [109] renderer, but it is a lot slower and displays a lot less frames per second. WebGL can be rendered in Google Chrome due to it using ANGEL to process WebGL into Direct X, improving the quality of WebGL content and allowing the content to appear on computers that do not support WebGL [112]. In a comparison test created by Airtight Interactions [113], results were very similar but none were decisive enough to be a clear winner. Both are extremely well suited to be used for development to the prototype, but both currently do not work across all platforms very well.

HTML 5 and WebGL were chosen for the development, as additional media can be used to support the application and enhance the interaction and visual appearance of the 3D dataset.

WebGL is also becoming relatively popular in its use in the cultural heritage field to display 3D objects [65, 89, 90, 91, 115, 116, 117, 118, 119]. The Smithsonian museum [89] and Europena [90], use WebGL to display parts of their collections online, accessible to all. X3DOM [65, 91], was created with cultural heritage in mind, to allow institutions to integrate their 3D content within their webpages or for mobile devices [65]. Schwartz et al.[115] provided research to fully visualise cultural heritage assets by combing bi-directional texture functions with WebGL [115]. WebGL is also being used for scalable interactive visualisations [116, 119], and exploring in its use with VR within cultural heritage [117]. Research is also being conducted to weigh the benefits of the use of WebGL and 3D printing within cultural heritage [118].

Figure 1: 3D dataset of a sculpture of Artemis (NML) rendered with different shaders. Clockwise from top left: Blinn, Mental Ray stone shader, subsurface scattering shader and toon shader. On the right, a photo realistic rendering of a 3D bust.

As WebGL is becoming more stable, it is offering museums and cultural institutes new and exciting ways to connect with website visitors. However, the audience that WebGL affects is still limited. WebGL Stats [120] measures how well WebGL is established and currently supported, giving a figure of just over 66.8% from across the web in 2013 [120]. This has since improved to 92.6% of visitors (83.1% stable and 9.1% using an experimental release), covering various web browsers and devices [121]. However, a separate report by Koan Interactive in 2012 investigated further and as they revealed that out of 1,872 visitors to their website only 24% could use WebGL [122]. Interestingly, they did report that 50% of total visitors could support WebGL, although 26% of visitors had WebGL disabled for an unknown reason [122]. If 24% of people visiting a museum's or cultural institutions website are able to use WebGL, it would clearly be recommended to have supported media alongside the 3D content. In addition, during the development of WebGL, browsers have all reported instabilities with graphic cards supporting WebGL. Mozilla Firefox has created a blacklist of graphic drivers [123] and other browsers may selectively choose to enable or disable WebGL depending on certain situations [124].

One of the popular frame works for cultural heritage is X3DOM [65], which have been used in collaboration with various institutions, including the Victoria and Albert museum, London [125], Europeana [90] and the European project, 3D –Coform [126]. This framework is based on the document integration model and integrates an X3D model into a webpage. The syntax to use X3DOM is based on XML, so it is possible for non-programmers to implement this framework, without having to learn JavaScript. As WebGL is only supported on webkit browsers, it uses Adobe Flash to allow it to work on multiple platforms [65]. However if a model with a high resolution is used, the Flash implementation has been known to routinely crash.

In an attempt to recreate the full visual appeal of cultural heritage artefacts, Schwartz et al. [115] combined WebGL with a Bidirectional Texture Function (BTF) [115]. BTF's that uses the advantages of PTM's [100] and surface light fields [98] to create artificial lighting by describing the appearance of a surface depending on both the viewing and lighting direction. Primarily this technique is used for flat surfaces, but can be used with more complex surfaces, such as those found on many cultural heritage objects. The framework created by Schwarz et al. [115] recreated the full visual appeal of an object by interactively rendering 3D geometry, constantly improving upon surface colour and reflectance information. This is achieved via a novel progressive streaming and compression algorithm to transmit the large BTF datasets. The surface information is applied to the model over time, allowing for a smooth interaction without a large loading time for the subsequent BTF texture files.

While these techniques can enhance and add a level of depth for a viewer for digital cultural heritage objects, it should be noted that the way in which audiences engage with information and media is changing. There is no longer a culture of “gatekeepers” or “digital curator” providing information; the internet has empowered viewers to choose how they engage with information [12, 127].

3.2.2 User Interaction

It could be assumed that visitors to a cultural institution's website may have a passion for cultural heritage, or that they will be seeking information regarding the institution and its physical exhibition spaces and objects. However, this has been found to not always to be the case [15]. Consuming information and wishing to understand content has led to a social drive to share information with our peers, for a variety of purposes. This sharing and dissemination brings content to a new wide and diverse audience [15]. This socially based process can be rewarding and beneficial for both the target audience and the cultural institution hosting the content if they can collaborate to create a dialogue. A cultural institution is able to provide the tools to enable visitors to engage with their collections, creating a community based around their collections. This increases the likelihood of future website visits for new content and repeat visits in general, but it also promotes the collection, the physical cultural institution and its brand [128].

Turner [129] describes an engaging experience as something we would like to repeat as something fun, interesting and rewarding. We disengage when it becomes mundane and boring. While this description is simplistic, this statement merits consideration; to engage with a visitor (on-gallery or on-line) the content presented must be enjoyable or provide satisfaction in some capacity. Time also plays an important role in establishing a relationship with an audience [130]. It is not possible to establish a relationship with a visiting audience by only providing a short novel experience; relationships need to be allowed to mature over time. Short, novel experiences could damage the establishment of a relationship, as the engagement would be fleeting, and would offer little in the way of substance to encourage repeat visits. In such a scenario, visitors may well seek the similar content elsewhere, represented in a way that can better cater for their needs.

While the above implies that the best approach for a cultural institution to take with respect to its digital collection would be the creation of a highly involved interface to interact with a 3D digital object, this would actually be very impractical. A simple to use and light interactive that has regular content updates and a high quality output can help to establish a lasting and trusting relationship with visitors [15].

3.3 Creation of a 3D interactive viewer for National Museums Liverpool

The work undertaken by the author of this thesis focused on the creation of a prototype viewer that identified and analysed the needs of users engaging with cultural content and visualising high resolution cultural content. The interactive viewer also provided users with appropriate tools to help them engage directly with the 3D model. The user testing for this project was split into two phases. The first phase of the user testing focused on understanding the needs of users and their expectations of how a 3D environment may work in 3D. This was explored with the use of low-level prototyping, predominantly paper prototyping to allow users to create a GUI and show how they expected to interact with the cultural content within the interface.

However, as it is not possible to explore the design ideas of a 3D environment in a 2D setting [131], the second phase revolved around a computer-generated prototype. This was developed on a Windows 7 machine using HTML 5 and WebGL. WebGL was chosen for the creation of the interactive prototype due to its popularity within cultural heritage, and its ability to display high resolution 3D models as mentioned on page 45. User testing was conducted with 8 participants (2M/6F) to evaluate the controls and usability of the viewer. Users came from a range of backgrounds including curators, conservators and non-specialists within the NML. Participants had differing levels of computer experience, ranging from having only used desktop applications, to using 3D manipulation software on a daily basis. The results from this study however may be limited, due to the physical age of the participants, not reflecting the true age range of physical and digital visitors. There is also a concern for bias in the results due to the participant's professional connection to NML and not showing the disparity of visitor's backgrounds. Visitor statistics can be seen in the appendices on page 155. The interactive applications success while considering the limitations of testing was measured and analysed during user testing.

3.3.1 Low Level Prototyping

The low-level prototyping stage was conducted with 8 participants within NML. The participants included curators, laser scanning technicians, conservatives and volunteers from within NML aged between 22- 38 with differing experiences with 3D graphic applications. The icons and designs for the paper prototyping design are displayed on page 161 and the participants' anonymous data is reported on page 171. The participants were selected due to the diversity that is present online, and to get many different viewpoints to inform the design of the interactive prototype. The paper prototyping stage in this phase consisted of handing design authorship to the users and asking them to create an interface with which they would like to interact with 3D cultural content. In doing this we enabled the participants to understand the restrictions that would be applied to the 3D prototype viewer and also allowed trends to emerge at an early stage on how people understood and categorised the information presented to them. It also allowed potential problems such as: usability issues, difficulties in navigation, object interaction and other errors to be identified early. However, to limit the scope of the interface design, the participants were constrained in the use of icons and in the width and height of the space in which the interface could be implemented as can be seen in figure 2. The user was then asked to voice how they would interact with the interface and what they would expect when they pressed a button. This approach was taken as it was very

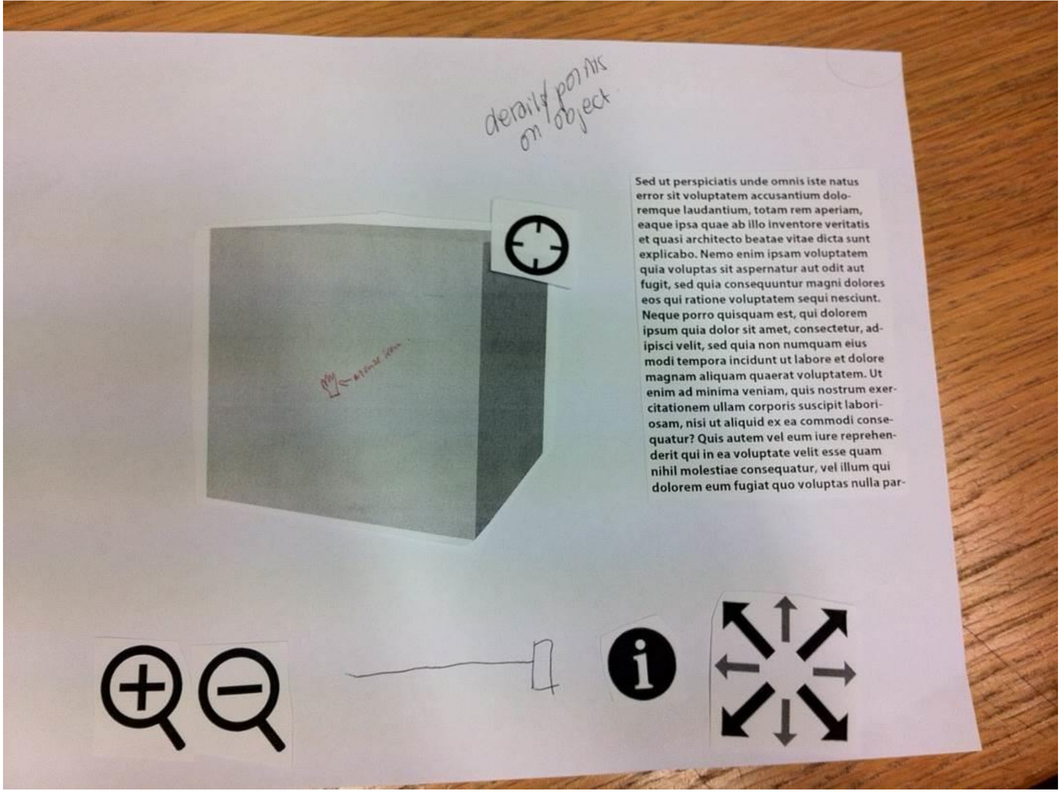

Figure 1: A representation of the paper prototype created by a participant during the paper prototyping stage, using various icons to represent how they would wish for the interface to look.

cheap to implement paper prototyping and allowed for fast iterations and designs to be generated [132]. It also generated a lot of user information, identifying usability issues and how they expect to interact with certain icons and data types in a 2D environment [133].

The created interfaces by the participants during this phase helped to reveal interesting trends and patterns. One of these was the organisation of controls, where the zoom buttons were always placed close together, and the icons for manipulating the models (arrows for directions) were also placed next to each other. This process is due to the cognitive function of "chunking" [134] classifying information into a group which you associate with that data. This was seen further when a variety of participants separated the controls, placing navigation controls outside the 3D environment, and placing icons that directly related to the 3D model within the 3D environment. The buttons placed within the 3D environment were the likes of the buttons for panning, information and one participant drew a light bulb as they thought it would be useful to have. An interesting point that split users was the presentation of information. Some participants wished for information regarding the cultural artefact appear in some way within the viewer while other users wanted it to be separate and present on the webpage; viewable but not to detract from the interactive experience.

One of the issues that identified early on was confusion over some icons, where users were unsure of their purpose. In the developed prototype, a language cue was placed alongside the icons, to avoid this issue and make the interface more accessible. When the participants were asked how they expected the buttons to work, they would describe the zoom buttons working as you would expect (+ sign to zoom in, - to zoom out), and buttons to control the object rotating and moving the 3D object. When asked how they thought the cultural model would move, for panning they could drag the 3D representative icon around the space. However, the limitations for low-level prototyping became apparent as they could not demonstrate the rotation of the model. Yet, they did identify that they would like the use of a mouse to be able to move and rotate the model within the 3D environment.

This stage enabled a suitable interface to be generated by exploring different design ideas and creating prototypes to engage with 3D artefacts rapidly. This allowed for more information to be captured directly during the testing, in contrast to asking users to write down what or how they would like to interact with 3D content. This low cost approach to user testing and prototyping also allowed participants to visualise how the interface would look, and how 2D interactive scenarios (i.e. clicking an information button) could appear without programming a 3D prototype. However, to fully explore the possible concepts for a 3D interactive viewer, a prototype was developed to allow users to interact with a 3D cultural content, and detail how information may be represented within this space.

3.3.2 Interactive prototype

The interactive prototype was developed in order to understand users' needs and their expectations when navigating within a 3D scene with 3D cultural heritage. The prototype was to be as accessible and usable as possible for a diverse audience. The data collected from the paper prototyping was used to influence the design of the prototype interactive viewer. The central design principle for the application was to create a usable system that is simple, fun and easy to learn, focusing on user centered design [131]. The Controls were denoted by various icons which showed their implied function, though a language cue was added underneath the icon to support the information that the icon was attempting to communicate [135] making it more accessible.

The information and pan controls were included in the "3D" space to associate their purpose with the object. The controls included a reset, zoom functions and orientation controls that were linked independently to the orientation and panning of the object. The GUI is shown in Figure 3 and a selection of the icons and the final design may be seen in the appendices in

B Example of icons used in the paper prototype and the prototype generated for the experiments in chapter 3 on page 161.

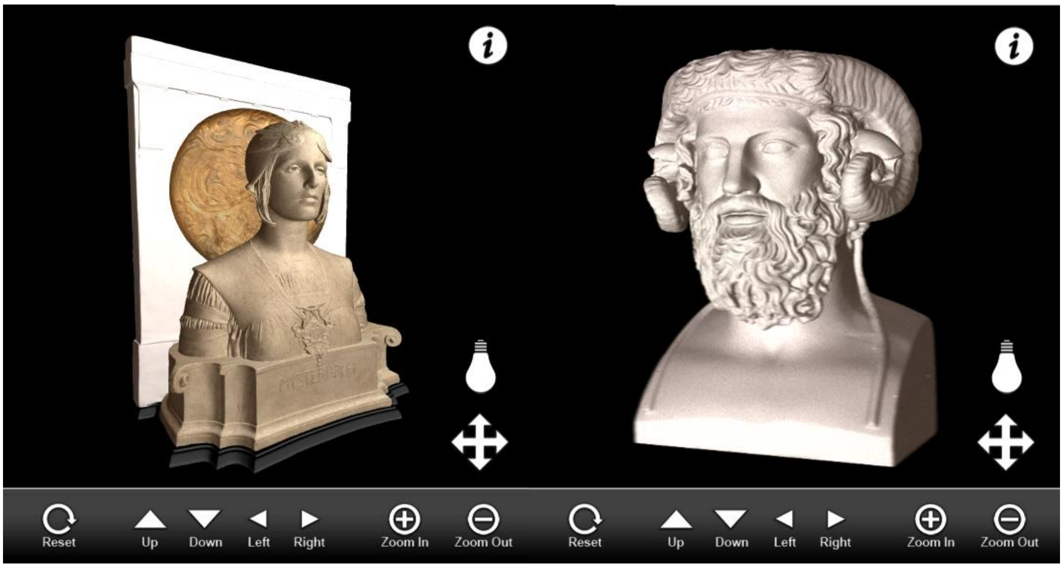

Figure 2: Zeus Ammon and Mysteriarch (National Museums Liverpool) in the interactive viewer

3.3.2.1 Implementation of the Interactive Viewer

The interactive viewer was developed with HTML 5 and WebGL with the Three.js library due to its versatility and features that are useful for engagement with visitors. Three.js was chosen due to the amount of control that it can exert over the visualisation and other parameters; which is not possible with the other libraries such as X3Dom [91] mentioned on page 47. Three.js was easy to learn and implement offering the ability to write custom, or implement pre-written shaders, and parse different file formats, such as the Open-CTM file format: which is used in the interactive viewer. Open-CTM is a file format that can be used to convert common file formats such as .OBJ files and .PLY exported from 3D software packages to a .CTM file type [136]. The advantage of the Open-CTM format is its ability to allow for a compression ratio that makes it a fraction of the size of the original file, which is extremely useful for scanned artefacts, which due to their accuracy have extremely large file sizes. It is also a simple and lightweight Application Programming Interface (API) allowing integration directly with WebGL. As the content is being created within a HTML page, it is possible to run the webpage online or offline allowing for it to be used within gallery interactive. This also allows for cultural content to be added easily and regularly to the website or interactive. The framework, however, is not supported on multiple platforms; it is only supported on webkit browsers (Google Chrome, Mozilla Firefox, and Apple's Safari) with a supported graphics card.

To enhance the participants experience with 3D digital cultural heritage [92], a scanned 3D dataset of the Zeus Ammon bust was used for the user testing. The 3D dataset was decimated to a resolution of 1 million polygons, to limit the file size that was to be uploaded while preserving the details of the 3D mesh. More information on this dataset can be found in section 5.2.2 Object Selection and preparation in chapter 5.

The design of the prototype interface was heavily influenced by the paper prototyping designs and feedback. The chosen icons for the GUI were the most popular icon for a given button featured in the paper prototype designs. The icon arrangement was heavily influenced by the paper design; the organisation of the icons and if they were placed within the GUI or within the 3D space as discussed in the low-level prototyping on page 49. To avoid the confusion over icons and their purpose in the viewer, a language cue was added underneath each icon. The icons can be seen in the appendices on page 161. The width and height of the viewer were fixed, as it needed to fit within the NML website if it was implemented. Due to this restriction, for the prototype; the information for the cultural heritage model was presented over the 3D model when the information button was pressed.

3.4 User testing of the interactive prototype

The aim of this user testing was to understand the needs of users and how they expect to interact with scanned cultural heritage datasets in new and novel interactives; making it usable for non 3D specialists and enhance visitor's experiences with digital heritage. The user testing would also address issues that could not be approached with low-level prototyping, such as: identifying usability and accessibility issues, problems with interaction, how the satisfied users were with the prototype and what tools or redesigns could be done to improve their participant's experience.

The usability and accessibility of the prototype interactive was studied with user testing with the same participants from the paper prototyping stage (2 male, 6 females), their feedback can be viewed from page 171. The participants came from a variety of backgrounds within NML, ranging from curators, conservators and volunteers. The participants had various degrees of experiences with 3D graphic applications and all had normal or corrected vision.

The data were collected via observations and asking users to complete a questionnaire at the end of the testing. The author of the thesis sat with each participant, for each part of the testing, recording their observations, difficulties and any comments made by the participant. During the testing participants were asked to voice their thoughts when using the interactive prototype. Apart from standard instructions given to each participant at the beginning of the testing, they were given no further explanation or assistance. Participants were prompted to think aloud, being asked questions such as "How do you expect that to work?" when the participant was quiet for a moment. If a participant asked for more information or help, they were reminded that the testing focused on the usability and accessibility of the interface, and needed to see if they were able to use it without further explanations.

During the testing, participants were asked to navigate to certain points on the Zeus Ammon bust within the interactive viewer. The interactive viewer did not provide instructions on how it was meant to work, which helped to identify areas that may cause confusion and help clarify areas that are clear and easy to understand. To test which was the preferred interaction style with 3D cultural content, the participant was asked to complete the above task with different control styles. The styles included the original interaction style, inverting the up axis, and the rotation of around the object changed to be around the centre of the screen, similar to Meshlab [137]. The participant also explored the use of the interaction style from the first task with 3D cultural content on a mobile device to investigate if it was transferrable to other devices.

Before testing begun, each participant was asked to read a guideline, to clarify the nature of testing and what it was meant to achieve, this can be viewed on page 165. The participants were also asked to evaluate the interface and answer a questionnaire, commenting on improvements that could be made and any additions they would like to enhance their engagement with cultural content. The participant was asked to notify the author of the thesis when they had completed the tasks and completed their questionnaire.

The study was conducted on a Dell Precision M4600 laptop with an AMD FirePro M5950 Mobility Pro graphics card, with 8GB of RAM and standard input devices (keyboard and mouse). The study was conducted within the Conservation Technology studio, a studio in one of the NML venues.

3.5 Results

The results of the users testing are reported here, describing the observations recorded of user's feedbacks of interacting with the prototype with the different interaction styles. The evaluation of the interactive on usability, accessibility and design recorded by participants via a questionnaire are presented here too. The full observations and questionnaire feedback from the participants can be found in the appendices on page 171.

3.5.1 Observations about navigation and interaction

Participants were tasked with four different scenarios to engage with a 3D digital heritage, with different interaction styles; two scenarios involved changing the Y and Z axis orientation in regards to the object; the 'world space' revolving around the centre of the viewer; and using a mobile device to interact with a cultural object. The participants were informed both verbally and through written instructions for each scenario, which can be found on pages 165 and 166 in the appendices. in the appendices. They were also informed that the aim of the testing was not to test their ability to use a computer system or engage with 3D models but to help spot any usability issues and determine their preferred interaction style. All 8 participants were able to complete the navigation tasks within the interactive viewer.

The first two tasks involved the user being asked to navigate to certain points on the object, mimicking real-world interaction with a cultural heritage object, using controls similar to ZBrush [138]. However, as the 3D dataset allows users to interact with a cultural heritage object in new ways, these tasks were designed to understand how users preferred to navigate around an object: moving the object in relation to the user or the user moving in relation to the object. These two tasks revealed a trend that all of the users preferred to navigate around the object, emulating real-world interaction (if they moved down they would be looking down upon the object) using the original controls. The inverted axis initially confused 5 out of the 8 users and they felt it was not as natural as the original navigation style, but they were able to become accustomed to the change in interaction style. As the participants felt that the original navigation was more natural and accessible, projects within NML that contained interactive 3D cultural content continued to use this navigation method. There were some issues initially regarding navigation but this was due to a lack of experience in using 3D software, which led to some of the tasks being completed in a slow time.

The third task was very similar to the first task but the navigation differed greatly. The navigation style used is commonly used in 3D software packages such as MeshLab [137], where the 3D "world space" revolves around the centre of the screen. It was found that all of the users did enjoy using this approach to navigate. However, when users used the pan option in combination with this style, all of the participants became either disorientated or mentioned this was not the way they expected to interact with the 3D cultural content. As the object had moved away from the centre of the screen, it started to revolve around the centre, which made users feel they had lost control of the navigation and that they could not reset the object to its original position.

The fourth task was performed with the original navigation style on a mobile device using the same participants. This was to investigate if the navigation style could be used in relation to different platforms such as mobile phones or tablet computers. Participants were asked once again to navigate to certain points on the object. Analysis of the results showed that all of the users were comfortable using the original interaction style on the device. However, 4 of the 8 users raised issues with regards to the gestures and finding the design slightly cramped on the mobile device. These would need to be taken into consideration if there was to be further development of mobile 3D content.

3.5.2 Evaluation of the Interactive Viewer

Through the observations, it was possible to identify usability issues and preferences when interacting with 3D cultural content within the viewer. However, the observations alone are not enough to assess the how accessible and usable the interface was. In order to have a more complete evaluation, the participants were asked to complete a questionnaire after using the interactive viewer. The questions related to the accessibility and design of the GUI and interaction with the 3D artefacts. The answers are discussed below in three sections: design and accessibility, controls and usability, and interaction with the 3D object. The participants' answers from the questionnaire can be found in the appendices on page 171.

It can be noted that the feedback generated from the paper prototyping is very similar to the paper prototyping exercise. The design of the GUI and arrangement of the buttons was complimented, yet there was still some confusion over some of the icons; over what their purpose was and the design of the icons.

3.5.2.1 Design and accessibility

The feedback from participants in the experiments, 7 out 8 users found the GUI to be clear and simple to understand. The icons with the language cues helped increase the accessibility by allowing users to understand the interface, avoiding the confusion that was s discussed on page 49. However, 3 of the 8 users reported confusion over the controls, unsure of what the function was or how it related to the viewer. One example is allowing the user to pan the object using the controls, expecting that to work exclusively with the mouse. However, all of the participants thought the functionality of the controls was clear and easy to learn. Feedback recorded during observation revealed that there was very little confusion over the navigation and users' expressions were never confused or dissatisfied.

3.5.2.2 Controls and usability

The use of the controls implemented in the GUI allowed participants to complete the tasks quickly. A positive side effect of the controls was that they involved the users more than a simple orientation of the object to points. The controls allowed the user to see more details in the model whether it was through a combination of controlling the pan and orientation. Comments from 3 of the 8 participants mentioned that the use of a raking light or other tools such as annotations would allow for a lot more information to be seen on the 3D model. All of the participants expressed a positive experience when using the prototype interactive. All of the participants again believed that the GUI was simple and clear, and clearly communicated what each of the controls did. The interface was also fun, intuitive and easy to learn, offering a natural interaction with cultural heritage objects in an unnatural environment. The switching between the control styles highlighted that the interface was an enjoyable experience to use.

3.5.3.3 Interaction with the 3D cultural object

Feedback for interaction with the 3D model was mixed, with 4 participants' content with the interactive viewer, the other 4 participants expressed small concerns they thought could improve the viewer. However, concerns with regards to the frames per second (FPS) was raised by 3 participants. A high FPS gave a smooth interaction; a lower FPS left users disorientated and confused due to the lack of response and could cause a jittery interaction. This can be highly impacting, especially for visitors who are using a computer with an old graphics card. If a visitor has a bad experience using the viewer, it could lead to fewer people using it. The Zeus Ammon bust, as an object wavefront file at 1 million polygons has a large file size of 82.5 MB, which could result in a large load time of the model. However when converted to an OpenCTM file [136] it comes in at a much smaller file size (3.48 MB) that can be reduced further if the polygon count was to be reduced. However, if the Zeus Ammon bust was reduced to video game ready quality of around 20-30 thousand polygons the artefacts would be severely degraded. This needs to be considered as this could be the first contact the visitor has with the object: original or digital. There needs to be a balance between visual quality and polygon count, which in turns affects the usability of the interactive viewer.

3.6 Discussion

This chapter has focused on the dissemination and interaction with 3D cultural content, exploring ways of disseminating and interacting with 3D artefacts platform agnostically. WebGL is one of the technologies available for the dissemination of cultural institutions 3D content, offering cross-platform access for both mobile and desktop devices. WebGL has also been adopted by many cultural institutions [89, 90, 91, 125, 119] and it continues to grow to include augmented and virtual reality as well [116, 117, 118, 119]. However, it is also important to understand how users interact with information, and how taking this into consideration can open up cultural institutions to new and diverse audiences [15]. To understand the needs of visitors, visiting a cultural heritage website or real-world venue can be difficult. However it is important to remember that the design and presentation of 3D cultural heritage needs to be accessible to all demographics as both digital and real-world venues will attract a variety of different demographics. The amount of computer knowledge and experience will also vary between each visitor meaning that tools and controls must be easy and accessible for all to use.

The use of low-level prototyping can be used to accomplish this, by allowing us to understand the needs of visitors and how they wish to interact with 3D cultural content, in new and novel ways, such as via a webpage. By understanding their needs it is possible to design and customise a 3D interactive viewer for cultural institution audiences. The results of the study carried out using low-level prototyping, helped to identify usability issues, and how the user expected the interface to work. It revealed common trends among the participants that came from very different backgrounds. Due to the diversity of participants, these trends have more meaning as they allow us to identify usability issues that may have been missed if a target user group was used. This has even more meaning when considering the diverse visitors that visit venues and websites each day.

However, the results of the paper prototyping were limited as they cannot explore the relationship between a 2D GUI and a 3D domain. To build upon and use the knowledge generated from the paper prototyping exercise, a HCI study was conducted focusing on investigating how users expected and preferred to interact with 3D cultural heritage models. This was investigated with the same participants from the paper prototyping and a 3D interactive prototype was built; influenced by the results and feedback from the paper prototype exercise. The study revealed that the preferred method was the original controls that offered natural interaction similar to that of 3D software such as ZBrush [138]. Changing the centre of rotation to the centre of the screen caused a lot of confusion when the object was moved off centre from the origin, as the object moved in correspondence with the world space. The feedback regarding the evaluation of the interactive prototype itself was positive. Participants thought the interactive viewer was easy and simple to use, with little confusion over controls. There were concerns regarding cultural artefacts polygon count, which in turn affected the fps. This can be addressed by reducing the polygon count and simplifying the 3D model such as seen in figure 4. However, there needs to be a balance between visual quality and the resolution of the 3D model, which is addressed more specifically in chapter 5 on page 82.

Figure 4: Mysteriarch at different resolutions (20 thousand, 200 thousand and 1 million polygons)

The

information generated from this chapter has been used to influence other 3D projects within NML. The knowledge has been employed within the image based viewer for the Pre-Hispanic Caribbean Sculpture project [139, 140]. The knowledge generated from this chapter has made it easier and more accessible for staff within NML, to both implement and understand the use of 3D content. It has also allowed the museum to publish and tailor their 3D cultural content for different audiences within both their gallery spaces and website.

It should be noted that the technology chosen to develop the prototype on page 45 has been continually developed and its use has exploded. WebGL is used to display 3D models for 3D printing [141], 3D portfolios for artists, disseminating user created content [95, 142] and rapidly adopted by cultural heritage institutions 90, 91, 119, 125, 139]. The development of viewers for both amateurs and professionals has also grown such as the Marmosets viewer [142], Thingiverse for 3D printing [143] or P3D to share 3D models [144]. They are all designed and tailored for their target audience and include the use of tools that will help make their interaction feel deeper and more meaningful; helping to establish a relationship with the users and bringing together a community around the display and sharing of 3D models. Sketchfab is one of these larger sites that is primarily used to allow users to explore millions of 3D objects [95] and has been used to display 3D cultural models from the NML's collections [145]. As one of the larger websites for sharing 3D objects, it claimed to have 50,000 users, 40,000 3D models and over 15 million webpage views in 2013 [146]. In three years the Sketchfab user base has grown to 500,000 users [147] with approximately 1500 uploads a day [147]. The Sketchfab community is an example of a community that has come together through sharing of 3D content and continues to grow to this day. This should also be possible for cultural heritage institutions that open up their 3D content to be shared and disseminated on the internet or via gallery interactives.

3.7 Conclusion

This chapter has presented a methodology for prototyping a 3D interactive viewer to display and engage with 3D cultural content, taking into account both user and stakeholder needs. The workflow created here involved low-level prototyping which was successfully employed to identify usability issues that may occur at the early stages of a project. It has demonstrated the importance of this approach to enable cultural institutions to quickly generate ideas, iterate designs and experiment very easily at an affordable cost.

To directly tackle usability issues in a 3D virtual environment, a 3D prototype was created exploring different interaction methods for cultural content. It demonstrated that the preferred interaction was similar to that employed in other 3D software like ZBrush [138]. It also showed that even a simple and easy to use interface offered deep engagement for the participants.

The issue of the quality of the displayed meshed, while raised as a concern in this chapter, is fully addressed in chapter 5 on page 82.

Chapter 4 Materials and visual quality assessment of 3D meshes

4.1 Introduction

This chapter introduces the issues of visualising a cultural heritage 3D artefact created from laser scanning. Recording an artefact via laser scanning will result in an extremely accurate dataset capturing the surface details but it may not record the surface material as well. The surface material may be recorded via a colour laser scanner or recreated from images taken during the scanning process. Yet the created texture from these approaches can look unrealistic, or it may contain visual or lighting artefacts. To combat this there are two approaches to accurately texture a model without needing knowledge of UV mapping: a) surface parameterization and b) 3D Texturing (Solid texturing). The former is an automatic method for un-flattening a model for texture mapping and can be used in conjunction with user input [148]. Solid texturing is a very different approach in regards to surface parameterisation, as of instead of a 2D mapping, it provides a volumetric colour information for the object without the need for a planar mapping [149]. The colour information can be generated procedurally [150] or by sampling images [151].

However, as this new 3D content is becoming a new media unto itself [152], there is a need to assess the visual quality of the displayed 3D object. This is in reference to how it is perceived by the general public, if it offers a good experience and if it is an acceptable representation of the original artefact as in accordance with the London Charter [13]. Non-contact laser scanning creates highly detailed 3D models, due to their accuracy this can result in a large file size that can impact on performance for real time interaction. A trade-off is needed between the visual quality of the 3D artefact and the processing time to offer a good experience with the 3D dataset. However, assessing the quality of a 3D dataset is extremely difficult. When a 3D dataset polygon resolution or texture resolution is reduced, artefacts may appear that may mar the visual appeal of the dataset, but this can be further complicated when you consider the different aspects of visualising a 3D dataset such that the lighting, and techniques may alter its appearance. There are however, methods to assess the quality of a 3D model, based off of 2D image quality metrics such as the Visible Difference Predictor (VDP) [153]. There are many different image metrics that have been adopted and expanded to assess 3D datasets, yet it is complicated by the fact 3D models can be viewed from different points of view, but also the lighting, shaders and textures may have an impact on the quality assessment. Though there are also methods based on subjective testing with human observers who give their opinion on the visual appeal and quality of the 3D datasets. The experiments conducted with human observers can be time-consuming and infeasible for some experiments depending on the criteria [154], but the data gained from these can be invaluable and help improve automated methods for quality assessment.

To visualise a 3D dataset in accordance with the London Charter [13], 3D datasets need to accurately convey to the user the material of the object and to also note if the visualisation is either speculation or scientifically accurate. The dissemination of these 3D datasets on either a gallery interactive or webpage can be hampered both by the lack of a means to apply a material to the 3D dataset and the high resolution of the 3D dataset. This chapter will review current literature on surface parameterisation, solid texturing and subjective quality assessment for the visualisation of 3D datasets. This is to help inform further research in these fields which are interconnected for the dissemination and sharing of 3D digital cultural artefacts.

4.2 Surface Parameterisation

Surface parameterisation is a method of applying a texture to a surface, via unwrapping and parameterising a 3D model. It has many uses in computer graphics such as detail mapping, detail synthesis, morphing and detail transfer, mesh completion, editing a mesh, remeshing, and mesh compression and many more [148]. It also offers benefits to the cultural heritage field as it allows for the parameterisation of 3D datasets and allowing for texture maps to be created for these datasets. During the scanning process, surface colour unless scanned with a colour scanner is lost. For a lot of the applications for the dataset this is adequate, however, for the visualisation it is very important to restore this lost data. The recreated surface colour and detail will create a context for the visitor, allowing them to understand the dataset more clearly.

Traditionally the approach for creating a textured a model is to have an artist create a polygonal mesh and then UV unwrap the model painstakingly by hand and add the desired texture effect. This approach is time-consuming, and the amount of effort increases exponentially as the resolution of the model increases. However, an artist will have spent a considerable amount of time crafting a digital model; they will understand where it is best to break up the model for texture mapping. Whereas 3D models created via laser scanning do not have this advantage as they are not created with texture mapping in mind. It should also be noted there are tools that can automatically un-wrap a model but they based upon surface parameterisation and can offer good results depending on the method used.

4.2.1 Brief introduction to surface parameterization

The basis for surface parameterisation is that two surfaces that share a common topology may be mapped to a common domain. This process if applied to a 3D triangulated mesh is sometimes referred to as a mesh parameterisation [155] where the 3D surface is normally mapped to a parameter domain. This is analogous to peeling the skin from an orange and then flattening the skin. This process was originally introduced to the field of computer graphics in the 1990's [156, 157, 158, 159] and since then, applications for surface parameterisation as well as research interest in the field has grown.

A good parameterisation has two main objectives; to reduce distortion during the parameterisation, and to guarantee a bijective mapping (avoiding a fold over of triangles during the parameterisation). Taking the above example of an orange, when you flatten the skin it will never lie truly flat and the same happens with texture mapping. The perfect parameterisation is known as an isometric mapping, which preserves both angle and area information. This is only possible on developable surfaces which have little to no Gaussian curvature. As most 3D models surfaces are not developable, parameterisations for this focus on the reduction of angle (conformal) or area (equiareal) deformation. However, methods that aim for a near-isometric parameterisation consider reducing angle and area distortion of equal importance.

4.2.2 Parameterisation of topological discs

The parameterisation of a mesh with a disk topology can be split into two groups. One group relies on a boundary being defined and fixed to a 2D convex domain, which is numerically simple and computationally fast [159, 160, 161, 162, 163]. Whereas, the second group computes the boundary during the parameterisation, which requires more time to process, but reduces distortion greatly, compared to methods that rely upon a fixed boundary [164, 165, 166, 167, 168, 169].

In 1963, William [170] introduced the earliest general framework for parameterisation. This framework has since been used and referenced many times in the literature. His approach applied to graphs, employing a two-stage approach that defined the boundary of the graph's vertices for the 2D convex domain, and then solved the positions of the remaining vertices through a linear equation. This approach can be directly applied to meshes, as proven by Floater [162]. If the weights in the calculation are positive, and the matrix is also symmetrical, it guarantees a bijective mapping. However, even though it provides a bijective mapping, it does not preserve any angle or area properties of the mesh [162].

Following work by Tutte [170], it was expanded upon to provide a general framework for parameterisation techniques [159, 162, 163], where research moved towards reducing the distortion of the angles and area within the mesh. However, it was found that the chosen weights can heavily influence the distortion and the bijectivity of the mapping. For angle preservation, the weights can be either addressed through the use of harmonic (Dirichlet energy), or conformal maps (Cauchy-Riemann equation). Eck et al. [159] first proposed the use of discrete harmonic maps by using a discretisation of the Dirichlet energy, proposed initially by Pinkall and Poltheir [158]. This approach creates a boundary free harmonic map, but it cannot guarantee a bijective mapping. Floater [162] also uses harmonic maps for the shape preservation method, however this is based on Tutte’s [170] barycentric coordinates. Floater uses positive weights and a symmetric matrix to guarantee a bijective mapping, albeit the smoothing over the mesh can vary. Floater (mean value) has since proposed a more simplified harmonic equation, which still guarantees a bijective mapping. However, the matrix produced as a result of the optimisation is no longer symmetrical, and the theorem discussed in [162] no longer applies.

For a greater reduction in the distortion during 3D mesh parameterisation, it is best if the boundary is calculated during the parameterisation. This is because boundaries can be non-convex, or have a border that is completely different from the fixed boundary. Lee et al. [169] increase the original 3D patches boundary by adding virtual triangles to create a more natural boundary for parameterisation. This is affixed to a convex domain and parameterisation is computed using the shape preserving method by Floater [162]. However, as the virtual boundary can move during the computation, it provides a better parameterisation. Levy et al. [164] proposed the use of conformal maps which use a discretisation of the Cauchy-Riemann equations based on least squares approximation to calculate the boundary during the parameterisation. Whereas Desbrun et al. [165] proposed the use of discrete conformal parameterisation (DCP) for the preservation of angles based on the Cauchy-Riemann theorems. Both of these approaches greatly reduce distortion but they cannot guarantee a bijective mapping [164]. Research has been extended for the least square conformal mapping (LCSM) method [164] to provide a hierarchical solver (HLSCM) to drastically increase the speed of the non-linear solver [171]. The Most Isometric Parameterisations (MIPS) method introduced by Hormann [172] optimises the non-linear function that measured the conformality of the mapping solution. It bases its parameterisation on the Shape Preserving Method (SPM) [162] and to guarantee a bijective mapping SPM moves vertices within the kernel of their neighbouring vertices. However, due to the process of constantly checking the vertex neighbourhood, it is very slow and the examples in Hormann's paper are limited to simple 3D surfaces. The Degener implementation [173] also measures the conformality of the mapping, however, it uses a hierarchal solver to speed up the non-linear solver.

The Angle Based Flattening (ABF) [154, 166, 174] approach differs from the above parameterisations methods by instead of using a UV coordinate metric, ABF bases its metric on the angles directly. ABF calculates the parameterisation of the mesh within an angle domain. The parameterisation looks for angles that closely match the angles within the 3D mesh and if they satisfy the constraints then it provides UV coordinates in a parameter domain. The produced parameterisations are bijective and have reduced stretch and area distortion, it can though suffer from global overlaps. A post-process can be applied to the parameterisation, reducing distortion further and removing global overlaps [166]. However, the ABF method is extremely computational heavy, making it impractical for meshes above 30k faces and prone to errors when calculating the UV coordinates. Sheffer and De Sturler [166] addressed these problems by implementing a direct and hierarchical solver depending on the mesh size and included a stable angle to UV coordinate converter.

The previous methods conserve angle, yet to preserve the area of the triangle during parameterisation, the use of stretch and distance preserving parameterisations are the most popular approaches. Though to guarantee a bijective mapping, a developable surface is required. Early attempts to implement this into a parameterisation [156, 157, 160] were met with difficulty due to the complex numerical calculations required to minimize the distortion metric. Sander et al. [161] introduced a solution to construct a progressive mesh, where all of the meshes in the sequence shared a common parameterisation. Sander et al. introduced a metric to minimize the stretch error during parameterization, to balance sampling rates across the mesh and also introduced a metric to measure texture deviation. After splitting the mesh into appropriate charts to parameterise, it calculates a stretch metric using SPM [162] to preserve the shape of the patch. It then resizes the mesh depending on the stretch and deviation metric to preserve area deformation. During the refinement process, vertices that share coordinates on multiple patch boundaries are evaluated to determine if they will cause a triangle flip over. This approach guarantees a bijective mapping, however to keep balanced sampling across the mesh, some areas may be enlarged to compensate for area shrinkage in other areas. Sander et al. [167] expand on their earlier work [161] focusing on a single mesh instead of a progressive mesh. This allowed them to optimise chart boundaries for arbitrary genus meshes, creating less distortion in the parameterisation.

Distance preservation was introduced by Zigelman et al. [175] and attempts to preserve the distance between vertices on the mesh, both locally and globally. It calculates a geodesic map of the mesh using a fast marching method, and then it parameterises the mesh using Multi-Dimensional Scaling (MDS). This approach, on a surface with very little Gaussian curvature, can provide a bijective mapping. However, complex surfaces suffer from fold-overs and overlaps, due to the nature of the preservation of area over angles. Zhou [176] uses a similar approach to MDS parameterisation for a mesh but also includes an optimisation process to avoid fold-overs. Zhou's method moves the vertices to the centre of neighbouring vertices staying in the 1 ring neighbourhood as implemented in [167].

To achieve a near perfect parameterisation, reduction of both angle and area deformation is equally important. Parameterisations which reduce both [165, 173] provide very good parameterisations for simple models, but still struggle with complex models. However, there is a common strategy being used in recent research, which first implements a parameterisation technique and then applies a post-process, to reduce the distortion of the parameterisation. Sheffer et al. [148, 166] implement a post-processing procedure after parameterisation of the ABF and ABF++ (an extension of the ABF) implementation [169], by allowing the parameterisation to relax and grow within a rectangular domain filling an empty room. It also applies an isotropic scale to the mesh to minimize overall error. Kraevoy et al. [177] and Lee at al. [178] parameterisation methods consider both angle and constraint points, during parameterisation. Lee et al, then applies the post-processing procedure from Kraevoy's work [177] to smooth the mesh parameterisation. During the parameterisation Steiner vertices are added to the mesh as well as a virtual boundary to achieve a more natural parameterisation. The post processing procedure removes the virtual boundary and Steiner vertices while trying to maintain an identical connectivity to the 3D mesh. However, if the Steiner vertices cannot be removed without causing distortion, these vertices will be added to the mesh, complicating its topology.

The primary focus for surface parameterisation has been the parameterisation to a disc, due to the classic approach for parameterising arbitrary genus models. There are many other techniques that could be used instead, such as the multi-charting approach [160, 179, 180], or approaches utilising cone singularities [181, 182, 183] or by mapping to a 3D parameter domain [184, 185, 186, 187]. However, it should be noted parameterisation to a disc works extremely well with constraint points.

4.2.3 Constraint Points

Constraint points allow a user to choose points on an image to match corresponding features on a 3D model. There are several approaches in the use of constraint points, and these can be split into two groups; soft and hard constraints. Soft constraints are an approximation of how the texture image should appear on the model, whereas hard constraints try to perfectly match texture coordinates to features on a mesh i.e. such as eye placements on a facial model.

Soft constraint approaches proposed by Levy and Mallet [160] and Desbrun et al. [165] investigated the use of a least squared system and LaGrange multipliers respectively, however, they cannot guarantee a bijective mapping. Zhang [180] proposed a solution by separating the model into separate regions based on basic geometry shapes, converting them to regions, and unfolding them with little stretch based upon the Green-LaGrange deformation tensor. However, the occurrence of fold-overs due to constraint points was not well documented within the paper. Hard constraints are a more widely debated subject, as it needs to accommodate user input for the exact feature matching between an image/photograph and 3D mesh. Eckstein et al. [188] first coined the term hard constraints for texture parameterisation and his solution attempted to tackle this problem. Eckstein uses a constrained simplification on the mesh, to guarantee a legal parameterisation for all multi-resolution scales for the mesh, and introduced the use of Steiner vertices to guarantee no fold-overs. This approach is theoretically sound and can handle large constraints, yet it is extremely complicated to implement and is not highly robust [177]. Kraevoy et al. [177] and Lee at al. [178] both implemented similar solutions, which split the surfaces into separate charts and adds a virtual boundary to these charts. A constrained Delaunay triangulation is applied to these charts, calculating the parameterisation before removing the virtual boundary. Due to the input of constraints from the user, the texture can be extremely distorted and requires a smoothing post-process to be applied. However, Kraevoy et al. [177] method cannot guarantee a bijective mapping as it does not consider consistent neighbouring ordering to find matching triangulations. The methods detailed here allowed for one image to be used, whereas Zhou [176] implemented a solution that allowed for more than one image to be used to create a texture image for a mesh based on hard constraints. It also includes an in-painting technique to texture areas that did not have an image assigned. To prevent seam artefacts appearing in the final texture map, Zhou subdivides faces that lie outside the seam region and create a boundary for the seam. Colours are then interpolated on the vertices that lie outside the boundary, using a Radial Basis Function (RBF) to fill in the seams.

Parameterisations based on RBF interpolation between constraint points have been proposed [189, 190, 191]. Such interpolations create a smooth parameterisation and reduce distortion greatly, without the need for implementing a post-process procedure to smooth the final result. Tang et al. [190] and Lee and Huang [191] implemented solutions based on solving a linear system calculating the analytic parameterisation.

However, both these approaches cannot guarantee a fold-over free condition. The subsequent parameterisations are heavily dependent on basis of the RBF function (i.e. thin spline plate), which can heavily influence the distortion [189]. Yu et al. [189] solve this by implementing an iterative step framework by deforming the mesh to constraint points. Yet when the algorithm detects a fold-over is approaching, it evaluates which triangles would be suitable to be subdivided. This method can guarantee a bijective mapping, however, the subdivision of triangles can add many redundant vertices and the computational times are influenced heavily by the amount of constraint points.

These last sections has discussed in length and focus on surface parameterisation with and without the use of constraint points. The methods that have been discussed here have managed to inspire and push the boundaries for UV mapping. State of the art techniques within surface parameterisation focuses on the use of Steiner vertices in attempt to guarantee a bijective mapping. The use of calculating a virtual boundary during the parameterisation for a more natural parameterisation has also led to work that wishes to preserve the shape and area of the triangles within the mesh to 2D parameter domain. Yet these methods still have faults, where it is still possible in some cases to not be able to acquire a bijective mapping; it can lead to redundant vertices being added to the mesh or increase distortion in the mapping depended on the method used. These can be rectified to a certain degree with the use of constraint points or a post process to reduce errors. Yet in most cases for the parameterisation method to even begin it needs a model with seams and atlas maps first to be generated; it requires someone with knowledge of UV mapping to implement this. As this literature is focusing for cultural institution environments, these skilled personnel are far and few between, more automatic approaches could be considered such as solid texturing.

4.3 Solid Texturing

Solid texturing is inherently different from surface parameterisation as it does not require UV coordinates to texture an exterior or interior surface. Solid texturing can represent the 3D models surface colour as volumetric 3D colour information, which can give the appearance as though the model is carved from the material. This approach allows for colour information to be ubiquitous and can be applied to any other model without the need for a planar mapping. This colour information can also be used to create a new a domain within the model for a new texture to be generated for the interior of the model. Solid texturing is a relatively new field of research, only recently been explored in the last 20 years. However, the advantages it can offer over 2D texturing are worth pursuing which include high quality subsurface scattering, synthesising textures for the surface and interior on the fly in physics simulations and the lack of needing a parameterisation. However, there are issues with regards to solid texturing which has hampered its use in industry.

The main issue itself is memory usage in comparison to a 2D texture. While a 2D texture uses a small amount of memory, synthesising a solid texture requires information to be stored on a regular 3D voxel grid, which is expensive to both synthesise and store even for a small resolution texture. There is also an issue with creating a texture that is free of visual and tiling artefacts on the 3D model itself. While certain methods can be used to hide repetitions and artefacts they may still be seen if the model is sliced along a certain plain. Furthermore, as the approach is mainly an automatic approach, certain features or artistic intentions of the user may be lost during the synthesising process.

Colour information can be calculated by either using procedural methods [192], or by adopting 2D texture synthesis for solid texturing to represent complicated materials. Procedural methods were explored early on in the research for solid texturing [192] where algorithms generated the colour for the volumetric dataset. Procedural methods are easy to implement with little memory storage needed in the computation of the colour information, and they can also be used in real time applications. Yet, as the colour is generated following the rules of the algorithm, the creation of textures is very limited and changing parameters can be very time consuming and inefficient. This limits the range of materials that can be represented procedurally to marble, wood and cellular patterns; whereas 2D texture synthesis used in conjunction with solid texturing can represent a range of textures and materials.

Yet, solid texturing can be split into two different categories: boundary independent or dependent. Where boundary independent methods calculate the colour for the volumetric dataset without considering the model's shape and boundary dependent textures are reliant on the boundary of the object to create the texture information. This texture volume can then be used to create a new domain for an independent internal material to be created. However, for the purpose of this literature review, it will focus only on boundary independent approaches as it is incredibly unlikely for cultural heritage institution to wish to use a 3D digital cultural object for a physics simulation or fracture simulations.

4.3.1 Statistic feature matching

Pioneering methods for solid texturing focused on using statistics that they were able to obtain from analyzing an input image, in order to create a synthesized texture based on these inherent properties. The most common methods used in 2D texture synthesis would include the use of histograms, spectral analysis of an input image and estimating 3D properties from measurements taken from 2D images. Early work within the field first concentrated on the use of Histogram matching [149] and spectral analysis [193, 194] which is still used in state of the art methods for solid texturing.

Heeger and Bergen [149], use a pyramid histogram of the input image to extract global statistics from the image at different resolutions and sizes. A solid block is initialized with noise and is iteratively processed by attempting to match the volumes noise and inputs images histogram via a matching operation from the captured pyramid histogram [149] to synthesise the solid texture. This approach focuses on creating a 2D texture for UV mapping a model, but it does allow for the expansion into solid texturing. It would use the exact same process as described in the paper but would require the use of multiple 3D steerable pyramids for the analysis of the initial image. It would then require a 3D steerable pyramid created from the 2d pyramids for the output histogram for the texture. It should be noted that the histogram matching chosen for the texture synthesis, is based on the Cumulative Distribution Function (CDF) and inverses this for its output histogram. This samples the image for colour distribution by mapping from a chosen "bin size" to the interval of [0,1] for each colour channel and the inverse CDF is remapped from [0,1] to "bin size". The histogram matching then calculates the colour on the output solid texture by replacing it with a colour from the input image which matches the CDF value. As each colour channel is calculated separately in an attempt to avoid artefacts, the final texture merges the different colour spaces together to create the final texture. This approach works well with stochastic textures but as the colour channels are calculated separately, there is a possibility of creating visual artefacts when the final solid texture is reassembled [149, 195].

Following a different approach, Ghazanfarpour and Dischler [193] proposed using spectral analysis of an input image to capture parameters and to create a solid texture from this information. The spectral analysis data from the 2D image is captured by using a Fast Fourier transform (FTT), and the captured parameters can be used in conjunction with a noise function such as Perlins noise function [192] to generate a procedural texture. This work was expanded the following year by Ghazanfarpour and Dischler to incorporate multiple images [194]. The multiple images are used to represent arbitrary slices through the texture block along an axis, helping to control the appearance of the final solid texture. The algorithm works on the basis that the input images that are used for slices along the axis values will not change or be affected by displacement along the axes, where a non-orthogonal view is then applied slowly to a solid texture to resemble the appearance of the input images along their axes. This is accomplished by iterating between a spectral and phase function on the captured FFT data, which is applied to an initial block of noise, which iteratively converges to resemble the input images along their various axes. However, there are issues with the phasing part of the algorithm which could cause visual artefacts and at times could not converge [196]. This method was later again expanded upon by Dischler et al. [196] where instead of a phasing stage, they used spectral and histogram matching to improve results and speed up the convergence for the solid texture. However, though the results produced via the method of spectral analysis are acceptable, they are only applicable to stochastic materials but were not able to preserve patterns or fine details in the solid texture [195].

Alternatively, Jagnow et al, [150] proposed a method that would synthesise a class of materials classed as "particles within a homogenous material". The approach is based on stereology, which extracts 3D information from 2D measurements on an image. The approach provided precise high quality textures for solid texturing by calculating the distribution of "particles" from an input image and replicating this to create a volume of various particle shapes (sphere, cube, ellipsoid) and distribution. After the size of particle distribution was calculated, the particles were spatially placed using simulated annealing by editing non-uniform rational B-Spline surfaces [150]. The process then applies a noise function for fine details that are similar in appearance to the input image. Cross sections taken at any point within the colour volume would replicate the distribution of particles as the input image. The work was expanded upon in 2008 by Jagnow et al. [197] to include a method for automatically approximating particle shapes and placing them within the volume. This work also researched how human's perception of a solid textured could be influenced by changing the particles shape [197]. This work produced realistic results yet, omitted finer details from the input image. This approach however, is only usable for a certain class of materials where it is possible to synthesise particles from the input image.

Work by Qin and Yang [198] developed a 2D texturing synthesis method called Basic Grey Level Aura Matrices (BGLAMs) which is based upon aura concepts [199]. The BGLAMs is an approach which samples all possible co-occurrences grey levels within the neighbourhood of an input image. It was expanded to solid texturing [200], and generates a solid texture by sampling the aura matrices of the input image(s). To generate the texture for the solid texture, the aura matrices are iteratively sampled to provide constraints for a white noise volume to be modified while the distance is decreased between the aura matrices. This is based on the assumption that two textures will resemble each other if their aura matrices distances are within a certain threshold of each other [198]. This method can produce pleasing results if a structured texture is used but it only works with grey scale images. To use colour information from images it would require a similar process to Heeger and Bergen's [149] applied to each colour channel separately and then rebuilt but this would lead to the same visual artefacts as those seen in [149].

While these approaches are able to produce acceptable results and helped to pioneer research into solid texturing, this was only possible with input images that had well defined parameters such as histograms or regular patterns. They are not suitable for the use of natural material textures that are irregular and contain macro and mesostructures within them. A more appropriate approach would be the use of methods that are commonly used in 2D pixel texture synthesis.

4.3.2 2D Texture Synthesis Methods

One of the most common approaches in 2D texture synthesis is the use of Nearest Neighbourhood matching. This approach calculates the colour for the pixel in a synthesized texture by determining the best colour match by analysing its nearest neighbour. The synthesized texture is then created pixel by pixel where the colour is determined by comparing against the input images neighbourhoods to determine if there is a similar neighbourhood that could be used to substitute the colour in the pixel. This approach was used for multi resolutions textures in combination with Gaussian Pyramids [201] and this approach was extended into the 3D domain. The approach is similar to 2D neighbourhood matching but instead of matching the colour for each individual pixel, it attempts to synthesize colour for a voxel instead, which presents issues with the extra dimension. The initial issue is the difficulty in comparing a voxels 3D neighbourhood against an images 2D neighbourhood. The second issue regards the difficulty in comparing different images orientated along the different viewing axis to calculate the colour for one voxel.

Wei attempted to solve these problems, by first extending fast texture synthesis [201] with one input source to multiple image sources [202]. This method primarily for 2D texture synthesis allowed users to provide multiple images and a weighted image to determine neighbourhoods and how the colour of the image would blend together to form the final colour. This was extended into 3D, by orientating the images along the various axes for solid texture synthesis. Colour candidates for each voxel are calculated by taking input pixels from each source image along their respective axis and solving an energy minimising function, which is the squared difference between in the image and the neighbourhoods. The final colour is chosen by using the weighted image to determine how the colours would mix together [202]. The algorithm works in a multi-resolution fashion by synthesising a block of noise, and iteratively processing each voxel and changing the colour information. This approach showed the difficulties in implementing neighbourhood matching in solid texturing; however, the resulting solid texture exhibited signs of blurring and is not able to preserve even simple patterns of the input image [151].

Johannes Kopf et al. [151] proposed an improved algorithm based upon histogram matching [149] and texture optimization [203]. The basis for this algorithm was to efficiently synthesise the input image by directly solving a global energy minimizing function. There was also a desire to implement an implementation that may use multi-layer textures (BTF) and recreate surface attributes from these and input images. The process for solid texture synthesis is based upon iteratively trying to solve the global energy minimizing function, bringing the volume closer to resembling the input image by switching between two phases: an optimization and search phase. A volume is first initialized by sampling random colours from the input images, where the optimization attempts to match the voxels as closely as possible to a neighbourhood candidate based upon the nearest neighbouring voxels by iteratively using reweighed least squares [203]. The search phase then searches for best matching neighbourhoods from the input image for each voxel within the volume dataset [151], and this is the most time-consuming operation within the function.

During the process of minimization, there is a chance that the formula may be caught in a local minimum causing the same local neighbourhood to be repeated throughout the synthesized texture [151]. Kopf et al. [151] introduced a reweighing scheme based upon histogram matching of both the synthesized texture and input image. The histogram of the synthesized texture would be iteratively changed so its global histogram matched that of the input image by reducing the weights for texels that may influence an increase in the differences between the histograms. As is seen in previous methods, the texture is synthesized by selecting arbitrary colours from the input images and converging from a coarse to fine level of detail and is carried out in a multi-resolution fashion using trilinear interpolation [151]. It is also possible in this implementation to include a feature map to retain features from the input image or create defined patterns [151].

While this method has produced better results than other implementations and is able to be used with multi-layer images (BTF, RTI), it suffers from long computation times. The time taken to synthesize a 128^3 solid texture with 3 channels per texel can take between 10 to 90 minutes

[151]. The time taken to synthesize a solid texture from a BTF dataset which can feature up to 9 channels is not mentioned. Since the method is a global implementation, the entire volume needs to be synthesized requiring large memory storage. This work was expanded upon further by Chen and Wang [204] by optimizing the synthesis problem by using a k-coherent search for the search phase. It also implemented a new histogram matching method based on position and index histograms to improve results and allow it to use highly detailed textures.

However, the time for synthesizing and the memory storage for a solid texture from a set of input images is a large problem and very limiting for the applications for solid texturing. Previous implementations [198, 202, 151], focuses on solving a global problem, synthesizing the entire colour volume as it is possible for an individual voxel colour to be dependent on all voxels within the colour volume. This approach results in large memory usage in comparison to a 2D texture map and requires complex computations to implement. Dong et al. [205] implemented a method that works in parallel with the GPU to rapidly synthesize solid textures interactively and voxels to be synthesized are a small subset of voxels close to the surface of the model instead of the entire colour volume.

Dong et al. [205] expanded upon Wei's neighbourhood matching algorithm [202] so it runs in parallel and the computation speed is increased by using a 2D K-coherence algorithm [206] extended to the 3D domain. The colour of a voxel is determined by candidates chosen from input images which have a similar neighbourhood to the synthesized texture. A possible candidate is comprised of three interwoven neighbourhoods from the 2D images. This, however, results in a large number of possible candidates by combining the triples of 2D neighbourhoods from the input images, which Dong et al. [205] attempt to limit by pre-processing possible candidates. Candidates are selected if they meet two important factors: colour consistency and coherence. Colour consistency for possible candidates is measured by first searching for consistent colour across the overlapped regions then compared against other triples neighbourhoods to calculate the smallest consistency error candidate; which is the sum of the squared colour differences in the overlap regions. After selecting possible candidates based on colour coherency, candidates are then checked if they are able to form coherent patches in all directions with other neighbouring candidates.

After this pre-processing step, each pixel for the input images has a set of 3D candidates, allowing for efficient solid synthesis and is performed in a multi-resolution fashion using a 3D volume pyramid. The best fitting candidates for each pixel are tiled across a block for the initial synthesis which is then iterated through two steps: up sampling and correction to increase the resolution of the block; a jitter step is applied after initial synthesis to introduce variance to the solid texture. Up sampling is used when the resolution is increased, and it is calculated by a simple texture coordinate inheritance. The correction step is performed on all synthesized voxels in parallel, where the input data is read from the previous step to stop neighbouring voxels from influencing each other. A best matching 3D neighbourhood is sought by using the voxels current 3D neighbourhood to define possible new candidates through the use of a k-coherence algorithm [206]. A candidate is defined as being a mixture of possible candidates sets referred to by their texture coordinates. The best matching is then sought by considering the distance between the current neighbourhood of the voxel and the new candidate sets using a summed squared difference of the colour values. The new coordinates are then substituted in for the next round of corrections.

This implementation helped to improve both the efficiency and synthesizing a solid texture on demand and producing high resolution solid textures. This method is still considered state of the art for raster images as it has been expanded to be used for translucent materials [207] and extended to allow users to provide a tensor field to guide the solid texture synthesis [208]. Whereas Zhang et al. [208] implementation differs slightly from other implementations as it has to use non-homogenous spatial distributions for synthesizing the solid texture as repeatedly using a synthesized cube across the volume with the tensor field would produce visual artefacts. Yet the approach is still limited as it takes a local approach where the search is limited to a small spatial window, this in turn limits the colour that may be used within the final solid texture.

These previous approaches have been able to provide excellent results in the short space of time that research has focused on solid texturing. The results have allowed for a higher quality of solid textures to be generated with the use of stochastic or patterned textured input images. Yet, time to generate and store these solid textures is still presenting issues. It may be time to consider using another format for the input image, not one based on raster images but vector images instead.

4.3.3 Vector solid texturing

A vector solid is a relatively new branch of research within solid texturing, where recently it was proposed to use vector images for solid texturing [209]. Instead of storing information per pixel, it uses primitive shapes and paths with various attributes for either colours or gradients to represent images. They offer advantages over the use of raster images such as being resolution independent, having a small memory footprint in comparison to raster images and they are easy to edit. Their first implementation for solid texturing was proposed by Wang et al. [209], where their algorithm created a solid texture using an RBF colour fitting function to create a vector texture representation. The approach, however, uses rastered images as the input images and requires the use of Kopf et al. [151] solid texture synthesis method to first create a synthesized solid texture. To generate colours for the vector solid, an RBF colour fitting function is then applied to a colour region within the synthesized solid texture to approximate the colour of that region. Each region is randomly selected from the volume but can be set to be more uniform for sampling. A minimization function is then applied to each region to reduce the error in that region and uses a teleportation scheme from [210] to avoid a local minima problem that could occur, and result in the texture to be repeated across the solid texture.

To store the new vector solid, only the signed distance field of the vector is stored, which is an N x N x N array of 32 floating point numbers. However, Wang et al. [209] noticed that regions far from the region boundaries do not affect the overall shape and thus can be set to a memory size of 4 bits while not degrading the overall quality of the boundaries. To reduce memory usage further they used Region Labelling, which stops RBF's from affecting other regions and that if another RBF does not affect another RBF they can share the same label. They use a Welsh-Powell heuristic algorithm [211]) to re-label these regions and can offer acceptable results for regions with up to 18 colours, while reducing memory usage to ~20% of a rastered solid texture [209]. To avoid anti-aliasing, the solid vector can also be generated in a multi-resolution fashion similar to mipmapping to avoid artefacts appearing when zooming out.

This method offers resolution independent texturing and low memory consumption for a solid texture. It does, however, have limitations such as needing a feature mask to define sharp features and currently is limited to one feature mask which can cause artefacts to appear. Also, high-frequency details from the original input image may be lost when converting the rastered solid texture to a vector solid using an RBF colour fitting function [210].

Zhang et al. [212] proposed another implementation with vector solids where they synthesized the vector solid directly from the input images, avoiding the heavy computations and storage needed for a voxel synthesis. They propose a gradient mesh approach using tricubic interpolation to capture the colour for the synthesized texture and feature maps to preserve details. The use of a gradient mesh approach to capture colour information offers advantages over RBF colour fitting, as it samples more regularly and reduces the colour errors that are approximated from the input image. It first initializes a block that assigns a feature vector to represent a colour from the input images; it then follows a two-step iterative approach to improve the results of the synthesis in a multi-resolution fashion from a course to fine level of detail. The two steps involve finding choosing similar neighbourhoods from the input image to match the vector patches within the vector solid and updating the patch to resemble the input image. This approach is based on Kopf et al. [151] neighbourhood matching and optimization approach, yet this differs slightly. Instead of trying to match neighbourhoods from the input image to a voxel, it attempts to find regions which closely resemble control points on the gradient solid. To search for similar patches, it uses the approach of Dong et al. [205] to look for consistent colour slices and uses a PCA projection to reduce the feature vectors dimensions to speed up computation time [213]. After each region has been found, it gives the control point an N x N patch of sample colours to update the representation on the gradient vector. However, due to few control points on the vector solid, there may be voxels that are unaffected throughout the vector synthesis while other voxels may contain multiple feature vectors which can result in blurred voxels. Zhang et al. [212] combated this by using a “bucket” for each voxel, and using a two pass quantization approach to avoid blurring of voxels while reusing buckets for areas that may contain no samples.

This approach offers improved results over Wang et al. [209] in both computation time and memory storage. It also offers an advantage by directly synthesizing a solid vector directly from the input images, and can be adapted to synthesise a synthesised raster solid texture directly. However, it is still limited in the detail it can display and is still reliant on feature maps to preserve sharp features. However, the loss of high frequency features on textures is a common problem for all vector approaches as this detail would be lost during the vectorisation process [212].

4.4 Quality Assessment for 3D models

With the possibility of using either surface parameterisation or solid texturing to apply a material to a 3D dataset, there is still the issue of the quality of the visualisation that may be output for interactions or visualisations. As one of the main goals of computer graphics is the creation of CGI that is indistinguishable from its real-world counterpart [214]. There are many new physically based rendering algorithms or artistic styles that can be chosen that can help suspend our disbelief, by rendering a 3D model and texture as realistically as possible. A scanned cultural heritage artefact can be sub-millimetre accurate capturing both the shape and fine details of the original artefact [4], resulting in a 3D dataset with millions of vertices and triangles. However, users that interact or view these 3D models are focused on their appearance and not interested in the number of polygons or texture resolution. It is the perceived quality of the final render that they consider more important [154] and as technology continues to grow this expectation is always increasing.

With this expectation for extremely high quality 3D output, there is also a need for interactivity without significant performance issues. As these scanned models are extremely accurate and require large amounts of memory, a trade-off is needed between the interactivity and the perceived quality of the 3D dataset. Even with the advances in computing power, network speeds, 3D datasets can still place a great strain on CPUs and GPU's for calculations as well as require large storage space. This can be reduced by compressing both the 3D datasets polygon count and texture resolution but it can introduce artefacts and mar the visual appeal in the final interaction. This section will deal with various methods that can be used to judge the quality of compressing and reducing the resolutions of methods. It should be noted, that there is still no defined method for judging the quality of 3D meshes and textures [215, 216]. It is even possible for some studies to contradict their selves such as [217, 218], which compared inverse tone mapping techniques for enhancing Low Dynamic Range images (LDR).

Due to this difficulty, the next section will review the literature on both 2D image metrics that have been applied to 3D datasets, as well as subjective testing with human observers, to provide ground knowledge and the best possible approach for assessing the quality of 3D datasets, and achieve the best perceptual experience for users. For a full review of current literature for the quality assessment of 3D meshes, Corsini et *al*. can provide a more in-depth analysis [219].

4.4.1 Quality Metrics

There are many quality metrics that can be used to judge the perceived quality of a 3D dataset, as there are many areas where a visual artefact can appear, not just in the final render but within the 3D dataset too. A 3D dataset can accrue visual artefacts from a range of operations on the mesh, from watermarking, compression, transmission and other operations. To detect these artefacts there are two approaches that can be used, one is a model based and the other is image based. A model approach focuses on using the model and directly, while an image based approach relies on images being taken at different angles to judge a 3D datasets quality. Image based can be split into two different approaches, one where the algorithms are dependent on certain views, while the other is independent. These metrics aim at predicting the visual quality or identify artefacts within sampled 3D model/image. The perceived quality itself, however, can be judged by human observers, in subjective quality experiments. These experiments involve human observers giving their opinion on the quality of the 3D meshes or identify artefacts that may be inherently difficult to identify in an automatic metric. The next section will cover subjective quality assessments. More information can be found within [219].

4.4.2 3D Quality Subjective Assessment

2D image metrics for the evaluation of 3D objects are able to predict fidelity and quality, they are not able to prove that their results though would be the preferred choice for human observers [220]. 2D image metrics are an automatic approach that compares images against a reference image, yet this does reveal much about the perceived quality. To properly assess this, a subjective quality assessment can be used, using human participants to give their opinions on the artefacts or their perceived scores for the 3D model.

Perceived tests, can also give a mean score to test a metric and for comparing and ranking different metrics between one another. The computed score of the predicted quality from the metric, and that of human observers when compared in correlation through the use of Kendall's rank-order correlation or Speamans rank order, can provide a quantitative evaluation method to assess the effectiveness of the used metric [221].

Although there are many recommendations for how to conduct the subjective evaluation test, there are many standards that define how a subjective test should be carried out for similar mediums. One of these standards is the ITU-R BT.513 for the evaluation of image quality on television screens [222]. There are also different experimental methods that may be used, such as forced binary choice, double stimulus, absolute ranking, and pairwise comparison. However, a study conducted by Mantiuk et al. [223] that focused on evaluating the effectiveness of 4 separate methods, found that pairwise comparisons were the most effective [223]. Pairwise comparison contained the least variation between scores during the testing, and it was also the quickest one to complete even with a large number of comparisons. This could be due to the fact, that it contained simple instructions, and asked the participant to choose which image they thought was better [223].

There have been many studies using subjective quality assessment experiments for both static and moving 3D objects, both stationary and animated models [220, 223, 224, 225, 226, 227, 228, 229, 230, 231, 232]. The first subjective tests used to assess the quality of 3D objects was conducted by Watson et al.[220], and Rogowitz and Rushmeier [225], which has gone on to inform many studies. They tested different algorithms for the simplification of 3D models at different levels. They both used a rating system that asked the participant to rate the object using a double stimulus versus the original [220, 225]. These are both very important studies as they were used in conjecture with 2D image metrics to predict the chosen quality at each level of simplification of each 3D model, and if they were suited for assessing 3D models [220, 225]. Rogowitz and Rushmeier [225], conducted two experiments, one asking users to rate still images of decimated 3D objects and then to rate a sequence of images showing a 3D model rating. This study alone showed how important that lighting can play in the perceived quality, and that perceived quality can be changed depending on whether an object was stationary or animated [225]. Watson et al. [220] on the other hand measured the performance of 2D image and model based metrics as well as conducting a subjective test rating the perceived quality of the 3D objects. The study showed that image Metrics were a good predictor for fidelity versus model based metrics [220]. However, Watson et al. [220] used 2D images of a static scene, the results could are only a good indicator of a static scene. Rogowitz and Rushmeier [225] extended this to animated and different viewpoints and found images metrics were not as good an indicator of fidelity as in a static scene. The study also revealed that evaluation results can be affected by the 3D objects; animal models and manmade artefacts [220]. This was also proved by Pan et al. in their study, which found abstract 3D models can help to avoid semantic interpretations [226].

Two more studies that focused on the use of subjective experiments to assess perceived quality for simplified models, are by Rushmeier et al. [224] and Pan et al. [226], however, they used textured models. These studies focused on the how texture and polygonal resolution may affect our perceived visual quality of the model and how effective texture can mask artefacts. The Rushmeier et al. [224] study discovered that the substitution of polygon resolution and texture resolution are object dependent and that low resolution textures can actually harm perceived quality of a 3D object regardless of polygonal resolution[224], where improving the texture resolution improves perceived quality. Also, the masking effect of texture on the contours and boundaries of a 3D object can affect perceived quality [224]. While Rushmeier et al. focused on the use of spheres for their testing, Pan et al.[226] used 3d objects and textures that were captured using a 3D scanner and proposed a subjective quality metric that would contribute to perceived quality of both the texture and polygon resolution. The captured data from the laser scanner was constructed to provide a ground truth 3D object. Simplifications were applied to both the captured 3D object and texture independently, to provide 3D models for the subjective test. During the testing, participants were asked to rate the simplified objects, in comparison to the ground truth object and the “worst” object which was the most simplified object. The study provided an insight into the relationship between polygonal and texture resolution; after a point polygonal resolution no longer affects perceived quality, yet texture resolution are perceived linearly [225]. However, the metric they proposed is not applicable for general testing for predicted quality, due to a number of parameters that are too specific for each 3D object.

During the simplification process of 3D models from their originals, artefacts and noise can be introduced to the new mesh. This can then be measured against the original mesh to evaluate the overall quality of the new mesh. The introduced artefacts or noise affect the surface and can either introduce roughness or smoothing of the surface. This introduced roughness can affect perceived quality of a 3D model, especially as it is very hard to predict the effect it may have on the overall quality of the mesh if it is in an area with a lot of surface details. To assess this for 3D models, Lavoué [228], proposed a local roughness metric to assess quality, and used this on 3D models that had separable regions of both rough and smooth areas. He carried out the experiment by applying different levels of noise to both the rough and smooth areas, in order to measure the masking effect of artefacts. The study showed that smooth areas were far more likely to reveal artefacts than rough areas. The provided metric for the study, provides a local roughness estimation of how likely an artefact may be hidden in certain areas without notice, and would be useful for designing future watermarking schemes or decimation algorithms [228]. As a result of this study, a database was created containing 26 models of varying polygon resolutions generated from 4 reference objects and publically available for use [232].

The above studies, asked participants in their subjective tests to evaluate perceived quality of 3D models presented to them, while a reference image or 3D model was presented at the same time. However, some studies used a method based on displaying a single 3D model or image, with or without a hidden reference displayed among the simplified models. To evaluate the visual fidelity of 3D models created from a watermarking algorithm, Corsini et al. [234] proposed two studies. They would focus on the various artefacts that may appear due to different algorithms used to watermark 3D models. Using the above testing method, they acquired a mean opinion score (MOS), to assess the perceived quality of various algorithms used to watermark each 3D model. They also proposed a perceptual metric which combines the subjective MOS with a global roughness value calculated per 3D object which is then derived into simple roughness difference based on the variance of the geometric Laplacian [228]. The provided metric was able to provide good results, predicting human perceptions of distortions on watermarked 3D models. Lavoué [235] also proposed a similar study to Corsini et al. [234], measuring the perceived quality of watermarked 3D models [234]. The database that was created in [232] was used in a subjective experiment using a single stimulus to acquire MOS for each 3D model which had either non-uniform noise or smoothing across the surface. The participants MOS for this study [228], were used to evaluate the performance of the mesh structural distortion measure (MSDM) metric, which has proved to be very similar to human judgement, especially in complex scenes[228]. This metric has since been improved upon in [235].

While either a single stimulus or double stimulus can provide a direct and applicable MOS for each 3D model per study, the ranking of the 3D models can be inaccurate. This is due to the difficulties of assigning a score to each 3D model, especially when quality differs among participants for experiments. This can be somewhat negated with the use of the simpler paired comparison method, which asks participants in studies, to choose their preferred choice between a pair of 2D images or 3D models. Silva et al. [230] proposed to use this method alongside, a rating subjective experiment to assess a simplification database containing 30 objects created from 4 reference objects. The reference models were simplified using three separate algorithms, and 65 participants took part in the subject assessment. For the rating experiment, they were asked to rate the models from 1 (very bad) to 5 (very good). The second experiment then asked to choose preferred stimuli in a forced binary comparison experiment [230], helping to actually rank the models. The rankings obtained from this experiment gave a valuable insight to perceived quality, as the ratings may not reflect the actual perceived quality of the 3D meshes. The results and database created via this study are also publicly available [236]. The experiment has also been repeated with a larger selection of models for just the paired comparison results by the same authors [232].

To assess artefacts that can be introduced by compression artefacts, [231] proposed a study using paired comparison against a database of 3D objects. The database contained 68 objects, created from 5 reference objects and different compression algorithms. Using a forced binary choice, the 69 participants, were asked to choose a preferred distorted 3D object between 2 stimuli and a reference object presented to them. The derived score for each object is computed from the users' choices [231]. The computed scores, user choices and the compressions database are also publically available for testing [236].

4.4.3 Merging image and 3D metrics

Apart from the work by Rushmeier et al. [224] and Pan et al. [226], subjective quality assessment rarely touches upon the use textures or other elements. This is important as a scene in which a 3D model is displayed for real-time interaction or applications can contain many elements; including the 3D models and texture; as well as lighting and shader models for the creation of the final render or tone mapped image. The metrics and studies in this chapter focus on either the 3D model or are tuned to detect specific artefacts. While surface details are already considered for the subjective evaluation, the lighting information of the final image should also be considered as well in the final evaluation of a 3D model, through a combination of 2D image metrics such as the HDR-VDP2 metric [237] and a subjective quality assessment metric.

4.4.4 Quality Testing Design

During the above studies, parameters became apparent that have a major influence on the perceived quality of a 3D object. These parameters could be simple things that were not considered such as slowly rotating an object, or changing the position of lighting. These simple parameters may not normally be considered, but can introduce significant bias to the scores or ranking of 3D objects. This is extremely important, especially as these parameters can all be controlled for a computer simulation. The parameters that most affect a study are listed below.

4.4.4.1 Lighting

The type of lighting and direction play a big factor in the perceived quality of an object, depending on where it is viewed from. Rogowitz et al. [225] found that lighting impacted on scores, where lighting from above was scored differently compared to the same objected lit from the front. Lighting from the front, creates a masking effect, where the preferred lighting comes from above and from the left [239].

4.4.4.2 Background

A object that is placed on a background of a uniform colour or simply being black can have an impact on user studies and perceived quality as it can mask the silhouette of objects . A simple black colour is used commonly in studies [220, 224]. A vertical gradient from blue to white has been used in a study by Corsini et al. [234], in an attempt to avoid the masking issue of a uniform colour.

4.4.4.3 Texture and Shading

To achieve photo/hyper-realistic effects in CGI and video games, complex textures, complex shaders and subsequent maps (normal, ambient occlusion) are applied to 3D models. Yet most subjective studies that wish to quantify perceived quality rarely consider texture maps or complex materials, they use either simple materials or none at all. This is to avoid artefacts being made more obvious due to the complex material or acting as a mask to artefacts [231]. Studies that have used texture include [224] and [226]. Due to the fact that texture can act as both an enhancer and masker of artefacts, they are not commonly used in studies as it allows more control of the output for perceptual studies.

4.4.4.4 Types of Objects

The types of objects that are chosen for subjective perceptual experiments themselves can lead to bias and need to be considered carefully. A study by Watson et al. [220] discovered that the results that were obtained for animal and man-made objects achieved different results. Objects that are abstract and not familiar to participants also may help to avoid semantic issues [224], as objects that are familiar may introduce bias [226]. A final observation is the shape and complexity of the object, artefacts that may arise to decimation and simplification may not be noticed due to the complexity or masked.

4.4.4.5 Interaction and Animation

While traditional 2D image metrics for a 3D scene or object, evaluate the quality from a fixed viewpoint, 3D objects can be manipulated into different positions and viewpoints quite easily. Studies by Pan et al. [226] and Rogowitz et al. [225], animated the 3D models within their test for participants to rotate, while [234, 236] all allowed for the free manipulation of the 3D object. During the testing by Rogowitz et al. [225], it was discovered that animation can have an impact on the perceived quality of an object. The rotating mesh obscured artefacts caused from the decimation process, so lower quality perceived meshes, were viewed as being of higher quality when moving compared to be at a standstill [225].

4.4.4.6 Display

The type of display and resolution can also be a factor in the perceived quality. This is important depending on how the 3D datasets are visualised such as via mobile devices, desktops or immersive environments with the use of the Oculus Rift, Vive and 3D screens. While they may offer a more immersive experience, they sacrifice quality of rendering versus an immersive experience. To use these devices in testing, the displayed area should be large enough to incorporate the full extent of the object, while not showing many objects within one scene. This may lead to confusion and to bias being introduced into the results.

4.4.4.7 Number of Objects and stimuli order

For tests that compare and contrast the reduction of polygons for 3D objects and/or texture resolution, the levels of reductions need to be selected accordingly. Too few levels for comparisons and the results may not provide good enough test data, while too many maybe unfeasible to expect someone to undertake. There are some studies, which used from 3 levels [220] up to 6 levels [232], but the selection depends on the type of test and how it is conducted. It also stands to reason, that how they are presented for comparison tests may affect the perceived quality, side by side or in succession. Rogowitz et al. [225] displayed to their participants the objects in succession, but allowed them to refer back to the reference object allowing for more in depth comparisons.

All of these parameters play a major part in the design and the studies for subjective quality assessment. There are some standards as well that may be worthwhile exploring such as ITU-R.BT.500 [222], which defines standards to use for subjective tests that involve 2D images. Yet, the standards for the evaluation of 3D models will evolve alongside the devices that may be used for testing offering a more immersive experience with 3D objects.

4.5 Discussion

This chapter has discussed the challenges of applying a texture to 3D cultural heritage objects, but also how to measure the overall perceived quality of the 3D object. This is especially important at this current time where 3D models are becoming a new form of media unto themselves, and their penetration into both cultural heritage institutions and society are becoming more prevalent. The importance of how these 3D models are displayed, interacted with and perceived will grow and expectations will be high from the public. It must also abide by the London Charter [13] as well. This leads to an interesting issue, of what would be more beneficial for use within cultural heritage, and how would you assess the quality of the 3D cultural object? The use of surface parameterisation is used extensively within the computer games and film industry, but is more involved and difficult. Solid texturing, offers an automatic approach that can be applied across models, but is extremely difficult to compute and implement. There are also many methods and approaches to assessing the perceived quality of 3D models, with each having their own pros and cons. The next sections will discuss in common areas that lie between both surface parameterisation and solid texturing: reusability, controllability and usability, appearance and distortion and which would be the best fit within a cultural heritage organisation. The issues of how you evaluate the perceived quality of 3D models and how you would design experiments would be discussed as well.

4.5.1 Reusability

A surface parameterization, and the achieved mapping from the 3D domain to the 2D is dependent on that object and the polygon count, it cannot be shared between different meshes. It is possible by following an approach by Sander et al. [161, 167] to construct a progressive mesh, which allows for UV coordinates to be shared if the mesh was to be simplified or decimated, which is important when you wish to display a lower resolution model, while keeping the same UV coordinates. This then allows for a texture to be created from a 2D image, or painting directly onto the model, and to be easily swapped out for other textures that can be created in applications such as Substance Painter[240] or created via PolyPaint in Zbrush [138] to name a few.

Solid texturing produced by methods explored in this literature review can be used to provide colour information for any potential dataset. As long as the 3D model is within the solid textures domain, colour information can be acquired for the model independent of its shape or complexity. The problem lies with the texture synthesis which is attempting to create a larger (non-repeating) texture image from a smaller input image. Depending on the approach that is used, it may result in repeating patterns or artefacts to appear to be displayed across the 3D model. There is also the memory issue, as a 2D image with RGB channels would require N*N*3 versus a solid textures N*N*N*3. For a 1k texture, a 2D texture would require roughly 341 kb, a solid texture would require closer to 3GB without some form of compression. Though this could be negated following an approach by Dong et al. [205], to synthesise it when required, but it still uses more memory than a 2D texture.

4.5.2 User Interaction and Controllability

For the implementation of an approach to create a texture or surface mapping within a cultural heritage institution, it should be easy to implement and easily change the texture depending on requests from curators or other professionals with the institution. Solid texturing offers an automatic approach, without the need of prior knowledge of UV mapping and only require a 2D image to be supplied to create synthesized solid texture. Although the Zhang et al. [208] approach allows the user to create a tensor field to guide the creation of the texture, it is still limited. Solid textures are normally tiled across a surface with some rotations and jitter applied to give variance to the texture, but tiling can still become apparent. Editing a solid texture can be expensive to compute and store but a vector solid approach [212] can give the user more control over the final appearance and in an interactive manner.

In comparison surface parameterization requires can be either an automatic approach or require a certain level of interaction. This interaction would include the selection of areas to localise the creation of seams, or generating atlases for the use in texture mapping [164]. There are approaches that allow for the use of constraint points on images. Where the mapping is controlled by the constraint points being either hard (accurate placement) or soft (an approximation) for where the points should line up on the 3D model in the final mapping. However, with a mapping generated via a UV map, it is then possible to quickly generate a texture from a 2D image or via texture generating or poly painting software [138, 240, 241].

4.5.3 Distortion and Mapping

While the final appearance of the 3D model and texture need to be evaluated, the generated surface parameterisation mapping or solid texture also need to be evaluated for visual artefacts and mapping distortions. The final appearance of the model that is textured using solid texturing contains no distortion, since the colour information is represented as volumetric grid of values. When a model is placed within this domain, the surface at these points would inherit this colour information through trilinear interpolation. Earlier methods did have visual artefacts with regards to colour as they attempted process each colour channel separately and then blend the channels back together [149, 198]. More recent approaches no longer have this issue and have been capable of displaying high resolution textures [197] and offering the option for resolution independent textures [208, 212]. While the time to generate and synthesis a solid texture and the storage required is still an issue, results are becoming more impressive and have scope for research opportunities.

The appearance of a model textured via a parameterization will always contain distortion unless the parameterization contains little to no Gaussian curvature, which is very infrequently seen in cultural heritage artefacts. Where an isometric parameterisation is not possible, parameterization methods attempt to reduce the angle or area distortion or a combination of both in the parameterization. However, in some cases it may not be possible to create a bijective mapping or avoid foldovers or overlaps of texture coordinates. While some methods propose using post processing method to smooth out the final mapping [165, 177], state of the art methods propose the use of Steiner vertices to guarantee a bijective mapping and use of atlas charts for the parameterization. Yet this can introduce seams artefacts to models visualization which Zhou [180] However, what used to be an issue for an artist is now easily fixed by the use of software such as Substance Painter [240], which can cover seam artefacts by letting the user paint directly on the seams or generate a texture via 2D texture synthesis.

4.5.4 Best choice?

Each approach for both solid texture synthesis and surface parameterisation both presents pros and cons, for the use within cultural heritage it would be only worth using one approach. However, which would be more beneficial and useful for the implementation and dissemination of 3D models? The approach of solid texturing would definitely appeal for some organisations, as it is an automatic approach and various materials could be created and applied between various objects. However, solid texturing does not seem to be supported outside the realm of research at the moment. There are many areas that could be improved, such as the search and optimisation phase [201, 205] memory needed for consumption and storage but it just is not supported within industry. The use of software for surface parameterisation has exploded in recent years and has become easier to use for non-skilled individuals. Zbrush UV Master [241] allows users to paint areas for where seams could be hidden and areas that they wish to preserve. With the press of a button, the model would be unwrapped and offer a close to bijective mapping in a timely manner [241]. The achieved mapping can then be transferred to other software so that a texture may be created for the model. The mapping may also be used with 2D texture synthesis with Substance Painter [240] to create a texture that is physically based on a material, has no artefacts or repetitions that can appear in solid textures and provide other maps to allow for normal, specular or ambient occlusion mapping.

4.5.5 Quality Assessment

Yet with the creation of a mapping or texture, there is a need to assess the perceived quality of the mapping and 3D model together. This can be achieved via an automatic approach such as 2D image metrics for 3D models to predict perceived quality and fidelity or human subjective quality testing. Subjective testing is the best approach and why 2D image metrics were not addressed in this chapter. Though the results provided via subjective testing is invaluable, some testing can be infeasible to conduct, but it can also offer data that contradicts other studies [217, 218]. However, it offers an opportunity for the general public to assess the perceived quality of 3D cultural heritage models.

For the use of assessing objects for their perceived quality, it should be noted that objects that are man made artefacts are perceived as having better quality than those of animals [220]. Also, objects that are familiar to participants may bias results due to semantic interpretations of the object [226]. Texture and materials are also a major factor in the perceived quality of a 3D object; this is due to the fact that they may hide artefacts and other details using the masking effect, [224, 206]. Animation and movement may also bias results [225], resulting in lower polygon models being perceived as having a better quality than when stationary. There are more parameters to be considered but it shows that designing an experiment and quality metrics to assess the 3D models is complex. There is a need for an approach that considers both the 3D model itself but also the environment in which it will be rendered/displayed. Furthermore, a standard for the testing of 3D models needs to be created, to lessen the chance of bias or errors being recorded in the results; a standard for images alone is not enough. There are many approaches for conducting an experiment and designing it but there is no standard approach. The use of paired comparison for images to obtain a ranking appears to be a route, though, that many studies are following [223]. This is due to the results recorded seemed to be very similar across participants in studies but it is also the fastest to be completed [223]. This is an extremely important point, especially if a cultural institution wishes to conduct such a study with visitors. The time needed to evaluate and assess the perceived quality would need to be completed quickly but also provide accurate and concise data.

4.6 Conclusion

With the growth of 3D contents as a media that is actively used both by the general public and cultural heritage institutions, the importance of how these 3D models are displayed is becoming increasingly more important, with regards to both a surface material and polygonal resolution too. While newer technology can acquire a 2D texture of the scanned model, there are still many scanned 3D models that do not have a texture. This chapter attempted to address this by reviewing state of the art literature on both surface parameterisation and solid texturing.

The use of surface parameterisation or solid texturing has vastly improved since their introduction to computer graphics in the 1990s. Surface parameterisation has been widely adopted within the CG and games industry as the traditional method of texturing a 3D model, yet it still requires persons to have knowledge of UV mapping and texturing to implement these methods. However, the technology is becoming more readily available for use as well as easier to use, offering many approaches for rendering and texturing 3D models. Though, as of yet no perfect solution yet that could be implemented within a cultural heritage setting.

However, when presenting a 3D cultural artefact to the public, the visualisation is one of the most important aspects to offer a good perceptual experience [154], there is a need to assess the perceptual visual quality of the 3D mesh. This chapter has also investigated various methods into subjective assessments of 3D models. This chapter has highlighted that there are many efficient metrics for assessing the quality of a 3D model, as well as many parameters that need to be considered into the design of a subjective experiment. There are also no set standards for the evaluation of 3D models, yet there are some such as the ITU-R BT.500-13 [222], for acceptable results. However, for everyday occurrences, it is very rare for a 3D model to be rendered with only a simple colour or texture. It is far more likely that it would be rendered within a scene with many varying affecting parameters, thus an evaluation model needs to evolve for 3D model assessment that reflects real-world situation. It needs to incorporate the assessment of the 3D model itself, its material, the lighting and the environment it is being displayed in.

Chapter 5 Subjective and objective assessment of 3D textured and non-textured Cultural Heritage Artefacts

5.1 Introduction

Cultural institutions are moving towards a paradigm of openly sharing and disseminating their knowledge, collections and their 3D digital artefacts (see chapter 3, page 44). Whilst there are still many restrictions in place that can hamper this (see section 2.2 Background, and chapter 3), there is still a need to investigate the best possible way of presenting these datasets for human consumption. One of the main questions being, how to offer the best perceptual experience for visitors to a cultural institution or museum, without the issues of sharing or visualising a full resolution dataset on a website or a gallery space?

The datasets that are created from laser scanning consist of points in 3D space that create a triangulated surface, and texture coordinates that allow 2D textures to be placed on the 3D dataset itself. These datasets themselves can be subjected to a number of processes to reduce the polygon counts, watermarking, compression and other operations [226]. These processes help to reduce the strain that can be placed on both GPU and CPU's for the display and dissemination of 3D datasets [226], with 3D models with polygonal resolutions commonly in the millions or more. Simplification and compression are one of the ideal solutions for the dissemination and display of 3D data sets, while maintaining the integrity of the 3D data sets and offering a solution to piracy. Yet, these processes may inadvertently cause degradation to the overall appearance of the 3D model, and this is true as well for the 2D texture maps which could be caused by a reduction in their resolution or the use of compression such as JPG. This may also impact on the interaction and engagement users may have with the 3D datasets, therefore there is a need to evaluate the visual appearance of the rendered simplified dataset. Especially when attempting to offer the best perceptual experience.

There are many metrics that evaluate the visual appeal of 2D images produced via computer graphics, such as the visible difference predictor (VDP) [234]. They focus predominantly on global illumination or tone mapping [220, 225, 226, 237], and how they affect the overall visual appeal of the image. They do not take into account the 3D model itself. However, the literature that does focus on 3D models, are primarily concerned with the surface of the 3D model and artefacts that may occur during various processes to the mesh. Little work has been done concerning the use of a combination of 3D model, textures and lighting for the final produced image. There is work that has been undertaken by Pan et al. [226] that focused on artefacts that were caused by the simplification of the 3D model and texture subsampling [226] with a subjective experiment. Work has also been undertaken by Guo et al. [221], that was similar to Pan et al. [226] via animated videos, but covered a wider range of geometry and texture processes [242]. However, both Pan et al. [226] and Guo et al. [221] restricted their work to limit and control the parameters of their experiments, and not evaluating real-world expectations. As 3D models and their textures are becoming more mainstream, the 3D models are viewed under different viewing conditions that include the 3D model, textures and different illumination conditions.

This chapter presents a large scale (70 participants) subjective experiment that tests for real-world expectations of the perceptual experience of 3D digital artefacts, using a pair wise experiment and a subjective questionnaire. This study will investigate the cost to texture and polygonal resolution that will allow for the best possible perceptual experience while reducing the 3D model and texture resolution. It will be conducted under real-world conditions, using rendering settings and techniques such as normal mapping, ambient occlusion, blur and image-based lighting. This also gives the results more weight, as it allows visitors to assess the visual quality of the 3D model, taking into account more than just the 3D dataset presented to them as discussed in section 4.4.3 Merging image and 3D metrics.

This chapter will, for the first time, also compare non-textured stimuli versus textured stimuli, measuring the effectiveness of texture when compared to surface details captured with non-contact laser scanning. It has been seen (see section 4.4.2 3D Quality Subjective Assessment), there are few studies that focus on the use of textured models when assessing visual quality. These are also limited to those, which use various texture resolutions while avoiding low resolution textures as they have been shown to degrade the perceived quality of a mesh regardless of the polygon resolution [224, 226, 242, 243]. Other studies [220, 223, 225, 228, 229, 230, 231, 232, 244, 245], use no texture and very simple shaders to display the 3D model, but no experimental work has compared the use of non-textured versus textured models. This could be due to researchers wishing to control parameters or believing that a non-textured model will be perceived badly [221]. It is important to measure the perceived quality of a non-textured 3D cultural dataset versus a textured dataset, as cultural heritage datasets will contain high level of details due to the accuracy of the dataset; whereas textures can mask and obscure details from the user [221, 228]. The comparison of non-textured and textured models in this thesis will reveal if the former are perceived worse than the latter; reporting how well the high level detail of the 3D dataset compares to textured models.

The parameters for this stage of the research are partially governed by the requirements of the engineering doctorate partner, the NML. The research needs to address the needs of the museum to provide an good an experience as possible, similar to that of the original artefact, and to be as physically accurate as possible in accordance with the London Charter [13]. The dataset also needs to be very manageable, so it may be transferred and used easily on multiple devices both on the galleries and with the museums website. NML has a collection of approximately 400 3D datasets that have been created with the use of non-contact triangulation based laser scanning to within 0.1mm accuracy of the original artefact [4]. The 3D datasets contained within the collection are varied and have been created from various materials. However, the technology that was used by Conservation Technologies to produce these did not capture surface material information (see 1.4 Problem Statement). This poses a unique challenge not just for NML, but also to other cultural heritage institutions that wish to disseminate their 3D cultural content that may not contain a material.

5.2 Methodology

5.2.1 Experiment design

The experiment to evaluate user’s perceptions towards textured and non-textured models at different resolutions was conducted within a gallery within the World Museums venue at the NML. This venue would allow for museum visitors to interact with an artefact and a digital representation of the cultural heritage artefact. A survey would then be completed with the visitor’s consent regarding their experience with the digital counterpart and the real artefact.

This would be conducted using an OpenGL interactive viewer that could be used both on touch screens and other interactive peripherals [119, 221, 242].

5.2.2 Object Selection and preparation

Conversation Technologies had created a loosely defined archive containing nearly 400 3D objects being stored within NML's archives. There is little to group them, except for what organisation they were scanned for and the definition of the object. The collection itself contains sculptures, busts, hogback stones, reliefs, archaeological finds, a tumour, a World War 2 bomb and more.

The models were selected for the experiments for a number of reasons; due to the complexity of their forms in terms of visual detail, how they are interacted with and the material they are made from. The statistics for each object and the reasons each object was chosen can be found in tables 1, 2, 3, 4. The experiments wished to investigate objects that offered a variety in surface detail, from very smooth surfaces to artefacts that contain a large amount of visual detail within the surface. Models were also chosen that offer different types of visual details, which could be familiar or iconic to a visitor yet still off rich, complex features. How users interacted with the artefacts was also a major factor in the decision process resulting in the objects offering either two dimensional or three dimensional interactions. A 2D interaction is used to describe the interaction for objects that predominantly lie flat on a surface (reliefs, rock art, paintings) and 3D are for those that require a visitor to move around the objects or handle them to rotate the object. Many objects that have been scanned have been made from a variety of materials, with metal and stone being the most common material; therefore the objects were chosen from these categories. Each object offers a unique experience during the testing allowing for a variety of interactions with the original/replica alone.

Images	Selected Objects	
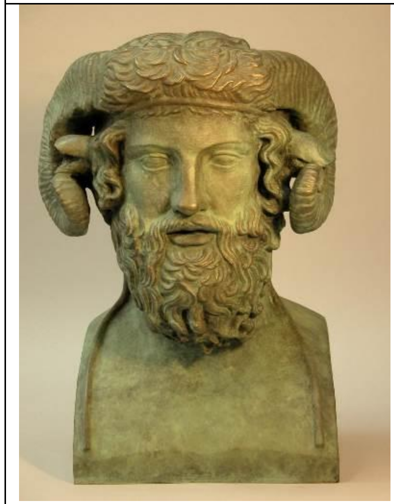	Name:	Herm of Zeus Ammon
	Accession Number:	NML Ince 163
	Approximate size (mm) H x W x D:	H470 x W380 x D270
	Accuracy (mm):	Sub-millimetre
	Material:	Bronze
	Scanner/Equipment:	3D Scanners Ltd. Modelmaker X (Faro gold arm) with a MMX 35 sensor
	Mesh Resolution (Vertices / Triangles)	2,523,883 / 5,040,129
	Contextual Information	The Herm of Zeus Ammon is an artefact within the Ince Blundell Collection within NML, made from marble. It is currently on display within the new Egyptian Gallery within the World Museum. Conservation Technology, made a replica of the marble bust using non-contact scanning laser scanning, to create a replica to be displayed at the entrance of the World Museum. Using additional funds, a master cast was created using a 3D printing process. A replica was then cast from the master pattern creating a bronze cast. The bust was finished by hand and displayed within the World museum to be handled and touched by visitors. As the original bust was not available and the marble replica could not be moved. The bronze replica was used for the study.

	Reason Chosen:	This object was chosen due to its size, form, material and volume. It is the second largest of the selected objects for the user testing with the recognisable form of Zeus Ammon. It also is made of bronze, offering variable colour information depending on the angle that the lighting hits the objects. Also due to its size, it offers the user the ability to interact with it in a "3D" manner, where they must manoeuvre the entire object to be able to see all the details on the object.

Table 1: Table 1: Zeus Ammon statistics, and reasons for being chosen

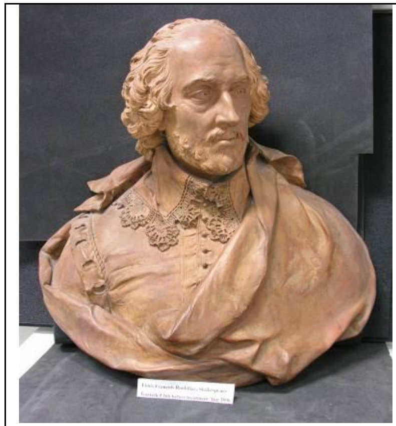

Name:	Roubilliac Shakespeare
Accession Number:	None
Approximate size (mm) H x W x D:	H620 x W590 x D290
Accuracy (mm):	1mm
Material:	Terracotta
Scanner/Equipment:	3D Scanners Ltd. Modelmaker X (Faro gold arm) with a MMX 35 sensor
Mesh Resolution (Vertices / Triangles)	3,399,624 / 6,799,264
Contextual Information	This iconic terracotta bust of William Shakespeare, created by Louis- François Roubiliac is currently on display within the Garrick club in London and not part of the NML collections. There are two other busts being displayed within the British Museum and the Shakespeare Memorial Library in Stratford upon Avon [246]. The bust used in the study is a replica of the Garrick Bust, created in Terracotta from the dataset created by conservation Technologies with non-contact laser scanning.
Reason Chosen:	This object was chosen due to its size and its recognisable form of Shakespeare. It is the largest of the selected objects, allowing for a 3D interaction where the user needs to rotate and pan the object to see the full details on the object. It also contains details on the back of the object, giving the user an

		insight into how the bust was made. The bust was also made from terracotta, offering a contrast to the material offered by the Zeus Ammon bust.

Table 2: Roubilliac Shakespeare statistics, and reasons for being chosen

Name:	Anglo-Saxon gilded cross headed brooch
Accession Number:	2006.65
Approximate size (mm) H x W x D:	H160 x W90 x D15
Accuracy (mm):	0.5mm
Material:	Bronze/corroded ronze and gold
Scanner/equipment:	3D Scanners Ltd. ModelMakerX (Faro gold arm) with X70 sensor
Mesh Resolution (Vertices / Triangles)	637,936 / 127,5876
Contextual Information	The artefact presented here was made by Conservation Technologies, who made a replica of an Anglo-Saxon great square-headed brooch (approximately 160 mm x 90 mm) for visitors to examine and handle at the Weston Discovery Centre, World Museum Liverpool. The surface of the gilded bronze brooch was too fragile to mould and so a non-contact approach was used. The object was laser scanned and a master pattern produced from the resulting computer model using a 3D printing process. A replica was then cast from the master pattern in a copper alloy. The replica brooch was gilded and finished by hand and a new clasp fitted to its back.
Reason chosen:	The Anglo Saxon Brooch was used in the user testing, due to its size, fine surface detail and variable surface material and its 2D interaction. The size of this brooch allowed for the artefact to be handled

		quite easily and for participants to see nearly to scale details on the 3D model presented on the monitor. It was also thin, so the interaction with the artefact would predominantly be 2D, allowing users to see it from just above or below to discover the full details of the object. The artefact also had a lot of variety in its material, from the shiny gold on the top side of the brooch, to rusted corrosion and copper patina forming on the original artefact, which was replicated on the 3D model. The Anglo Saxon Brooch also had a replica created from bronze, in the Weston discovery centre, to allow participants to handle the replica and compare that against the original and digital artefact.

Table 3: Anglo-Saxon gilded cross headed brooch statistics, and reasons for being chosen

Name:	Temple Wall Relief Carving of Tuthmose I
Accession Number:	56.22.141
Approximate size (mm) H x W x D:	H360 x W350 x D150
Accuracy:	0.2mm
Material:	Limestone
Scanner/equipment	Minolta VI900
Mesh Resolution (Vertices / Triangles)	249,966 / 498,383
Contextual Information	The artefact used in the study was a replica created from limestone and finished by hand by Conservation Technology to allow visitors to touch and interact with the Egyptian relief. The original Limestone carved relief artefact depicting King Tuthmose I, was being conserved within the conservation department within NML and is currently being displayed within the new Egyptian Gallery in the World Museum. The artefact was laser scanned, and the resulting dataset was used to create a CNC milled replica which was put on display within the Weston Discovery Centre within the World Museum
Reason chosen:	The Egyptian relief was chosen for its bold details and the finer details that show that the original artefact was supposed to have the arm crafted straight down. The relief is also flat, offering the user

		to be able to observe the artefact from above to fully appreciate the visual appeal. Though with a raking light the user would be able to see further details on the model that are not visible under normal lighting conditions. The Egyptian Relief was also made from limestone, with some residual paint and dirt that are present on the surface due to the bombing during the 2nd World War. This was fully replicated on the 3D digital replica. While the Egyptian relief was created through a mixture of photogrammetry, and scanning. However, the file and its texture were created solely for the use with VRML [247], and there was not a possible way to convert it to an OBJ with the texture file. However, an STL file was generated to create a replica of the object and this file was used for the experiment.

Table 4: Temple Wall Relief Carving of Tuthmose I statistics, and reasons for being chosen

5.3 Stimuli Preparation and Texturing

Although these 3D digital representations were created from real objects within the collections of NML, their surface material was not captured alongside the surface geometry. Each model also has a very high number of vertices per object, ranging from 249,966 to 3,399,624 vertices. This number is clearly larger than some studies that have taken place to evaluate the perceptual visual quality of 3D textured 3D models [226, 242]. This large number also makes it extremely difficult to create either a manual or automatic surface parameterisation of the mesh, especially for those discussed in chapter 4 on 60.

As these original 3D datasets are overly large, an approach similar to that of Pan et al. [226], was taken that would allow for each models polygon count to be reduced, making them more manageable. Each model was simplified using the Quadratic Edge Collapse Decimation [242] used within Meshlab [137], as it allows the preservation of boundaries, normal's and texture coordinates in the decimation process and is readily available in Meshlab [137]. Each full resolution selected model would be decimated by 10 % using the Quadratic Edge Collapse [242], preserving the boundaries and normals of the mesh. Three independent reviewers would then be asked to compare between each of the decimated objects and this continued until there was a unanimous decision between where all agreed upon noticing a difference between the decimated meshes from the original mesh. When a difference was noticed, a level up from that decimation level would be appointed as the high resolution stimuli for the study. Table 5 reports the polygon resolutions of the stimuli, and table 6 reports the differences in percentage between the original dataset and the new 100% stimuli. As can be seen in table 6, the reduction in resolution is over 50% for each model, and in the case of the Ammon, Shakespeare and Brooch, they were reduced by up to two thirds.

With these new resolutions for the 3D models to be used within the experiments, they were then processed with Zbrush UV Master [241] and further pre-processed using UnFold3D [248] to achieve as close as possible to a bijective mapping as possible. To achieve a realistic result and adhere to the London Charter [13], a physically based approach for the texture creation was taken. Using photographs taken during the scanning process, the material was recreated within Substance Painter [237], to be as physically close as possible to the artefact using various 2D synthesized materials with Substance Painter [237]. A normal, diffuse, specular and ambient occlusion map was generated with Substance Painter [237] at a resolution of 2048 x 2048 pixels. The results of the texture maps and final geometry resolutions can be seen in figure 5. It should be noted as the textures were created using Substance Painter [237], the textures will not be identical to the artefact but are a close substitute.

5.3.1 Pair wise Stimuli Generation

For the pair wise experiment, the high resolution object was decimated using the Quadratic Edge Collapse [242] preserving the boundary, normal's and texture coordinates to a further 70, 40 and 10% of the high resolution, creating 4 polygon resolutions. As these models could possibly be displayed on websites using WebGL [97], the models were also further compressed with the OpenCTM file format [136]. This allows for the 3D models to be compressed, using a lossless compression using a variety of differentiation operations, followed by lossless entropy encoding using LZMA encodings [136]. Due to the differing sizes of polygon resolutions, instead of a manual approach to adjust and provide different desired visual qualities ("Excellent", "ok", "poor") for each model. The percentage for the decimations was chosen to keep the percentages uniform across the 3D objects and not to introduce bias into the results by deciding what is an “excellent” or “poor” representation of a 3D dataset.

The texture of the objects was only subjected to a loss of resolution down from 2048x2048 to 1024x1024, 512x512 being saved as PNGs to avoid any compression artefacts. A resolution of 256x256 was not chosen, as it has been shown that low resolution textures can harm the perceived quality of the 3D object regardless of the polygonal resolution [224]. 3D models without textures at various polygonal resolutions were also used in the paired comparisons, to investigate if higher polygonal resolution models could be perceived as a better perceptual experience than 3D models with textures. A total of 64 stimuli were generated for the pair wise experiment with differing polygonal and texture resolution (4 polygonal resolutions * 4 texture resolutions * 4 different objects). Table 5 and 6 provide full details on the different parameters for the experiment.

Vertices/Triangles	Ammon	Shakespeare	Relief	Brooch
100%	655,906 / 1,297,076	517,686 / 999,738	121,191 / 249,193	194,110 / 382,762
70%	461,131 / 907,952	367,492 / 699,816	87,812 / 174,435	136,693 / 267,932
40%	265,955 / 518,830	216,343 / 399,894	50,433 / 99,677	79,223 / 153,104
10%	68,977 / 129,706	60,037 / 99,972	13,054 / 24,919	20,909 / 38,276
Subjective	363,654 / 713,392	292,115 / 549,856	68,855 / 136,590	107,978 / 210,518

Table 5: The polygon resolution for each object at differing decimation levels

Objects	Original Resolution	Decimated Resolution	Percentage Reduction from Original
Ammon	2,523,883 / 5,040,129	655,906 / 1,297,076	25.98%
Shakespeare	3,399,624 / 6,799,264	517,686 / 999,738	15.22%
Egyptian Relief	249,966 / 498,383	121,191 / 249,193	48.48%
Anglo-Saxon Brooch	637,936 / 127,5876	194,110 / 382,762	30.43%

Table 6: Original and new high resolution decimation

Figure 3: Digital textured Representations

5.3.2 Subjective Stimuli Generation

The subjective part of the experiment plays a very important part in this research, it is not only attempting to quantify what is a good perceptual experience for users, but also allows for a direct comparison of the digital artefact versus the artefact. The stimuli that were generated for the subjective experiment, not only took in the perceived quality of the 3D mesh and texture, but also how it may be perceived in the environment. This allows cultural institutions to fully exploit the 3D dataset, communicating the significance of the historical artefact, while presenting a 3D digital artefact that offers a perceptual experience that is very similar to the artefact.

Section 4.4.2 3D Quality Subjective Assessment highlights the fact there is little work that focuses on subjectively evaluating the quality of 3D meshes with textures or other elements. The 3D model is not the only element within a scene; it includes different types of textures, lighting, and shader models. These all combine together to create the final render or tone mapped image. Metrics used to evaluate the perceived quality of a mesh focus on either the surface features and/or artefacts that may occur due to certain processes, and rarely consider how the final render would look when it is presented in a real-world scenario. This could include a sky box, global illumination, tone mapping, reflections and lens flares calculated within the scene.

That is why the stimuli generated for the subjective part of the experiment would need to consider these elements as well; therefore a 2D image metric was used to evaluate the meshes within the scene compared to a reference model and texture. Through a combination of the meshes generated using the simple method of Pan et al. [226] for the paired and a 2D image metric to generate the subjective model polygonal and texture resolution. The 2D metric chosen was the state of the art 2D image metric, HDR-VDP2 [237] image metric. HDR-VDP-2 is a 2D image based metric for real-world scenes, catering for complications and multitudes of parameters. HDR-VDP-2 is capable of measuring the visibility and quality metric, detecting differences in images across a variety of lighting conditions [237]. This would allow for scaling to be controlled via a metric that does not take into account the subjective nature of human evaluation and can be used tested within the virtual environment.

There are metrics which evaluate images that contain HDR images and can be explored in [216]. This section is concerned with the HDR-VDP-2 metric, which offers a more complex visual metric, similar to the Human Visual System metric [237]. The HDR-VDP-2 quality metric provides the capability of being able to detect the differences in images between large ranges of lighting conditions. The metric evaluates two images of either HDR or LDR as inputs (one being a reference image, and the other with simplifications or distortions artefacts) and evaluates the images to predict the differences between the images and provide a Mean Opinion Score to rate the visual quality of the image [237]. It is also one of the state of the art 2D image metrics, providing improved results over the HDR-VDP [249] and other metrics [196]. It also provides Matlab [250] code that can be run to compare images that have been created from applications or experiments.

Using the provided code for the HDR-VDP-2 Metric, 68 Images were captured of the different polygonal and texture resolutions and compared using the code. The metric also allows for different types of colour encoding. The monitor that was used in the experiment was calibrated for the RGB.BT.709 colour space and the captured images were compared using this colour encoding within Matlab. However, the images were also compared using the sRGB colour encoding and these can be found in the appendices. The captured comparisons were then compared for each model, to choose a possible resolution for the subjective test. These can be found in the appendices in section F Images computed for experiments conducted in chapter 5 on page 188. Using the brooch as an example, which can be seen in figures 6 and 7, there is little difference between 100% geometry and 70% respectively, but there are major differences in at 40% which is illustrated in figure 8. This was common across all of the objects, thus the subjective resolution was chosen to lie between the 40% and 70% polygonal resolutions at 55%. The chosen texture resolution was 1024x1024 as there was little difference between the texture resolutions. These chosen resolutions were created through the use of Pan et al. [226] using three independent observers, and the use of the HDR-VDP2 [237] 2D image metric. The resolutions of these subjective datasets are roughly 25% of the original 3D dataset, allowing them to be more manageable both memory and computational wise. These 3D datasets may offer a good perceptual experience for users, but are far removed from the accurate 3D original datasets. This provides a solution for cultural institutions that may be worried about losing control of their IP when disseminating and sharing their datasets.

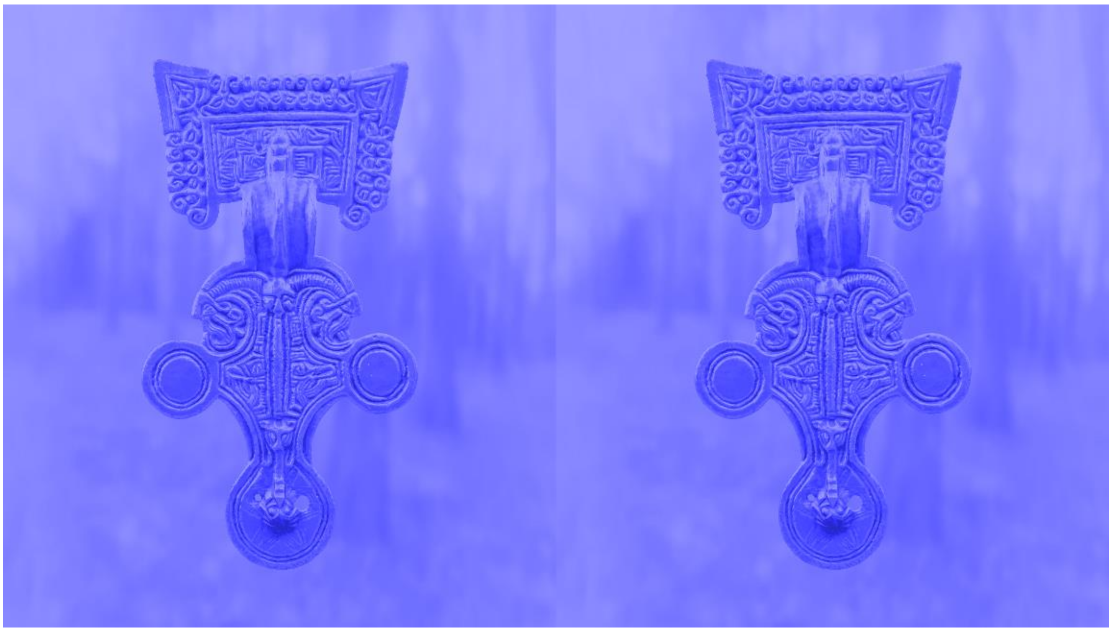

Figure 4: Reference image and 100% polygonal and 1024x1024 px texture resolution

Figure 5: Reference image and 70% polygonal and 1024x1024 px texture resolution

Figure 6: Reference image and 40% polygonal and 1024x1024 px texture resolution

5.3.2.1 Rendering Parameters

In trying to achieve as physically realistic rendering as possible using OpenGL as well as providing a real-world scenario regarding the display and visualisation of the 3D datasets, an OpenGL renderer was created for the experiment. It would allow for two models to be rendered side by side for comparison, using the same lighting and shaders to create the final tone mapped render. However, in the creation of such an experiment, there are many parameters that needed to be considered that could bias results and affect how people may perceive the visual quality of the created datasets.

5.3.2.1 Lighting

As discussed in the lighting section of the quality testing design in chapter 4 on page 76 lighting is extremely important for the visualisation of 3D models and the perceived quality of a 3D model. For this experiment the lighting was created using a HDR image based lighting that offering static realistic lighting illuminating the 3D model without obscuring the details. The HDR image that is used to illuminate the object is also used as a sky box.

5.3.2.2 Background Image

As discussed in the background section of the quality testing design in chapter 4 on page 76, a background may affect the perception of a 3D model by obscuring the boundaries of the 3D model. Due to the lighting being image based there could be confusion if there is no background that reflects the lighting environment and or is different from the actual image that is being used. The HDR image used for lighting is also used as the skybox image.

5.3.2.3 Material and shading

While there are many graphical applications that use a variety of complex or simple techniques to contribute towards a 3D models surface material (texture, bump mapping, ambient occlusion), they are rarely used in the subjective evaluation. This is discussed in further detail on page 76 in chapter 4 in the texturing and shading section. This study will be using testers who are visiting a museum, and asking them about their perception of the digital object with regards to the original. As we are ascertaining what a good experience is, the use of techniques that are common in computer games (Image based lighting, tone mapping, normal mapping, Screen Space ambient occlusion and more) will be used. This is to provide a good perceptual experience for the user, as well as provide a similar experience to when they have interacted with 3D models in games for instance. Thorn et al. [243] Found that via 2 studies that offering a non-shaded and Lambert shading model showed, that non-shading allowed users to confidently choose higher quality 3D meshes However the Lambert shading Model, created a masking effect, where participants were not as confident in their choices regarding what was the higher quality mesh [243].

5.3.2.4 Interaction

To interact and evaluate a 3D model properly, users will be allowed to view the models from different viewpoints. This will be achieved by allowing the user to freely rotate and control the camera [225, 230, 243, 244]. More details can be found in the interaction and animation section in chapter 4 on page 76.

5.3.2.5 Types of Objects

As is mentioned in section 5.2.2 Object Selection and preparation, there are many reasons for the chosen 3D models. However, there is the possibility of bias being introduced to the experiment due to the object, which was discussed in chapter 4 on page 76**.** Yet for these experiments the 3D models that are being used within the experiment are man-made, and easily identifiable in function and form, so this may introduce bias into the results. However, this is of little concern, as most objects within cultural institutions collections will be manmade.

5.3.2.6 Masking

A texture or texture technique that is applied to a 3D model may occlude geometry and affect our perception for said 3D object. Studies by Pan et al., [226] and Rushmeier et al., [224] studied this effect and how it affects our perception. This study will be investigating something similar, but focusing more on the aesthetic appearance and how our perception is affected and if a simple colour can be better choice for representing a 3D cultural heritage objects over say a textured 3D model.

5.3.2.7 Levels

As this study is testing both the reduction of polygons and texture resolution and as well as non-textured objects, the levels of reductions need to be selected accordingly. Too few levels for comparisons and the results may not provide good enough test data, while too many may provide unfeasible to expect someone to undertake. There are some studies, that use 3 levels [220], 4 levels [243], and up to 6 levels [245]. However, due to the number of comparisons between resolutions and textured and non-textured, there are 4 levels for the decimated polygon models (10%, 40 %, 70%, 100%). This offers a good polygon reduction, especially after the using the method by Pan et al. [226] to reduce the polygonal resolution of the original 3D dataset. The texture resolution would follow similarly; starting at a resolution of no texture, 512x512, 1024x1024, and 2058x2058 pixels.

5.3.2.8 Screen and Model Resolutions

The screen that is to be used in the experiment is a touch 27 inch touch screen monitor, offering a resolution of 1920 x 1080 pixels. The models are scaled to fit within the screen showing as much as possible of each model within the limited space of 960x1080 pixels, offering the best viewing of the 3D datasets and the textures. The monitor is also configured to a RGB.BT 709 colour encoding profile, as is recommended in the ITU-R-BT.500-13[222].

5.3.2.9 Restrictions

Taking into consideration the cuts that have been placed on NML budget and available equipment, the experiment was conducted with some restrictions in regards to objects and equipment. The experiment was conducted within one of NML's museums galleries. For the rendering and display of the 3D models, a laptop with an Intel Core i7-2640M CPU at 2.8GHz with 8GB of RAM and a Nvidia Quadro 1000m graphics card. The display used is a touchscreen monitor, that is becoming more common in museum galleries using a ITU.BT 709 colour encoding profile.

The interactive viewer was created with C++, OpenGL and GLSL which would incorporate: diffuse, normal and specular mapping, tone mapping, image-based lighting, ambient occlusion and motion blur when an object is moved and auto keying for the lighting. The interactive viewer would also allow the user to rotate, pan and zoom into the 3D model.

5.3.2.10 Perceptual Quality Estimation

There are many ways to estimate the perceptive quality of 3D models and textures, where some are limited in that they are not able to take into account the aesthetic or naturalness appearance of a model. The evaluation of 3D models is still a young field that has explored both automatic and subjective approaches. A subjective approach concerns us, where many approaches assess using different data metrics including: absolute ranking, double stimulus, ranking or pair wise comparisons. A study by Mantiuk et al. [223] compared these various approaches based on their sensitivity and time to complete. The forced pair wise comparison, offered the lowest variation in score between participants and the shortest time to complete each experiment. This is due in part to the simplicity of the task of just comparing 2 objects, and why it has been chosen for this study.

Even though there will be a subjective comparison method of ranking for objects, it will be difficult to determine how real the user felt it compared to the original. This is because that the information captured through paired methods is comparative [243]. Therefore there will be a second stage to the experiment, where a chosen resolution for the model would be chosen via a metric such as HDR-VDP2 [237] metric to measure noticeable changes with a subjective evaluation with a questionnaire.

5.3.3 Experiment Design

The study is split into two parts; one is a binary forced choice comparison experiment, and the other is a subjective experiment each slightly differs in what is asked of the participants in the study.

5.3.3.1 Setup

The experiment would be set up in one of the gallery spaces within one of NML's venues, allowing some privacy yet access to the collections. The set up for the experiment would be a touchscreen monitor to display the 3D models and a mouse and keyboard for input if they do not wish to use the touchscreen. The artefact the 3D dataset represents will also be close at hand for the users to compare against.

The interface for the experiment is minimal, showing only the 3D models within the virtual environment. The user would be able to rotate the models with the left mouse button, pan the 3D objects with the right mouse button. The user would be allowed to use the mouse scroll wheel to zoom in and out on the objects. The lighting image used is fixed and cannot be changed, but illuminates the 3D model, highlighting the details of the 3D object. To select the preferred model, the participant would select their choice on the keyboard, to choose either the right or left 3D model. When a participant chooses their input, the models will automatically change to the next one, and will continue through all of the comparisons. When the comparisons end, it would change automatically to provide instructions before the subjective experiment. There are also are clear and simple instructions provided at the start of the experiment taking the participant through the procedure for the experiment. They would also be provided with a simple demonstration comparing a textured and non-textured object, allowing them to adjust to the controls. There would also be a demo that would allow the participant to engage with the 3D models, and manipulate and choose one. This would not be counted towards the data. There are also black screens between experiments, to allow users to rest their eyes, and view instructions for the next part of the experiment.

5.3.4 Pair Wise Experimental Design

The first experiment is a forced binary comparison test, as this has been shown to acquire more accurate results, and be less time consuming [223]. Participants are asked to compare two randomly selected models, and choose either the right or the left one based on a simple question. "Compared to this artefact, which one do you prefer?" The semantics for this question are simply trying to reduce the bias in the results. The question also encapsulates what the participant is expected to do for this part of the experiment; comparing different resolutions of polygonal models and textures/without a texture and what they believe is more visually appealing and engaging for them. This experiment captures data on how important the texture and polygon resolution is in relation to our perceptual experience with the 3D dataset. The users are not given a time limit on deciding between the right and left, and can freely manipulate the 3D model. The user then selects their desired choice using input from a keyboard that is provided.

However, the main issue with forced comparison tests is the large number of comparisons that need to be made comparing polygonal resolutions and texture resolutions (4 decimation levels, 3 texture levels and no texture. 4 total texture levels.) in a forced comparison test. Using the formula for forced comparison tests: 0.5(N*(N-1)), this would result in 120 comparisons for all possible permutations if N = 16 and 240 comparisons to test for the reversed permutations as well. This would be impossible and unfeasible to ask museum visitors to complete, as it would take too long to complete in a reasonable time. Yet, the number of comparisons can be reduced by using a near completion comparison table based on an efficient sorting algorithm [251] or a self-balancing binary tree [223]. An efficient sorting algorithm such as a quick sort, or a self-balancing binary tree can reduce the number of comparisons down to 64 comparisons approximately using $N \, Log_2 \, N$, if N = 4. Used within the experiment the sorting algorithm sorts which items need to be compared to each based on the previous comparisons in the same experiment session. This is due to the sorting concentrating on comparisons around very similar images, which are more sensitive to subjective variations. The algorithm would be based off a very simple assumption: If A > B and B>C, then A is greater than C automatically, allowing for reduced comparisons. However, the recorded data will need to be screened for discrepancies that may have been the result of pressing the wrong button or randomly choosing during the experiment. The ITU-R-BT.500-13 [222] provides a guide on how to remove the results with discrepancies. This involves, checking the data to make sure the participants' scores do not lie ±2x outside the standard deviation range, rejecting the results when 5% of their data lies outside this range and if the values for the other values do not exceed the bound of absolute difference range by 30% [222].

With enough participants, the incomplete comparison table becomes more accurate and starts to reflect the completed comparison table [251]. This is a factor for accurate results, but is limited by the amount of time that a visitor may wish to spend their time helping with the research. To address the issue of accuracy of the recorded data from the self-balancing binary tree, a small control group was used to record a full comparison matrix, resulting in 120 comparisons. The group consisted of 5 people, allowing for the experiment to be completed fairly quickly and efficiently, and offer an accurate comparison table.

5.3.5 Subjective Experimental Design

A second yet shorter experiment is also to be completed by the participant, asking the user to interact with a 3D model (the stimuli generated based on the images calculated using the HDR-VDP2 2D image metric [237]) and asked open ended questions for the user to complete. The data collected would provide an insight to what is an acceptable reduction in both polygonal and texture resolution for a good perceptual experience compared to the real-world artefact. This would give invaluable data on the trade-off for offering a good perceptual experience without having to use the full polygonal model and texture resolution. The participants are also asked to answer the following questions:

- *How does this 3D model and texture compare to the real object on a scale of 1 to 10? 1 being the worst and 10 being the best.*
- *What do you think this 3D model is made out of?*
- *How important is the texture for you when interacting with this 3D model?*
- *Would you like the option to choose to display and remove the texture from the 3D model?*
- *What would you prefer interacting with: the original/replication or the 3D model?*
- *After this experiment, would you like to learn more about the collections, or the 3D models that the National Museums have?*
- *Are there any additional comments you would like to make, either about the first or second part of the experiment or anything about your time here today?*

The questions are designed to assess how the user perceives the quality of the decimated texture and mesh and if they understand what the material is made of. The questions regarding the use of textures, is to provide evidence of whether or not a texture is important in the presentation of the 3D dataset. Further providing data to help disseminate and share 3D datasets. However, unlike the previous experiment, there will be no need to filter the results. This is due to the simple nature of the questions, which are not misleading but are more open to interpretation.

5.3.6 Participants

To assess that the use of technology is warranted to display 3D cultural heritage and positively impact visitor's experiences with these digital artefacts [252], unique visitors were chosen who were visiting the World Museums Weston Discovery Centre. The participants are visitors are individuals with an increased interest in the past in general [253] and/or have a desire to connect with their cultural heritage and learn more about history [254]. A total of 70 participants took part in the experiment, an equal split of males and females (35M, 35F). The participant's age ranged between 18 to 60, with 31 participants aged between 18 to 25, 21 were aged between 26 to 33, 11 between 34 to 41, 4 between the ages of 42 to 49 and 7 participants aged between 50 plus. Each participant had either normal or corrected vision. The users had a mixture of experiences with 3D graphics but mostly having very little experience with computer graphics. It should be noted, however, that there may be limitations and biases introduced into the results of this study, that do not reflect the full diversity and variety of visitors to NML. The main bias is due to 74% of the participants being aged 18 – 33 which only represents 25% of NML visitors. The study also has no visitors above the age of 60, where visitors aged 65 and up represent 15.7% of visitors. However, participants were chosen and approached if they were by themselves, to reduce the bias in the results due to possibly being influenced by others. This limited the number of participants that could be approached as they only represented 12.6% of total visitors to NML venues and only 8.5% of World Museum visitors. Visitor statistics can be viewed in full in the appendices on page 155.

The average time to complete both experiments for a single object was approximately 6 minutes, with 3.5 minutes spent on the forced comparison test. The average time to complete could be slightly exaggerated as the timer did not stop for the overall time, until the OpenGL application was closed. The applications were rarely closed until the participant left as they either forgot to press the enter key to end the application, or asked questions after completing the subjective questionnaire.

The studies were conducted on in two weekends in November, to maximise the number of participants for each study, where a different object was used on different days. 15 Participants rated the Anglo Saxon Brooch, 15 rated the Egyptian Relief, 20 reviewed the Zeus Ammon Bust, and 20 rated the Shakespeare bust. The Brooch and relief received fewer participants, due to time constraints in the galleries and difficult nature of asking visitors to participate in studies without an incentive. They conducted the experiments on a laptop with an Intel Core i7-2640M CPU at 2.8GHz with 8GB of RAM and an Nvidia Quadro 1000m graphics card and using a 27 inch touch screen monitor within the Weston Discovery Centre. The study was conducted with the participant viewing the monitor that rendered the 3D digital artefact, with the artefact it was based on to their right next to the monitor within touching distance. The keyboard to navigate was placed in front of them. The experiments were all conducted on different days with different models. Day one was conducted with the Anglo Saxon Brooch, where the results can be seen in section G: Anglo Saxon Brooch on page 341 within the appendices. Day two consisted of the Egyptian relief, the results can be found in the appendices on page 343 in section G: Egyptian Relief. The third day consisted of the Zeus Ammon Bust, and the results are broken down in section G: Zeus Ammon in the appendices. The Shakespeare bust was compared on the last day, and the results may be found in section G: Shakespeare Bust of the appendices on page 341.

5.3.7 Computing Scores

As each participant is shown a pair of 3D models to compare, they are forced to choose what they perceive to be a better representation (even if they do not see a difference). With no time limits on how long to make their decisions and the simplicity of the study, it can be assumed that the results provide accurate and reduced variation of results between participants [223]. However, as this study is using a reduced comparison test versus the full design, the data can be noisier and rarely reflect the true ranking of the actual full comparison table. However as has been shown by Silva et al. [245], a large number of observers can converge to be similar to the full design, while reducing the number of comparisons and time taken to complete the experiment. However, calculating the scores using a reduced comparison test, it must be assumed, that the quality of the estimates may not entirely be balanced: Model A is better than Model B and Model B is better than model C, therefore Model A is better than Model C. This does not need to be the case in the full comparison table. It is a very common violation of the assumption in the full comparison table, especially when images that are viewed are similar.

Assuming the above, then it is trivial to work out the global ranking and order of the 3D models, their position in the rankings of the model would be equivalent to the number of votes that each model has received. The highest rated model would have a score of 15 when compared with the other 16 stimuli, this is also similar to how Mantiuk et al. [223] conducted their study.

As the reduced comparisons are using the votes for each model as a way to ascertain their global ranking, a method of scoring is needed for the complete comparison table. A way of unifying these scores is also needed, and is suggested to use the Z-scores for pair wise comparison [223]. Which will be calculated, but to allow for direct comparisons, a preference score or the use of Thurstone's Law of Comparative Judgment, Case V [255], would be used to obtain scores for each models ranking. The preference score would be calculated for each mesh, where the formula is:

$$ps = (ta - tb) / (ta + tb)$$

Where *ta* and *tb* are the number of times the participant preferred mesh A over B. The scores for both the reduced and full completion test would then be processed in a one way ANOVA, to calculate a correlation between texture and polygonal resolution, and if these closely match each other. A post hoc Tukey honestly significant difference (HSD) test would be applied to the results to show the significant results between the stimuli. A Bonferroni correction or a Pearson correlation tests could also be used.

However, there may be some scores that may not be coherent within the data, due to not understanding the experiment, pressing wrong keys or randomly guessing. With an experiment with a low number of participants, these data can easily be observed. However, the ITU-R-BT.500-13, provides a way to screen for these discrepancies. It involves counting the number of trials in which a participant's results, lie ±2x outside the standard deviation range, rejecting the participant's data, when 5% of their data lies outside the range and the calculated values for the other trials do not exceed the bound of the absolute difference range by 30% [222].

As is common with most studies, a 95% confidence interval for the mean is expected when displaying the mean scores in a graph comparing the data. The confidence interval is the range in which the true mean scores resides with a 95% probability [222]. If the confidence interval overlaps with another comparison, even if it is small, there is a small chance that the true mean lies within that range. We can then assume that there is not enough evidence to significantly decide which one is perceived as being significantly better. However, this does not mean that both produce equally good quality, just that there is no statistical difference in quality and more participants are needed.

This confidence interval range itself would also need to be adjusted due to the large amount of multiple comparisons. This is due to a large amount of comparisons can lead to a type 1 error and reject the hypothesis (a lot of false positives). This could be corrected using the Bonferroni Correction or a Tukey HSD criterion.

5.4 Results

The following section discusses the results of the study, across the four 3D cultural artefacts and how their polygonal and texture resolution relationship effects human's perception of 3D digital cultural artefacts. The detailed result provide an understanding of how people rate the different resolutions, but also allows to compare the subjective score using the method from Pan et al. [226] and the HDR-VDP2 image metric compares to the full comparisons. Other studies focus on the either geometry or texture artefacts [224, 226, 242], this study focuses on the whole scene, evaluating both the 3D model and the scene in which it is rendered. While it does not provide a metric to measure to provide a quality score, it provides results that reflect real-world expectations, with visitors who have little to no experience with computer graphics.

5.4.1 Screening Users

With a small group of users it is possible to screen the data for outliers that may influence the results and introduce inconsistencies. This may result in inaccurate results or impossible results, due to not understanding what was expected of them or randomly giving answers. However, for large numbers of participants this becomes harder to spot. The ITU-R-BT.500-13 [222] provides a screening guide on how to spot and remove these outliers. This procedure involves, checking the data to make sure the participant's scores do not lie 2x outside the standard deviation range, rejecting the results when 5% of their data lies outside this range and if the values for the other values do not exceed the bound of absolute difference range by 30% [222]. This means that if one of their scores is 2x above or below the standard deviation, the result is removed. This procedure has been applied to the participants of the study. The results of the Zeus Ammon Bust resulted in 5 user's results being removed out of 20 participants. The Brooch study resulted in 4 out of 15 participant's results being removed as well as 6 out of 15 for the relief study. The Bust of Shakespeare resulted in 3 out of 20 participants being removed. There could be a number of reasons for the numbers of participant's data being removed, from a lack of instruction or not fully participating with the study and randomly choose their answers. However, it should be highlighted that the participants were naïve museum visitors, who had very little experience with 3D graphics compared with a lot of studies, that use participants who have experience or relevant knowledge of computer graphics, which may reflect why they have low rejected data.

To back up the accuracy of the recorded data from the reduced forced comparison tests, a complete design for the 4 cultural artefacts has been conducted with a control group of 5 users to back up the results of the reduced comparison test.

5.4.2 Observers Agreement

With data being removed from the test data for the comparison test, it is therefore essential to scrutinise the agreement between users and their scores for the models. This is accomplished by computing Kendall's coefficient of concordance (Kendall's W), using the scores produced by participants, which assess the agreement among participants. The produced W coefficient lies between zero and one, where zero means there is no agreement among the participants, and one there is a unanimous agreement. The results are considered significant if the P-value is extremely low ($P<0.01$), and the null hypothesis is rejected so that there is no agreement between participants. Table 7 presents the overall Kendall's W coefficient and P-value for each object, for both the reduced and full comparison table.

Object	Kendall's W	P-Value
Reduced – Brooch	0.712591152	5.57877E-18
Reduced – Relief	0.517138707	4.82166E-09
Reduced – Zeus Ammon	0.252183007	9.07919E-07
Reduced – Shakespeare	0.344270303	2.55924E-12
Full – Brooch	0.675294118	9.43502E-06
Full – Relief	0.755294118	9.41871E-07
Full - Zeus Ammon	0.590117647	0.000100159
Full – Shakespeare	0.598588235	7.95747E-05

Table 7: Computed Kendalls W between users

The table shows that there is a strong agreement between the participants for the Anglo Saxon brooch and Egyptian relief for both the full and reduced comparison matrix ($W>0.5$) and this is confirmed with the low P-values. For the full comparison of the Zeus Ammon and Shakespeare Bust, there is also a strong agreement ($W>0.5$) and confirmed with low P-values. Yet, for the reduced comparison, this is not the case. Zeus Ammon has a W coefficient of 0.25 and Shakespeare has a coefficient of agreement of 0.34, this could be due to the number of comparisons being very similar and the nature of the reduced comparison test. While they have low agreement, among the users for the experiment, the small P-values still reject the null hypothesis and the results are significant.

5.4.3 Confidence Intervals and significance

The results that were recorded for each experiment will be displayed via a box plot to visualise the data and their confidence intervals. The confidence intervals are expected in addition to most displayed data as it provides the range of values in which the true value may reside with a 95% probability [222]. However, the confidence levels do not explain if the differences in scores are statistically significant or not. If there is a small overlap between confidence intervals, it is highly unlikely that the true value will lie in this region. Thus if there are overlaps between confidence intervals, it is not possible to say which is perceived as being of better quality, as there is not enough evidence to support this with an $\alpha = 0.05$. However, this does not prove that the overlaps produce as equally as good experiences, just that there is no statistically significant difference in their quality.

As there are a lot of comparisons between different conditions, the confidence intervals and the data needs to be adjusted post hoc after the One Way ANOVA. By increasing the number of comparisons between stimuli, there is an increasingly small chance of a Type 1 error being made (rejecting the null hypothesis), and the alpha is no longer $\alpha = 0.05$. The α approaches the sum of probability where is $\alpha = 1 - (1 - 0.05)^c$, where C is the number of comparisons. The data for this study was adjusted using the Tukey HSD criterion [249]. However, this could also be adjusted using a post hoc Bonferroni Correction to control the Type 1 errors.

5.5 Paired Comparison Results

In order to analyse the perceived quality of the 3D cultural artefacts with reduced texture and polygonal resolutions, a forced binary comparison test was completed with 70 human participants. Participants evaluated 4 levels of polygonal resolution against 3 texture resolutions and an additional no texture level. This resulted in 16 stimuli per Object to produce a global quality ranking of the perceived quality of the meshes. The comparison experiment was conducted using a reduced comparison test using a self-balancing tree similar to that used by Silva et al. [245] and Manituk et al. [223]. The four objects used to create the stimuli as mentioned above are an Anglo Saxon Brooch, Egyptian Relief, a Zeus Ammon bust and a Shakespeare Bust. Sixteen Stimuli were generated per object and as can be seen in table 8, it lists the distortions that have been applied to each stimulus, and provide a guide when looking at the graphs.

ID	Geometry Resolution	Texture Resolution
1	10% of reference resolution	None
2	40% of reference resolution	None
3	70% of reference resolution	None
4	100% of reference resolution	None
5	10% of reference resolution	512 x 512 pixels

6	40% of reference resolution	512 x 512 pixels
7	70% of reference resolution	512 x 512 pixels
8	100% of reference resolution	512 x 512 pixels
9	10% of reference resolution	1024 x 1024 pixels
10	40% of reference resolution	1024 x 1024 pixels
11	70% of reference resolution	1024 x 1024 pixels
12	100% of reference resolution	1024 x 1024 pixels
13	10% of reference resolution	2048 x 2048 pixels
14	40% of reference resolution	2048 x 2048 pixels
15	70% of reference resolution	2048 x 2048 pixels
16	100% of reference resolution	2048 x 2048 pixels

Table 8: Details about the distortions applied to objects

Due to the reduced comparison matrix being generated with a self-balancing tree, a control group of 5 participants was used to create a small complete comparison matrix. As a reduced comparison matrix can provide noisy data and require a large number of participants to give accurate results. There is also the possibility of data being removed using the ITU-R BT.500-13 guide [222]. Figure 9 and figure 10 were generated showing the mean for each object and their stimuli. As can be seen in these figures, they are very similar; with more participants in the reduced comparison test it would be possible to achieve more accurate results.

The perceived quality of the generated stimuli differed across each object. Each object will be discussed in turn, starting with the Brooch, Relief, Ammon and Shakespeare.

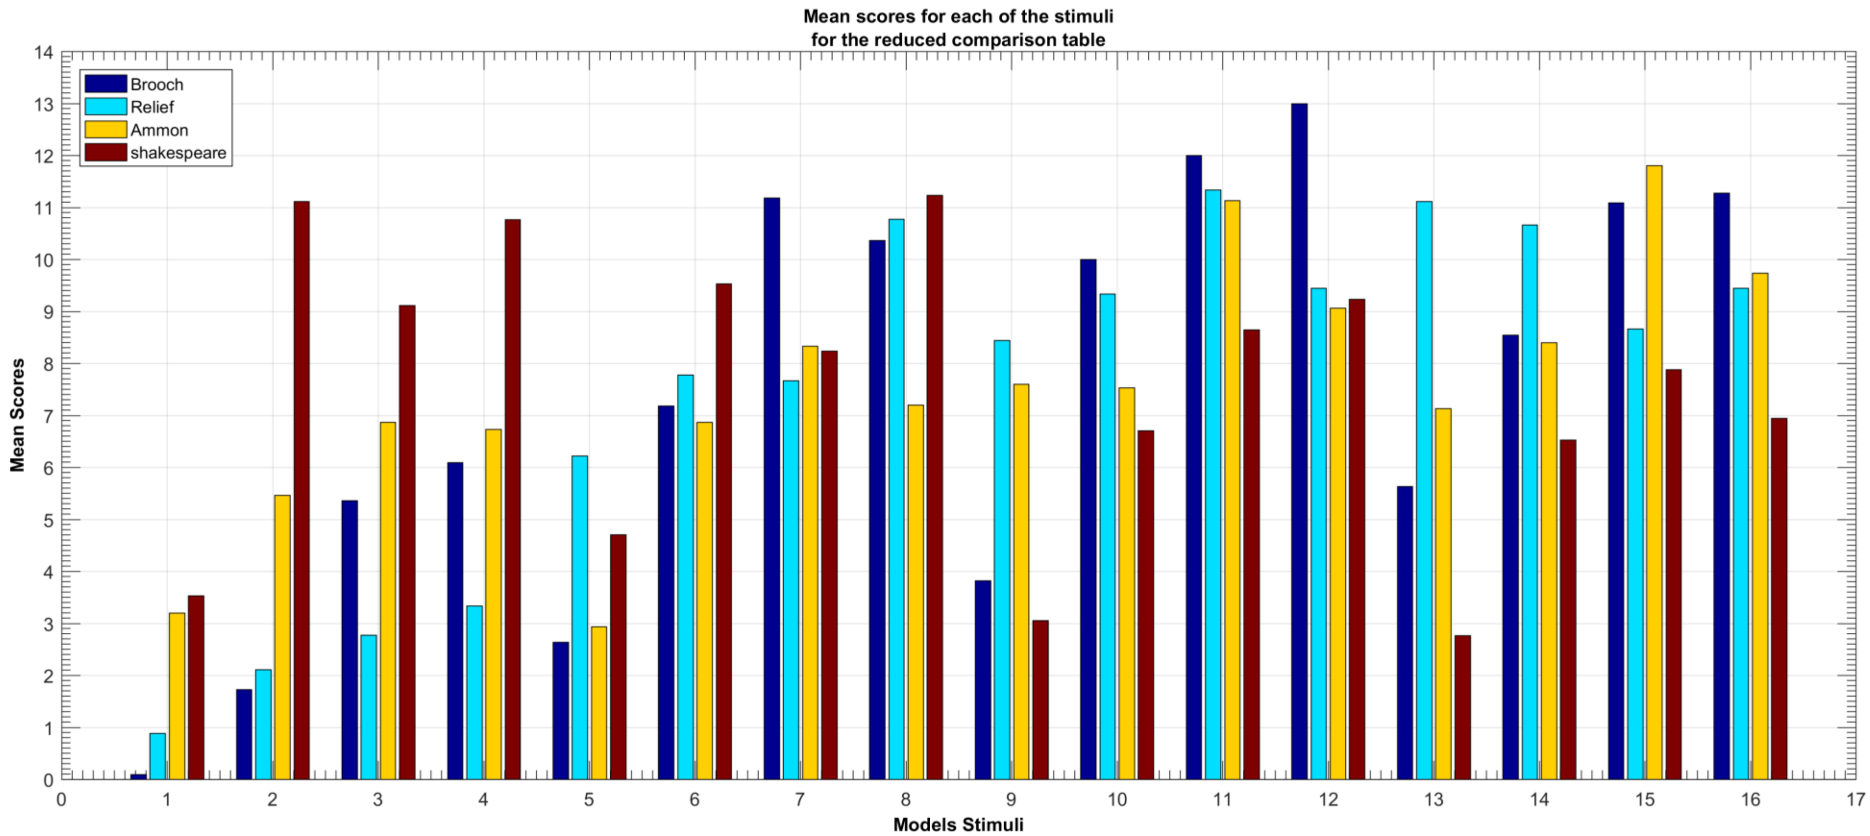

Figure 7: Mean scores for the reduced comparison test

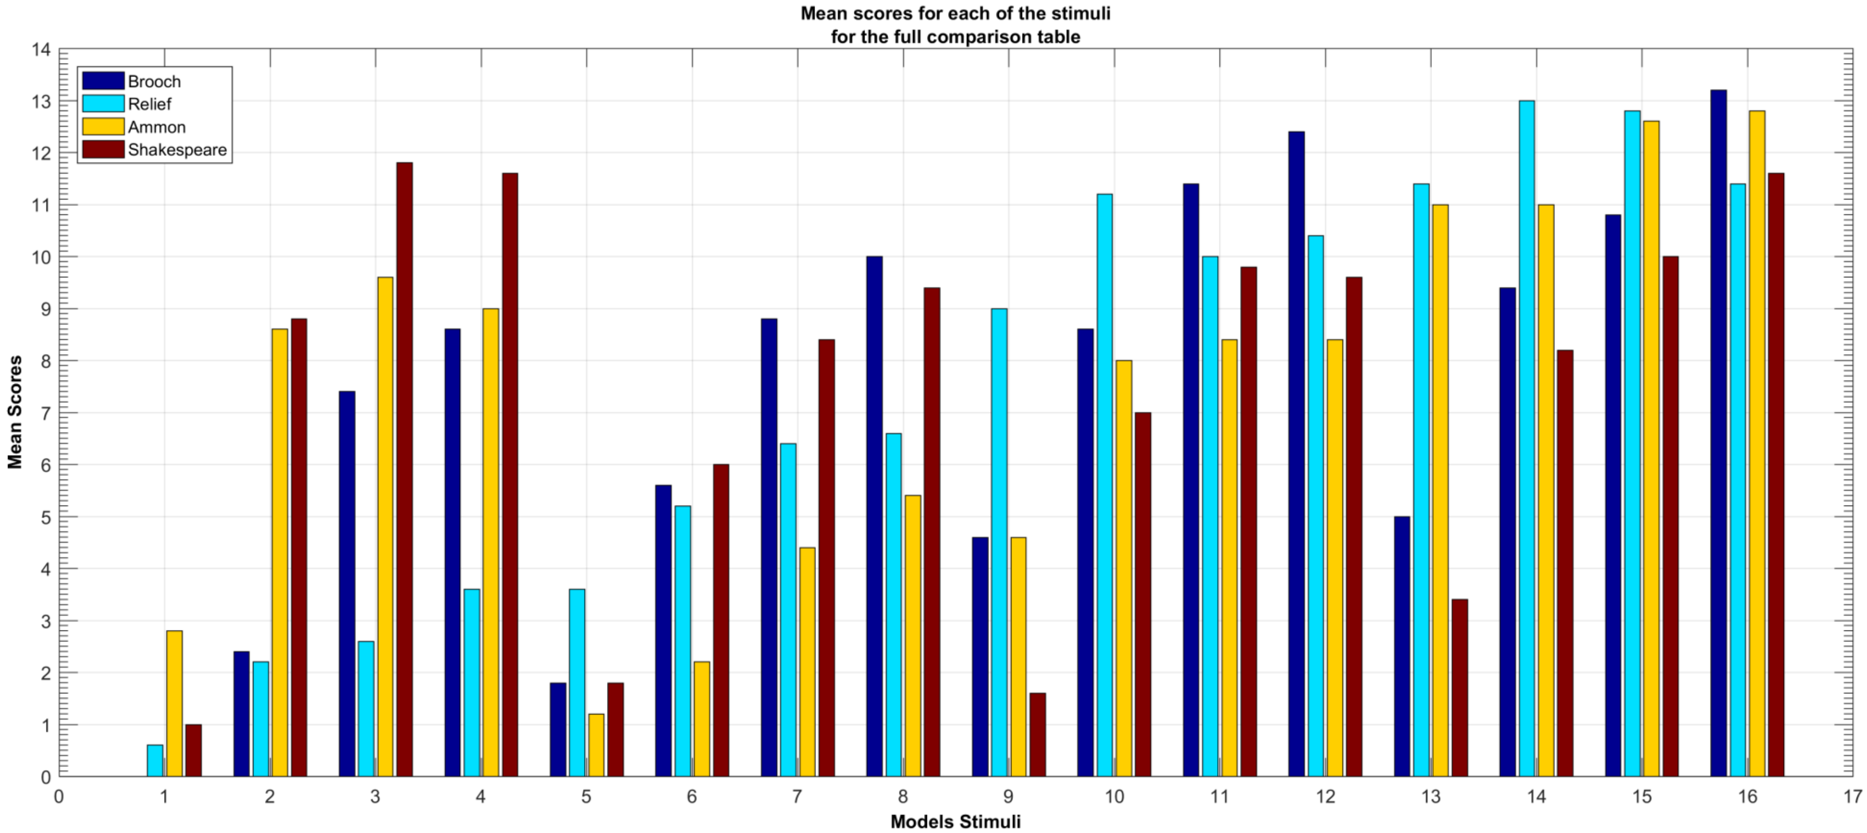

Figure 8: Mean Full comparison score

5.5.1 Anglo Saxon Brooch

The Anglo Saxon Brooch forced paired comparison experiment was completed with 15 museum visitors, with their data being screened using the ITU-R BT.500-13 guide [222] to remove outlier data. 4 participant's data was removed from the results analysis, due to being out of the range that was acceptable. A Kendall's coefficient of concordance was calculated to assess the agreement among the participants, resulting in a strong agreement with a W = 0.712 and P <0.01. A One Way ANOVA was used, to calculate a correlation between texture and polygonal resolution, resulting in significant results with a P value <0.05. A post hoc Tukey HSD criterion test was also applied to identify significant differences between the individual stimuli. The results of the One Way ANOVA is represented as a box plot in figure 11, the mean scores and confidence intervals in figure 12 and the results of the post hoc Tukey HSD criterion can be seen in figure 13. Table 9 presents the mean scores of each stimulus in the Tukey HSD test.

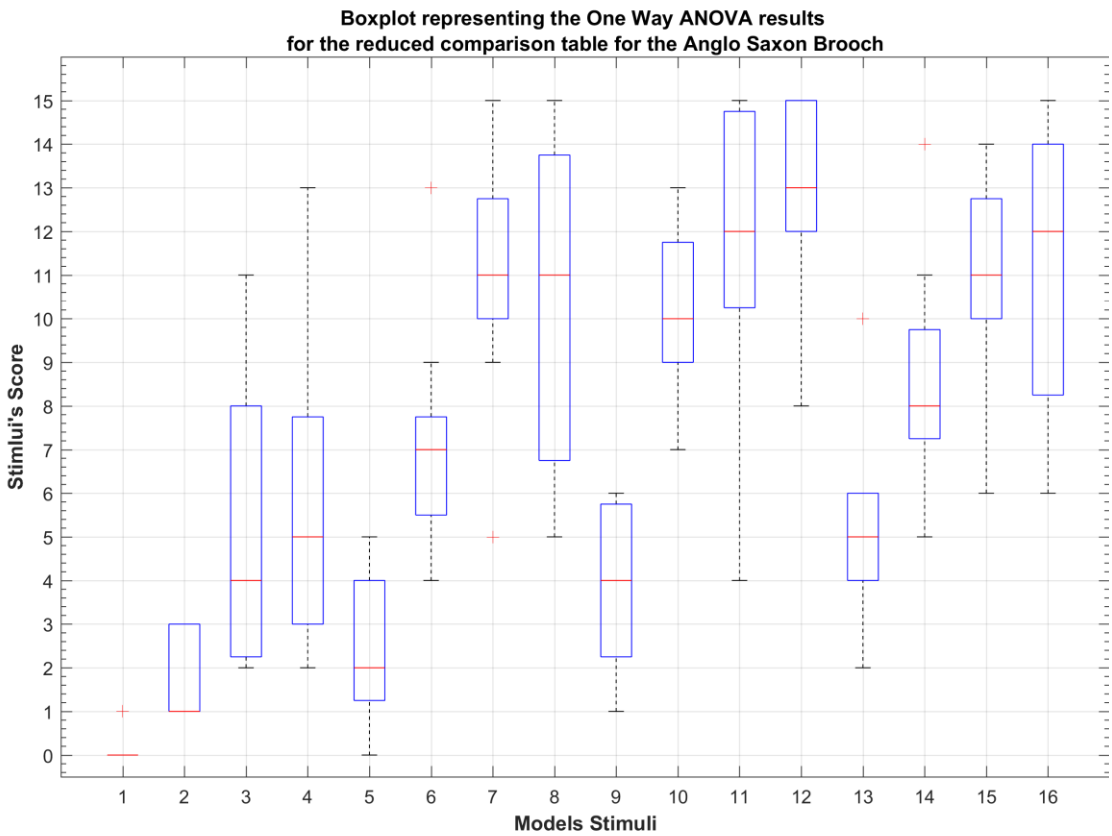

Figure 9: Results of the reduced comparison One Way ANOVA represented as a Box plot

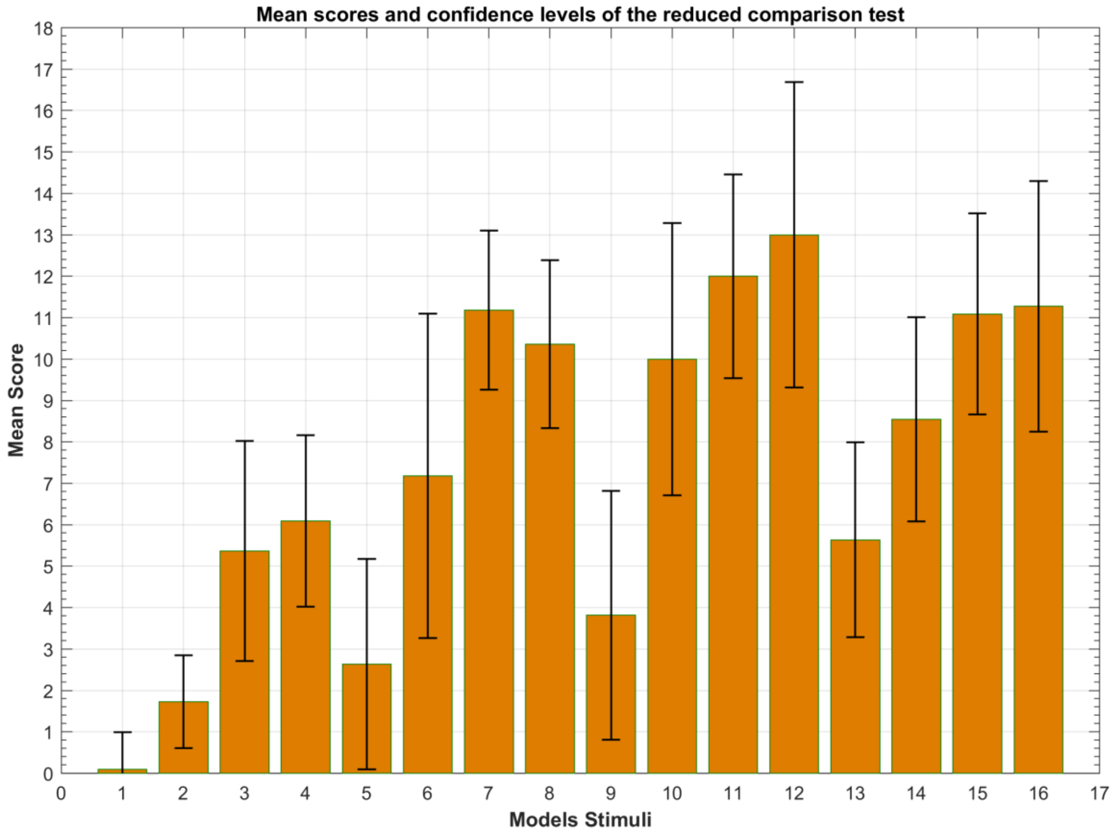

Figure 10: Mean scores and confidence intervals of the reduced comparison

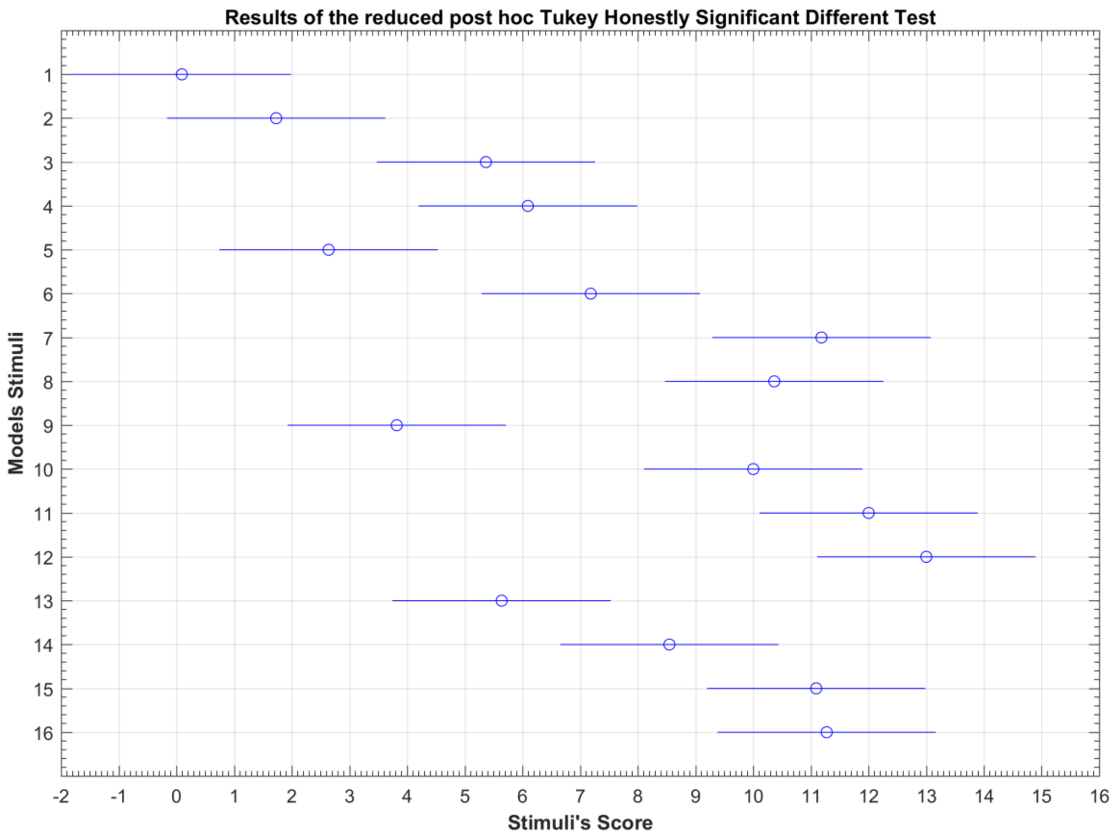

Figure 11: Results of a post hoc Tukey Honestly Significant Different Test

At a quick glance, it appears that the lowest geometry with no texture is perceived (Tukey HSD, Score = 0.9091, P <0.05) as the worst quality, with a higher polygonal and texture resolution being (ID 12) perceived as the best quality (Tukey HSD, Score = 14.89, P<0.05). The increase in both texture and geometry resolution appears to affect perceived quality.

However, when looking in more detail, the belief that an increase in both texture and polygonal resolution is not so clear-cut. The lowest perceived model is the lowest geometry with no texture (Tukey HSD, Score = 0.9091, P <0.05), however, it is not significantly different from the 40% stimuli with no texture or from the models with the lowest polygonal resolution and 512x 512 and 1024 x 1024 k texture resolution. While their mean scores are higher, their confidence levels overlap; there is no evidence to significantly decide which one is perceived as being of better quality. The 70% and 100% (ID 3 and 4) models (ID 3 HSD, Score = 7.257, ID 4, HSD Score = 7.984), while having mean scores less than meshes with textures, they are significantly better than models 1 and 2, yet it is not significantly different from Models (5, 6, 9, 12, 14) with texture resolutions of 512x512, 1024x1024 and 2048x 2048 with polygonal resolution of 10% and 40%. The highest rated mesh is model 12 (ID 12, HSD Score = 13) with a 100% polygonal resolution and 1024 x 1024 texture resolution, yet is not perceived as significantly different from models (7,8,10,11,15,16), which have a polygonal resolution of 70% and 100% apart from model 10 which has a polygonal resolution of 40%. There is no perceived difference between meshes with a polygonal resolution of 70% or greater with a texture applied. Though there is no significant difference, in the One Way ANOVA and Tukey HSD test, it does suggest that the increase in perceived quality is related to the polygonal resolution over texture resolution. However, there is not enough evidence to suggest that an increase in texture resolution increases the perceived quality of the 3D object.

ID	Tukey Honestly Significant Different Mean Score (reduced)	Tukey Honestly Significant Different Mean Score (full)
1	0.0909	0
2	1.7273	2.4000
3	5.3636	7.4000
4	6.0909	8.6000
5	2.6364	1.8000
6	7.1818	5.6000
7	1.1818	8.8000
8	10.3636	10.0000
9	3.8182	4.6000
10	10.0000	8.6000
11	12.0000	11.4000

12	13.0000	12.4000
13	5.6364	5.0000
14	8.5455	9.4000
15	11.0909	10.8000
16	11.2727	13.2000

Table 9: mean Scores from the Tukey Honestly Significant Different Mean Test

These results are supported in the full complete comparison matrix, completed by the control group, providing similar results to the reduced comparison test. Scores can be seen in table 9. There are only slight variations between the comparisons, such as the Highest Mesh (ID 15 HSD Score = 13.2, $P<0.05$) is perceived as better in quality from meshes with 70% and 100% polygonal Resolution with texture resolution greater than 512x512px. The 70% and 100% models with no texture are also not significantly different from meshes with textures, except the 70% and 100% resolution with 1024 x 1024px and 2048 x 2048px, which are perceived as better. However with these results, while there may be an overlap between confidence intervals, there is no evidence to significantly decide which one is perceived as being of better quality. However, this does not mean that both produce equally good quality, just that there is no statistical difference in quality and more participants in testing are needed.

The results of the full comparison tests One Way ANOVA is presented in Figure 14, and its post hoc Tukey Significance Difference Test is Presented in figure 15, with the mean score itself in table 9.

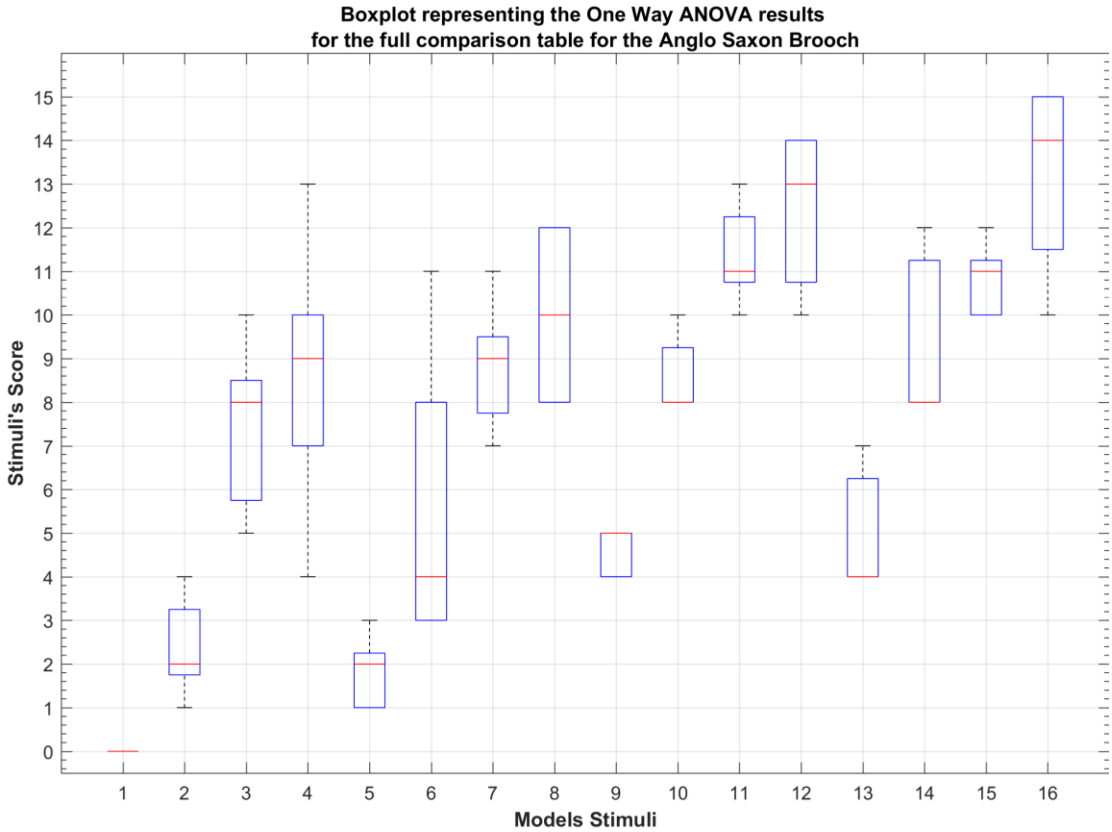

Figure 12: Results of the full completion matrix One Way ANOVA

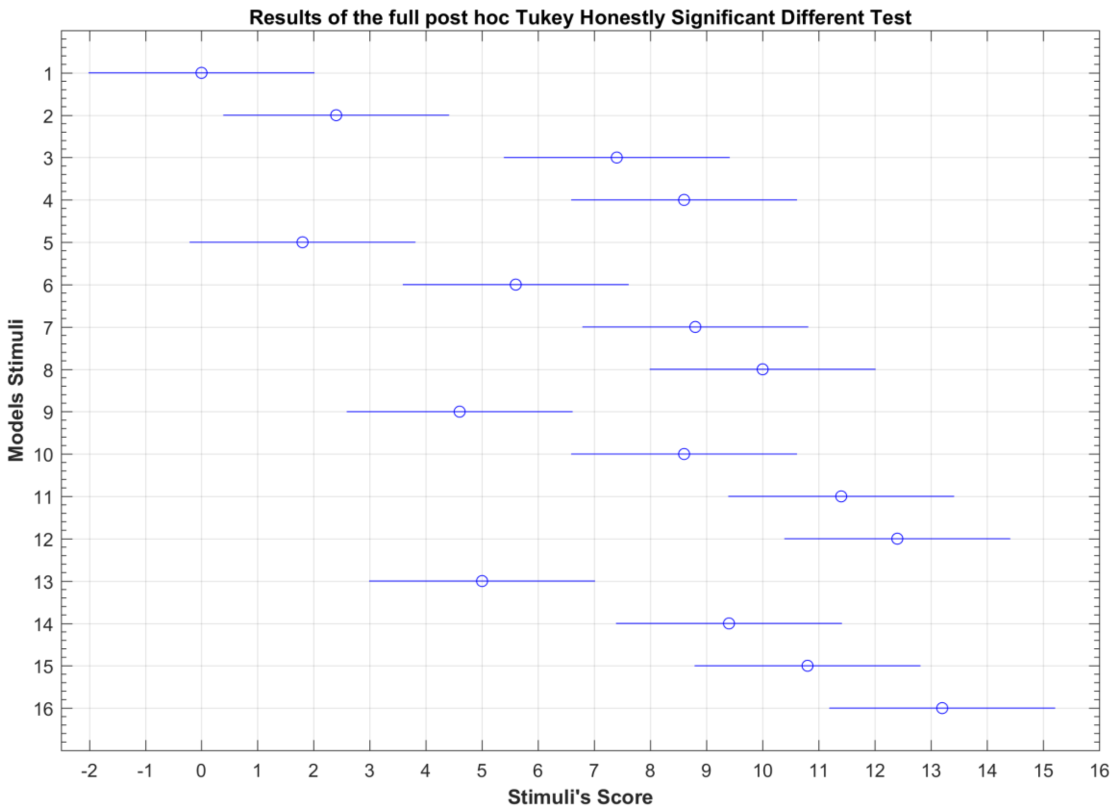

Figure 13: Results of a post hoc Tukey Honestly Significant Different Test

5.5.2 Egyptian Relief

The Egyptian Relief forced paired experiment, was generated with photogrammetry by Conservation Technology. This artefact was chosen due to it being too large and heavy to be handled or manipulated by the general public, unlike the digital model. The forced comparison experiment was completed with 15 participants but the usable data from this experiment was reduced to 9 using the ITU-R BT.500.13 screening guide [222]. Their data was removed from this analysis, yet the participants still had a strong agreement among themselves. The experiments Kendall's coefficient of concordance $W = 0.517$ and $P < 0.01$. The full comparison matrix, had a stronger agreement with $W = 0.755$ and $P < 0.01$. A One Way ANOVA was used, to calculate a correlation between texture and polygonal resolution for the Egyptian relief, resulting in significant results with a P value <0.05. The post hoc Tukey HSD test was also conducted on the data to identify significant results between the different stimuli. The results of the One Way ANOVA is represented as a box plot in figure 16, mean scores and confidence levels in figure 17, and the results of the post hoc Tukey HSD criterion can be seen in figure 18. Table 10 presents the mean scores of each stimulus in the Tukey HSD test.

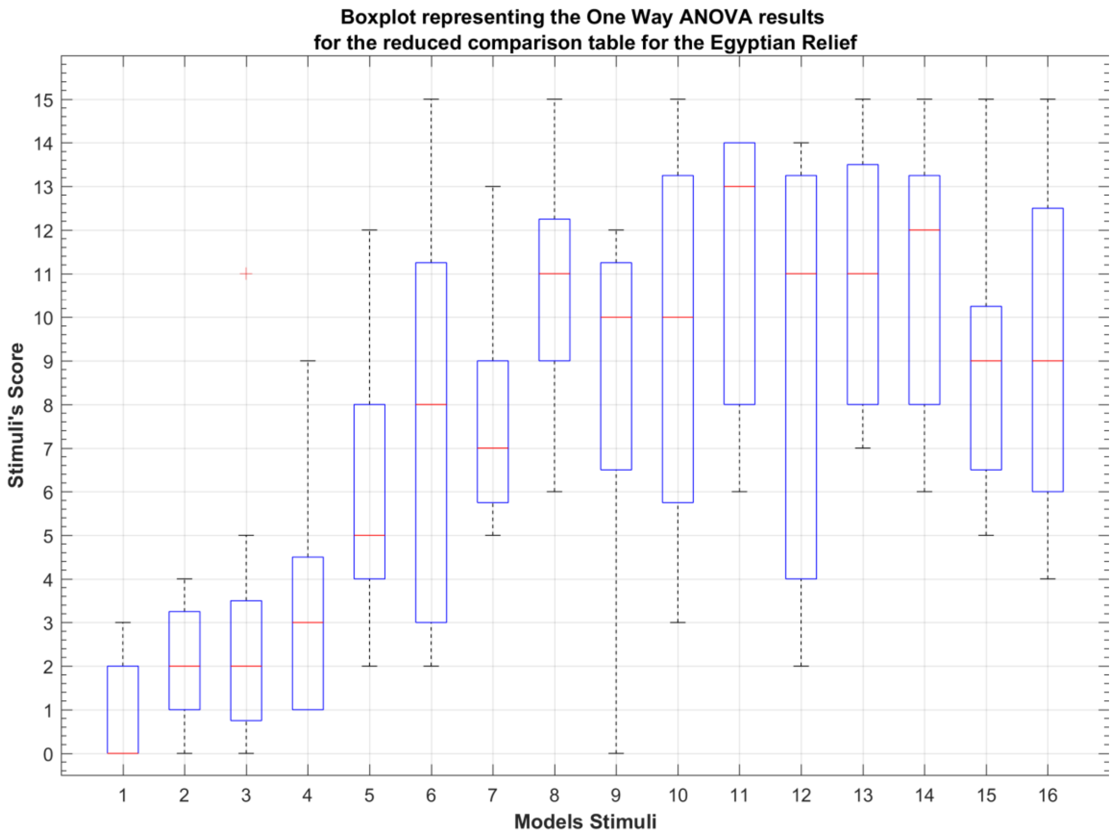

Figure 14: Boxplot representing One Way ANOVA for the resuced results of the Egyptian Relief

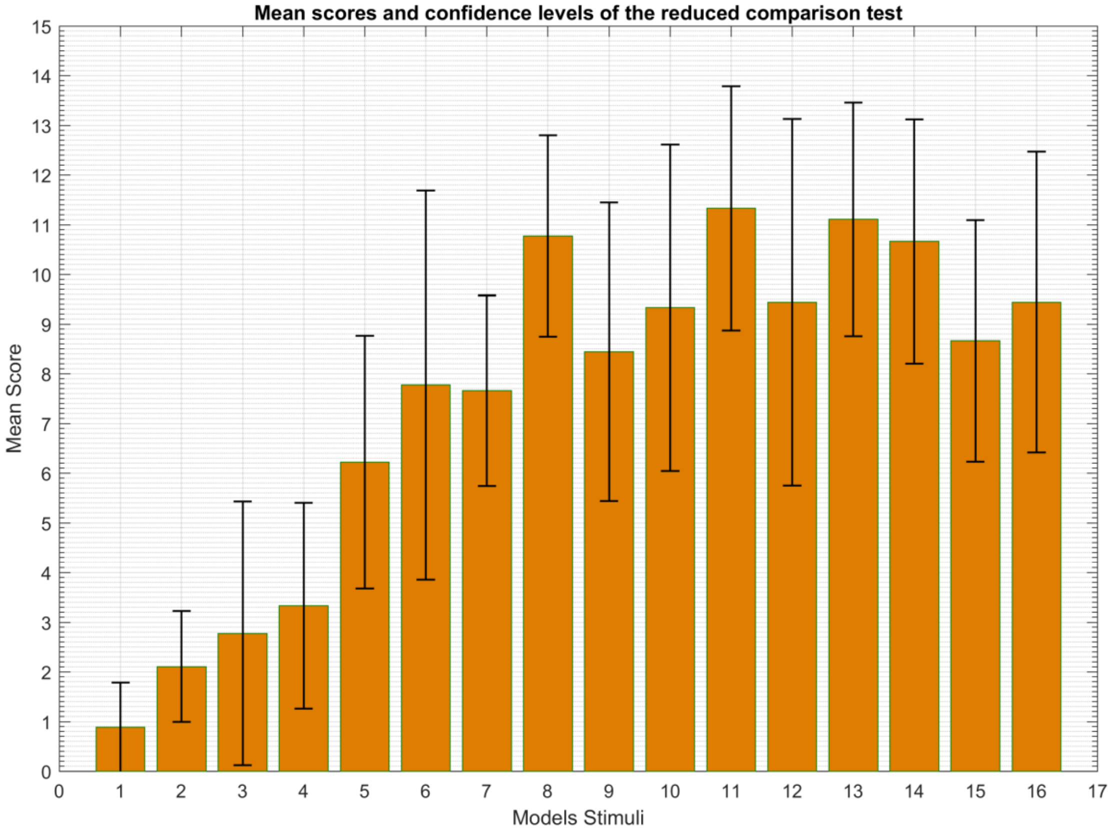

Figure 15: Mean scores and confidence levels of the reduced comparison test

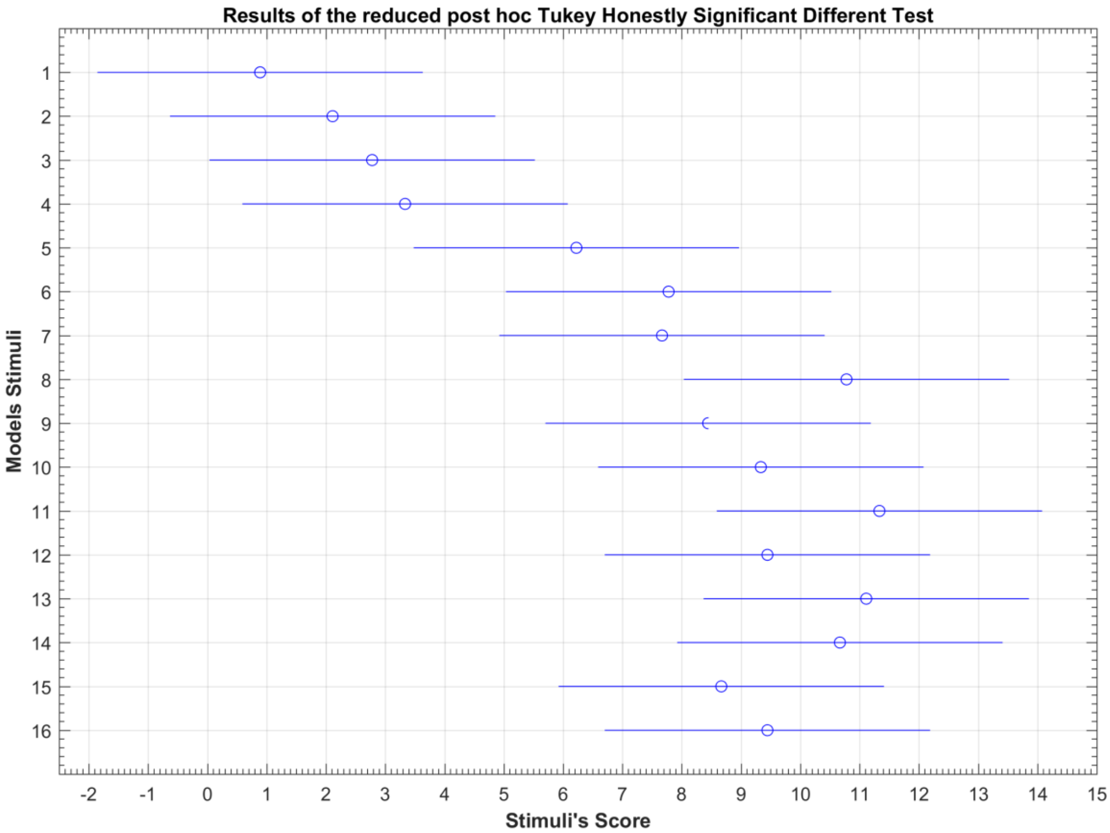

Figure 16: Results of the post hoc Tukey Honestly Significant Different Test on the reduced data

A first glance at the data gives the appearance that the worst perceived 3D model, is the lowest polygonal resolution with no texture (ID 1, Tukey HSD score = 0.8889, P <0.05). The results also seem to concur with the results from Pan et al. [226], where increasing the texture resolution increases perceived quality linearly [226]. Polygonal resolution also appears to plateau, where participants cannot tell the difference between the resolutions. This is supported by the results which where the best perceived model is ID 11, which has a texture resolution of 1024x 1024px and a polygonal resolution of 70% (ID 11, Tukey HSD score = 11.3333, P<0.05).

However, the results are not identical to those produced by Pan et al.[226], as it appears in the data produced for this comparison test, suggests that texture can also plateau. The lowest perceived models are the models with no textures, models ID 1,2,3,4. They have the lowest mean score, with all models being significantly perceived as better than model ID 1 and 2 apart from ID 5 (Tukey HSD, Score = 6.222) which has 10% polygonal resolution and 512x512 px texture resolution. Model 3 is similar, but there is no evidence to suggest that models with the texture resolution of 512 x512 except ID 8 (Tukey HSD, Score =10.3636, P <0.05), are perceived as better. The models with textures have higher mean scores from the Tukey honestly significant test as can be seen in table 10 compared to the non-textured models, yet all of their confidence intervals overlap. There is not enough evidence to suggest that an increase in texture resolution, increases the perceived quality of a mesh. It is only possible to say that a texture improves the perceived quality for this specific mesh. There are no significant differences between texture and polygonal resolutions.

However, the full comparison matrix provides evidence that increasing texture resolution to a point increase perceived quality. The confidence levels are smaller, and show that textures over 512x512 are perceived as better. Model ID 8 (Tukey HSD score 6.8, p<0.05) is significantly better than meshes 1, 2, 3 and is significantly worse than meshes 10, 12, 13, 14, 15, 16. However except mesh 9 (Tukey HSD score = 8.4444, P <0.05), there is no significant difference between meshes at all polygonal resolutions and 1024 x 1024 px and 2048 x 2048 px texture resolution. This suggests that for this 3D model, texture resolution plateaus similar to polygonal resolution, where users cannot tell the difference in texture resolution. Increasing either texture or polygonal resolution after a point, will not increase the perceived quality of the 3D model. However with these results, while there may be an overlap between confidence intervals, there is no evidence to suggest which polygonal and texture resolution is perceived as the best quality or if they are perceived equally. More participants would be needed for testing to increase the accuracy of the results.

The results of the full comparison tests One Way ANOVA is presented in Figure 19, and its post hoc Tukey Significance Difference Test is Presented in figure 20, with the mean score itself in table 10.

ID	Tukey Honestly Significant Different Mean Score (reduced)	Tukey Honestly Significant Different Mean Score (full)
1	0.8889	0.6000
2	2.1111	2.2000

3	2.7778	2.6000
4	3.3333	3.6000
5	6.2222	3.6000
6	7.7778	5.2000
7	7.6667	6.4000
8	10.3636	6.6000
9	8.4444	9.0000
10	9.3333	11.2000
11	11.3333	10.0000
12	9.4444	10.4000
13	11.1111	11.4000
14	10.6667	13.0000
15	8.6667	12.8000
16	9.4444	11.4000

Table 10: Mean Scores from the Tukey Honestly Significant Different Mean Test

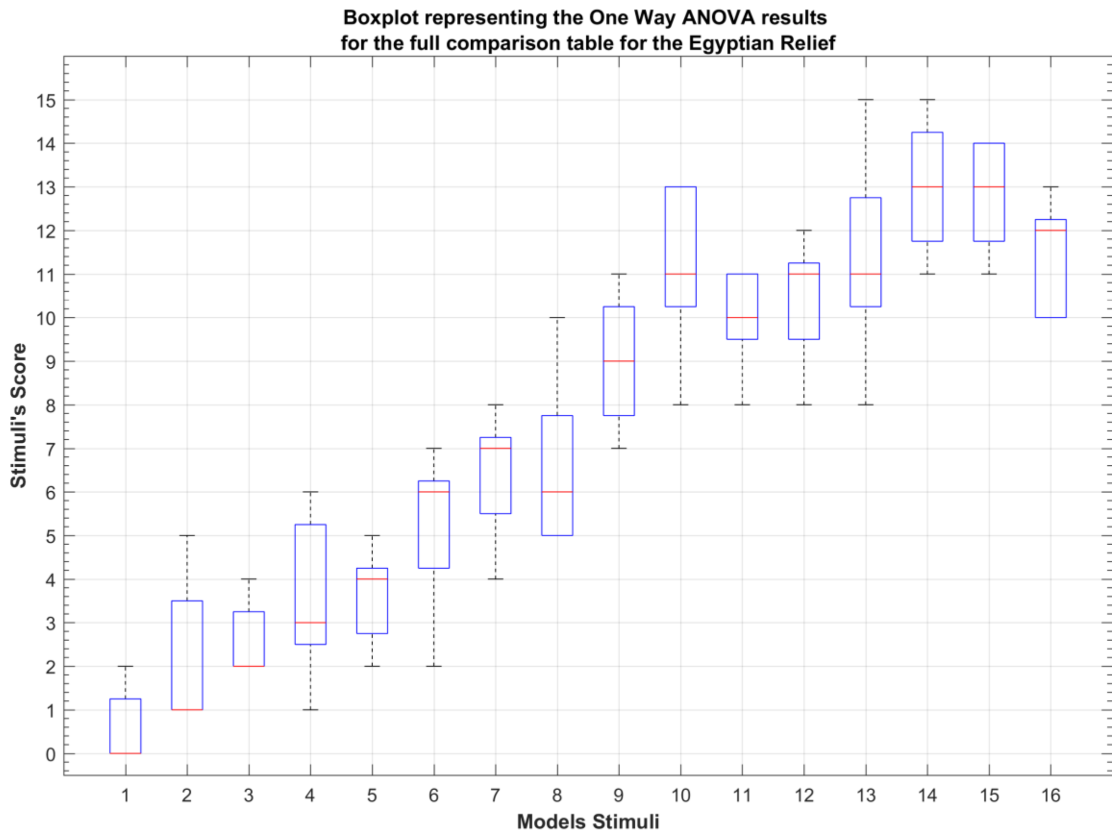

Figure 17: Boxplot representing the One Way ANOVA of the full data

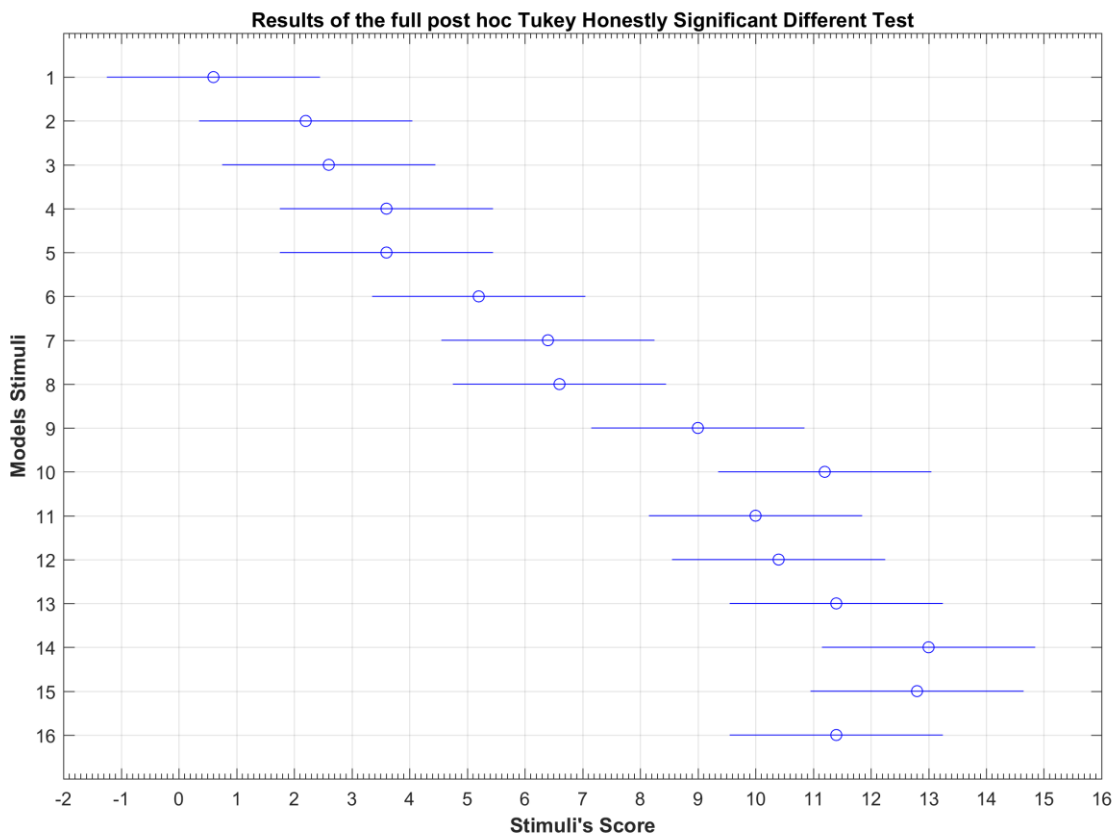

Figure 18: Results of the full post hoc Tukey Honestly Significant Different Test

5.5.3 Zeus Ammon Bust

The Zeus Ammon bust was the first of the "3D" experiences, where the user had to significantly interact with the model to see the full details of the 3D stimulus. 20 participants took part in this experiment. However, 5 participants data was removed after using the ITU-R BT.500-13 screening guide [222]. Unlike the "2D" interactives, there was low agreement among users regarding the perceived quality of the meshes. A Kendall's $W = 0.252$ with $P<0.01$, was computed for the reduced comparison matrix. However, the P-value <0.01 for both the One Way ANOVA and post hoc Tukey HSD, showed there were significant results to calculate the correlation between texture and polygonal resolution However, the low Kendall's W has led to a large number of overlaps between the individual stimuli. The lack of agreement could be due to a number of reasons including; the nature of the self-balancing binary tree, which can cause noisy data, especially when models appear very similar.

However, the Kendall W calculated for the full comparison matrix, has a strong agreement with $W = 0.59$ with $P<0.01$. A One Way ANOVA and post hoc Tukey significant difference test was conducted on both datasets. The results of the One Way ANOVA is represented as a box plot in figure 21 for the reduced comparison matrix, the mean scores and confidence intervals are shown in figure 22, and figure 23 presents the post hoc Tukey HSD test. Figures 24 and 25 present the results of the One Way ANOVA and post hoc test. Table 11 presents the mean scores of each stimulus in the Tukey HSD test.

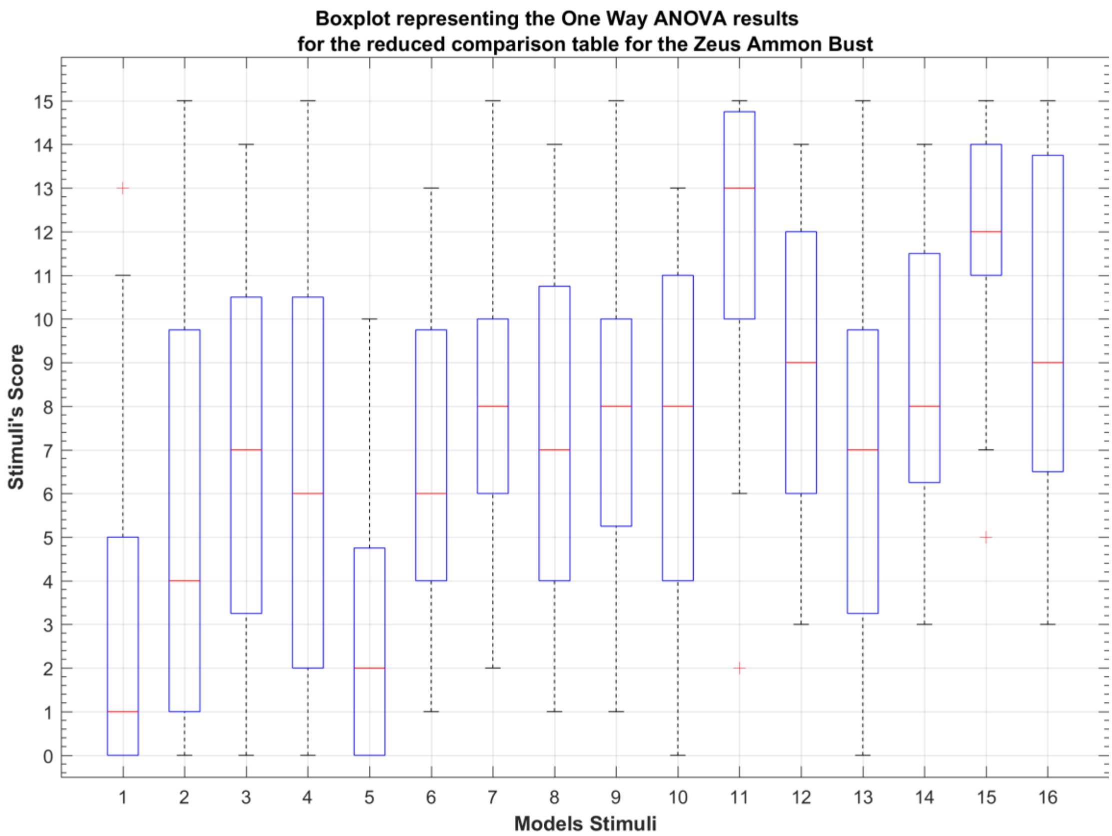

Figure 19: Boxplot representing One Way ANOVA for the reduced results of the Zeus Ammon Bust

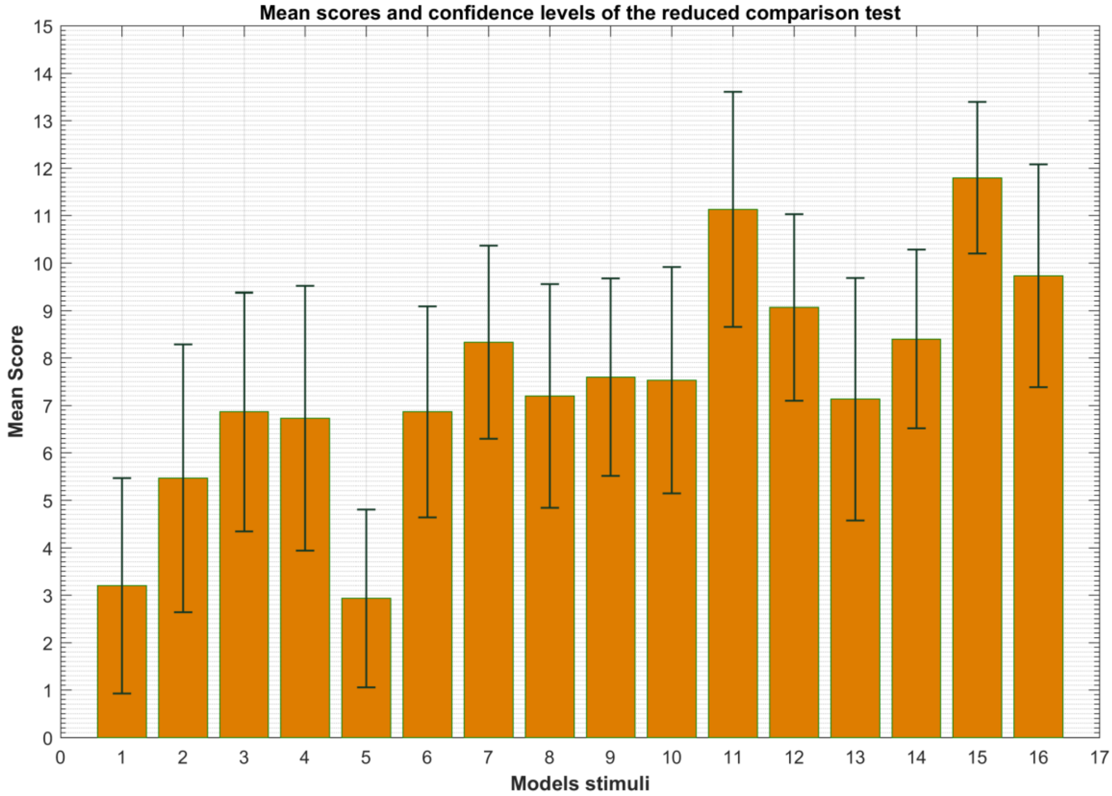

Figure 20: Mean scores and confidence intervals of the reduced comparison test

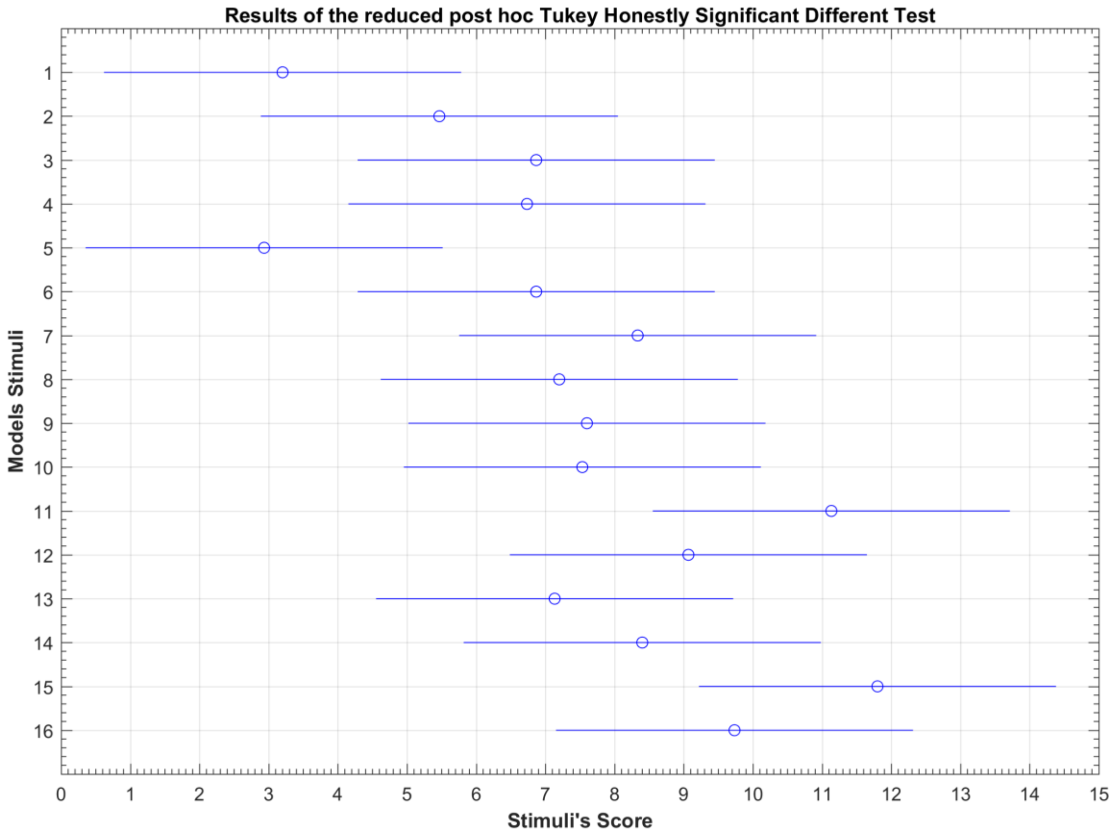

Figure 21: Results of the Post Hoc Tukey HSD for the reduced comparisons

The data that is shown in the reduced One Way ANOVA is very confusing, with a lot of overlaps and no emerging trends or patterns. The results reveal that the stimuli with the worst perceived quality are stimuli 1 and 5 (ID 1 Tukey HSD score = 3.2, P <0.05, ID 5 Tukey HSD score = 2.9333, P <0.05). It also shows the best perceived stimuli are 11 and 15 with 70% resolution and texture resolution of 1024x 1024px and 2048x2048px (ID 11, Tukey HSD score = 11.1333, P <0.05, ID 15 Tukey HSD score = 11.8, P<0.05) with the smallest confidence intervals. The post hoc Tukey HSD does not reveal much information either. It reveals that even though stimuli 11 and 15 are the best perceived stimuli, they are only significantly better than stimuli 1, 2 and 5. They are otherwise are not significantly different from the other meshes, and there is no evidence to suggest they are perceived as the best quality stimuli. However, the worst stimuli is actually stimuli 5 (Tukey HSD score = 2.933, P<0.05), which is significantly worse than stimuli 7, 11, 12, 14, 15, 16. There is not enough evidence to suggest that an increase in texture resolution increases the perceived quality of a stimulus. There is also not enough evidence to suggest that increasing polygonal resolution increases perceived quality either. It is only possible to say that a texture improves the perceived quality for this specific mesh. There are no significant differences between texture and polygonal resolutions apart from at the lowest and highest polygonal and texture resolution.

However, the full comparison matrix does provide evidence that increasing texture resolution to a point increases perceived quality. The results also provide evidence that polygonal resolution plateaus after a certain point, with scores similar between the polygonal resolutions at different levels of texture resolution. The lowest perceived stimuli is the 5th (10% polygonal resolution and 512 x 512px texture resolution) (Tukey HSD score = 1.2, P <0.05), yet it is not significantly better than stimuli 1, 6,7,8,9 (ID 1 10% polygonal resolution texture), (ID 6, 7, 8 - 40, 70 ,100% polygonal resolution and 512x512px texture resolution), (ID 9, 10% polygonal resolution, 1024x1024px texture resolution). The full comparison also reveals that there is no significant difference between 1024x 1024px and 2048x 2048px texture resolution. There are no significant differences, and there is not enough evidence to suggest that 2048x2048px texture resolutions are perceived as better than those of lower texture resolutions. The full resolution also reveals that there are no significant differences between the polygonal resolution and a texture resolution i.e. 13 -14, there is no significant difference between them, and the same for 9 – 12. What was not expected was that stimuli 2, 3, 4 have high Tukey HSD scores, similar to those of stimuli with 1024x1024px texture resolutions but are not significantly different from the other stimuli except 1, 5, 6. Stimuli 3 (Tukey HSD score = 9.6) though is perceived as better than stimuli 7 and 9 as well. This trend suggests that the high resolution polygonal details captured in the mesh are either perceived quality is as good as textures and the best way to display the model. It also suggests that the textures are creating a masking effect on the 3D model, obscuring details the details of the mesh, reducing the perceived quality of the mesh. However, there is also evidence suggesting that increasing the texture resolution increases the perceived quality of the stimuli. However, due to the confidence level overlaps, between the stimuli, it is not possible to suggest what increases perceived quality of the mesh, there is not enough evidence to suggest that users perceive models with texture resolutions greater than 1024x1024px as better quality than those without textures in this experiment. More participants would be needed for testing to increase the accuracy of the results.

ID	Tukey Honestly Significant Different Mean Score (reduced)	Tukey Honestly Significant Different Mean Score (full)

1	3.2000	2.8000
2	5.4667	8.2000
3	6.8667	9.6000
4	6.7333	9.0000
5	2.9333	1.2000
6	6.8667	2.2000
7	8.3333	4.4000
8	7.2000	5.4000
9	7.6000	4.6000
10	7.5333	8.0000
11	11.1333	8.4000
12	9.0667	8.4000
13	7.1333	11.0000
14	8.4000	11.0000
15	11.8000	12.6000
16	9.733	12.8000

Table 11: Mean Scores from the Tukey Honestly Significant Different Mean Test

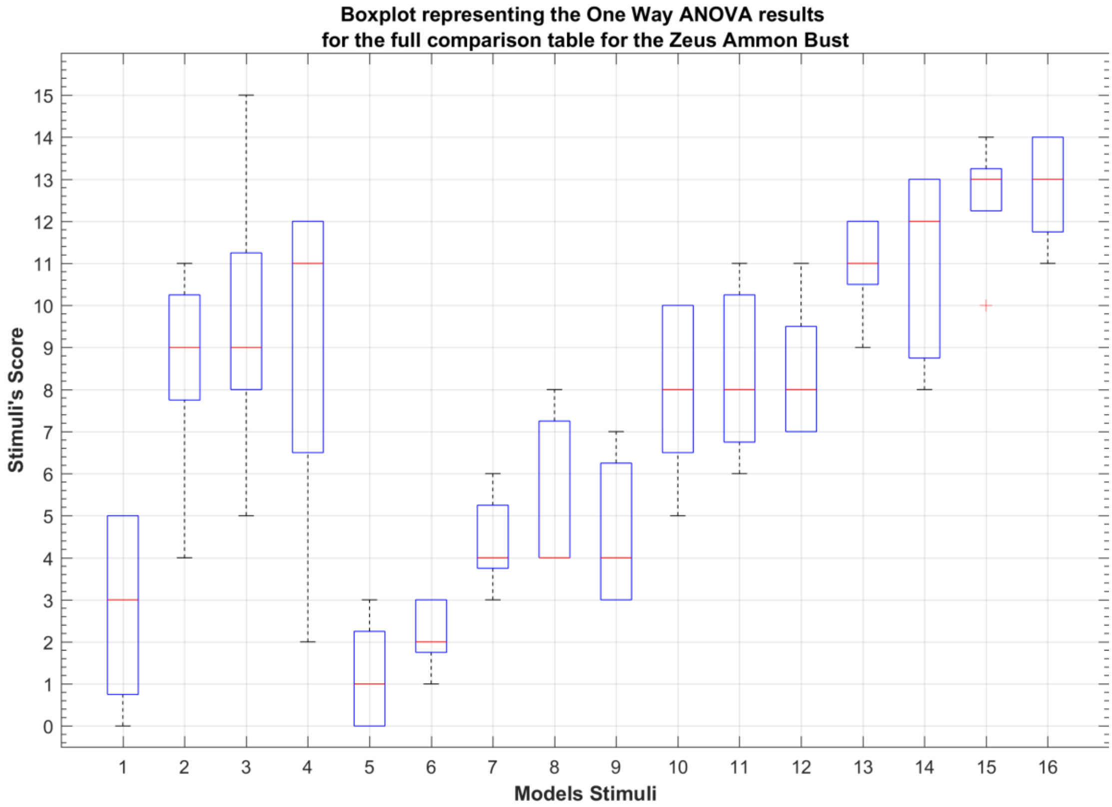

Figure 22: Boxplot representing One Way ANOVA for the full results of the Zeus Ammon Bust

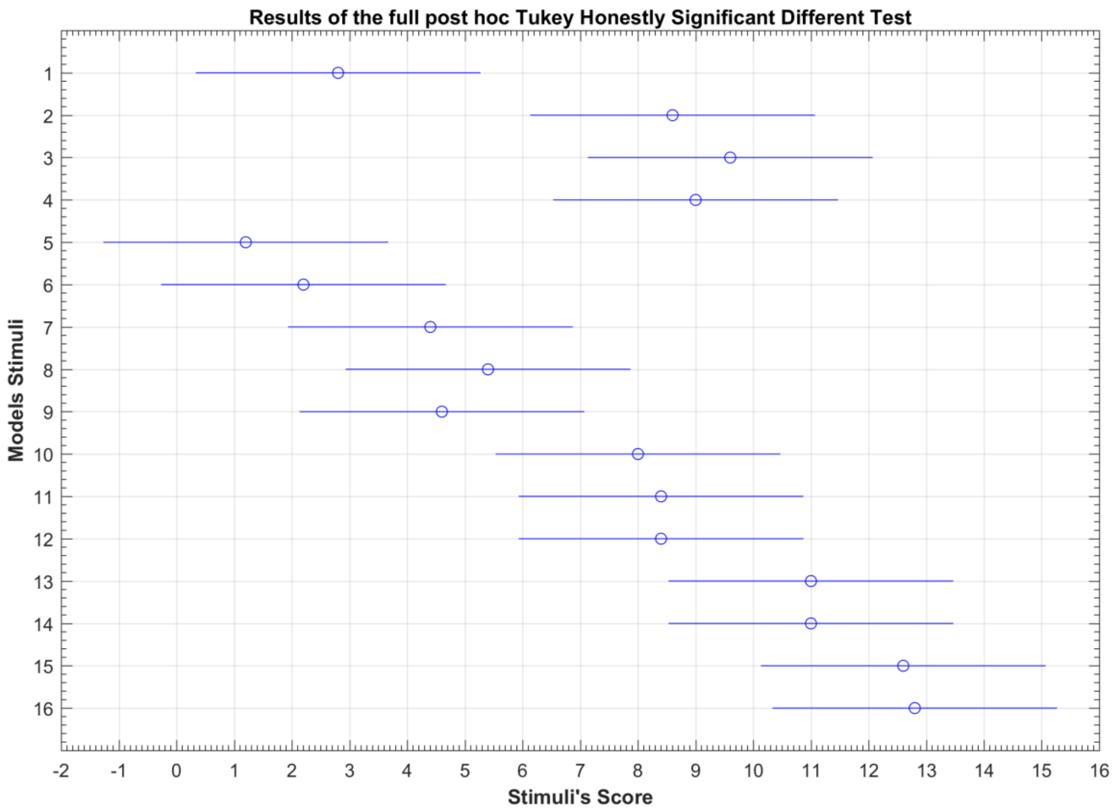

Figure 23: Results of the full post hoc Tukey Honestly Significant Different Test

5.5.4 Shakespeare Bust

The last comparison experiment that was conducted was with the Shakespeare Bust. As with the Zeus Ammon bust, the Shakespeare experiment involved 20 naïve participants, with only 3 participants' data having to be removed following the ITU-R BT.500-13 screening guide [222]. Similar to the Zeus Ammon experiment, there was a low agreement among participants, with a Kendall's *W* = 0.344 with P<0.01 computed for the reduced comparison matrix. However, the P-value <0.01 for both the One Way ANOVA and post hoc Tukey HSD, showed there were significant results to reject the null hypothesis that there is no correlation between texture and polygonal resolution and its effect on perceived quality. However, the low Kendall's W has resulted in large overlaps between data on and this can be seen in figure 26. Once again, the lack of agreement among users could be due to the nature of the self-balancing binary tree, which can cause noisy data, especially when models appear very similar.

However the Kendall *W* calculated for the full comparison matrix, has a strong agreement with *W* = 0.598 with P<0.01. A One Way ANOVA and post hoc Tukey significant difference test was conducted on both datasets. The results of the One Way ANOVA are represented as a box plot in figure 26 for the reduced comparison matrix and 29 for the full comparison matrix. The mean score and confidence intervals can be seen in figure 27. The post hoc Tukey HSD figures are presented in figure 28 for the reduced matrix, and 30 for the full comparison matrix. Table 12 presents the mean scores of each stimulus in the Tukey HSD test.

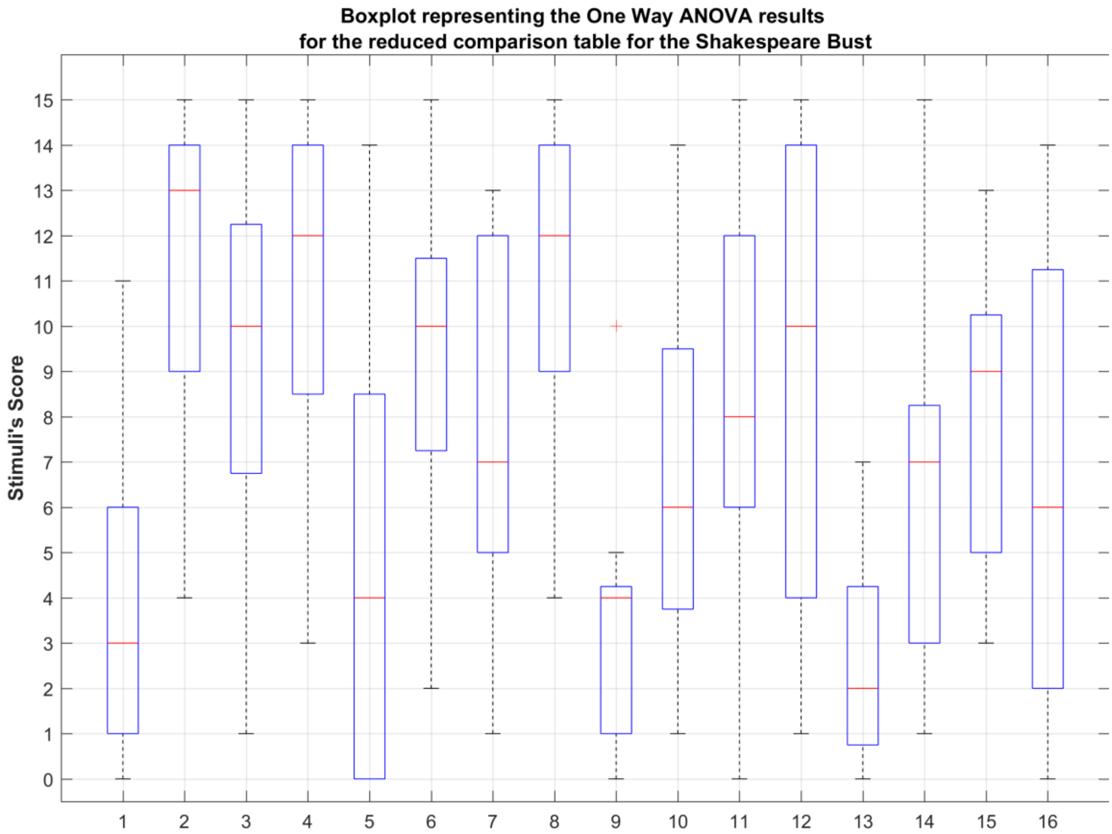

Figure 24: Boxplot representing One Way ANOVA for the reduced results of the Shakespeare Bust

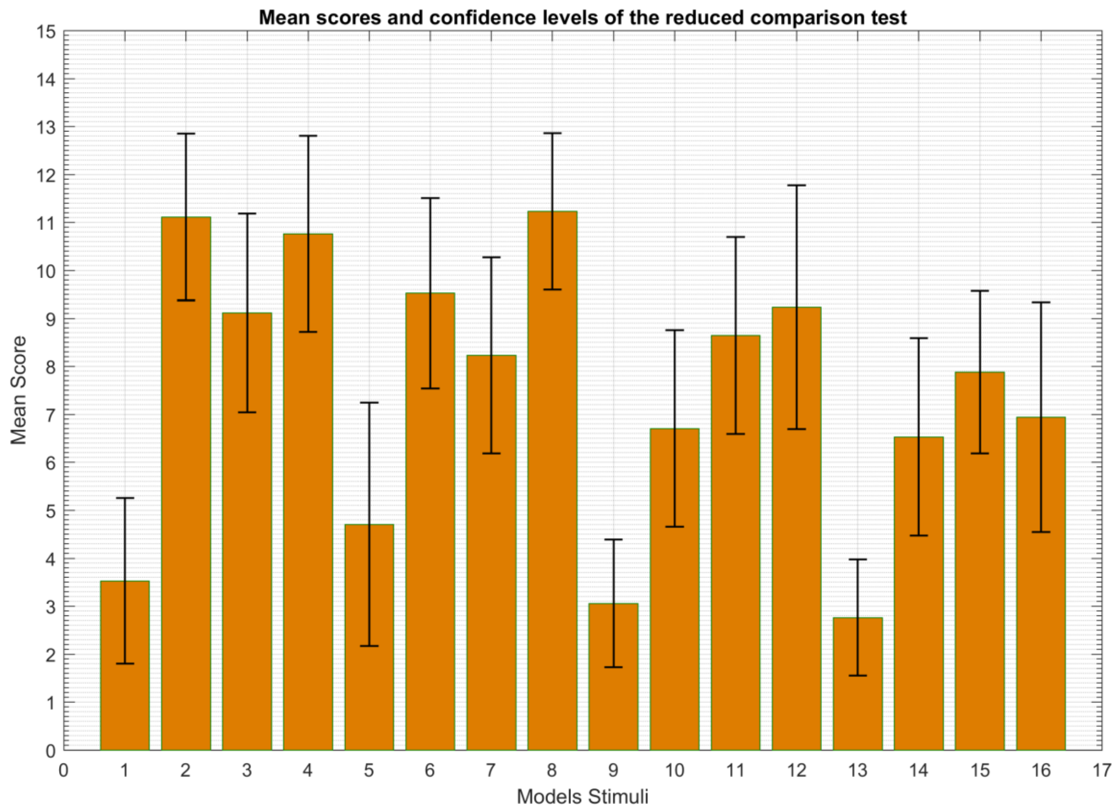

Figure 25: Mean scores and confidence levels for the reduced comparison table

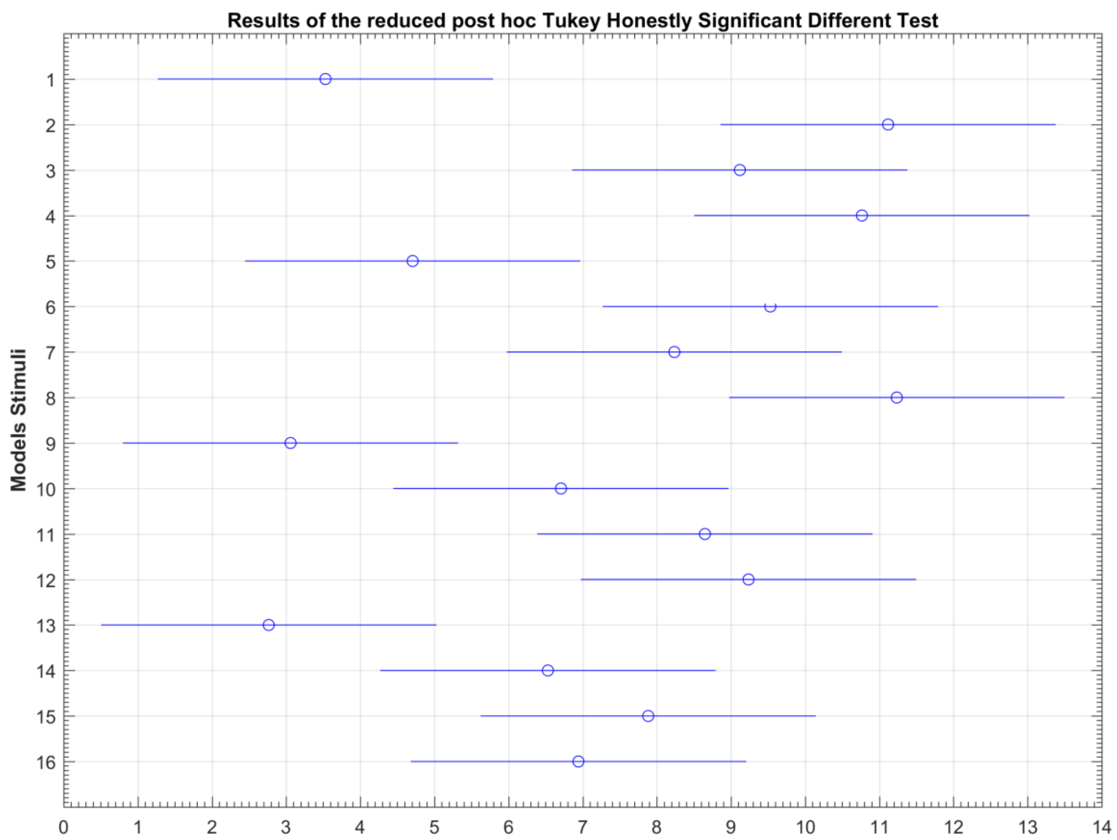

Figure 26: Results of the reduced post hoc Tukey Honestly Significant Different Test

The data for the Shakespeare bust, like the Zeus Ammon bust, contain a lot of overlap among the stimuli shown in the reduced One Way ANOVA, though it is possible to observe a trend. The stimuli that are perceived as the worst are 1, 5, 9, 13 (ID 1 Tukey HSD score = 3.5294, P<0.05, ID 5 Tukey HSD score = 4.7059, P<0.05, ID 9 Tukey HSD score = 3.0588, P<0.05, ID 13 Tukey HSD score = 2.7647, P<0.05), which contain polygon resolution of 10% and they range across all of the texture resolutions. A trend also emerges, where the perceived quality appears linked to the geometry resolution, rather than the texture resolution. This is further supported in the full table comparison One Way ANOVA. The post hoc Tukey HSD, does reveal that there are no significant differences between the stimuli with textures, with a polygonal resolution greater than 10% except for 10 and 14 (ID 10 Tukey HSD score = 6.7059, ID 14 Tukey HSD score= 6.5294) where 10 is perceived as being worse than stimuli 2 and stimuli 14 is significantly worse than 2 and 8. It is also noted that the stimuli without textures and polygonal resolutions greater than 10% have high mean scores, but there is not enough evidence to support that they are perceived as equal or better than meshes with textures. There is not enough evidence to suggest that an increase in polygonal resolution increases the perceived quality of a stimulus. It is only possible to say that meshes greater than 10% of the original mesh are perceived better than the lowest polygonal resolution.

These results are supported by the full comparison matrix, in figure 29 and 30, which are very similar to the reduced comparison, except the highest resolution (ID 15), which was perceived as the best stimuli. The full matrix supports that increasing the polygonal resolution affects the perceived quality of the stimuli. In the post hoc Tukey HSD, there is no significant difference between the stimuli, where polygonal resolution is greater than 10% regardless of the texture resolution. The stimuli with no textures and polygonal resolution greater than 10%, have the highest scores (ID 3 Tukey HSD score = 11.8, P <0.05, ID 4 Tukey HSD score = 11.6, P <0.05), where they are in some cases being perceived as better quality than meshes with textures. However, there is not enough evidence to support that they are perceived as the best representation of the cultural artefact. This does suggest as found in the Zeus Ammon experiment, that it is possible texture is creating a masking effect, obscuring the details on the mesh. However with these results, while there may be an overlap between confidence intervals, there is no evidence to suggest which polygonal and texture resolution is perceived as the best quality or if they are perceived equally. To more accurately ascertain, which is perceived as the best, more participants would be needed for testing to increase the accuracy of the results.

The results of the full comparison tests One Way ANOVA is presented in Figure 29, and its post hoc Tukey Significance Difference Test is Presented in figure 30, with the mean score itself in table 12.

ID	Tukey Honestly Significant Different Mean Score (reduced)	Tukey Honestly Significant Different Mean Score (full)
1	3.5294	1.0000
2	11.1176	8.8000
3	9.1176	11.8000

4	10.7647	11.6000
5	4.7059	1.8000
6	9.5294	6.0000
7	8.2353	8.4000
8	11.2353	9.4000
9	3.0588	1.6000
10	6.7059	7.0000
11	8.6471	9.8000
12	9.2353	9.6000
13	2.7647	3.4000
14	6.5294	8.2000
15	7.8824	10.0000
16	6.9412	11.6000

Table 12: Mean Scores from the Tukey Honestly Significant Different Mean Test

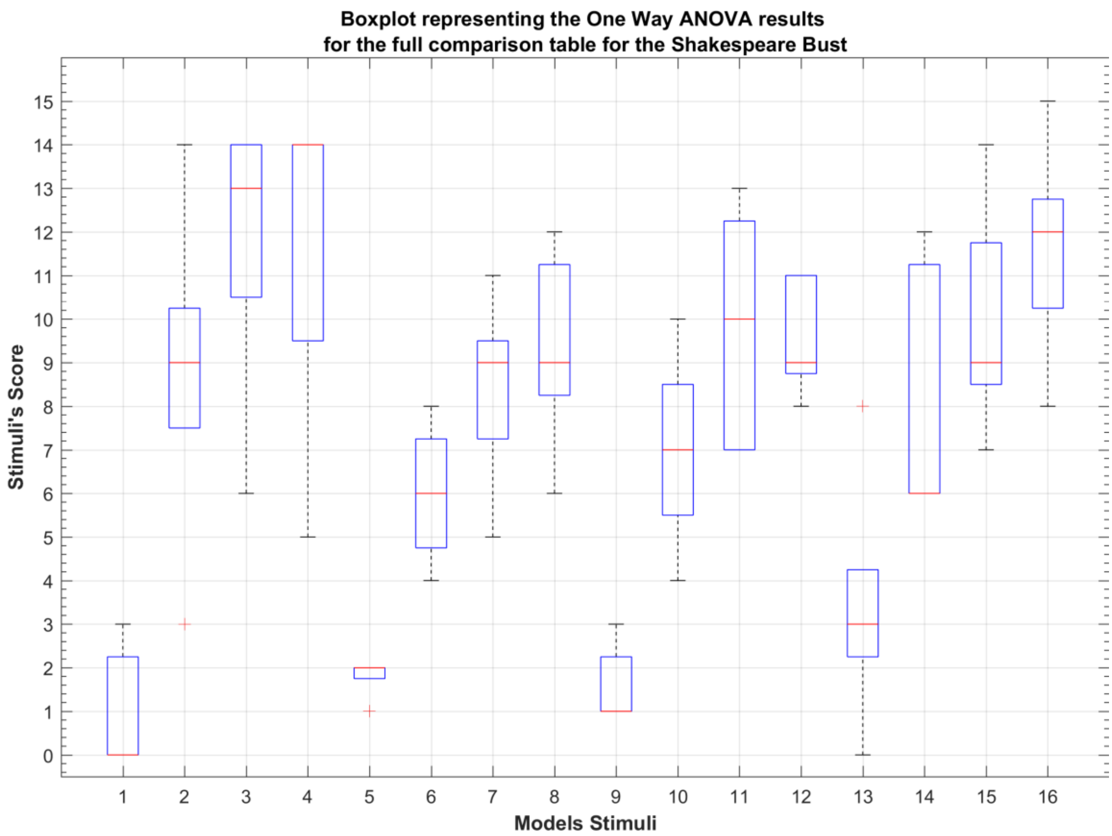

Figure 27: Boxplot representing One Way ANOVA for the full results of the Shakespeare Bust

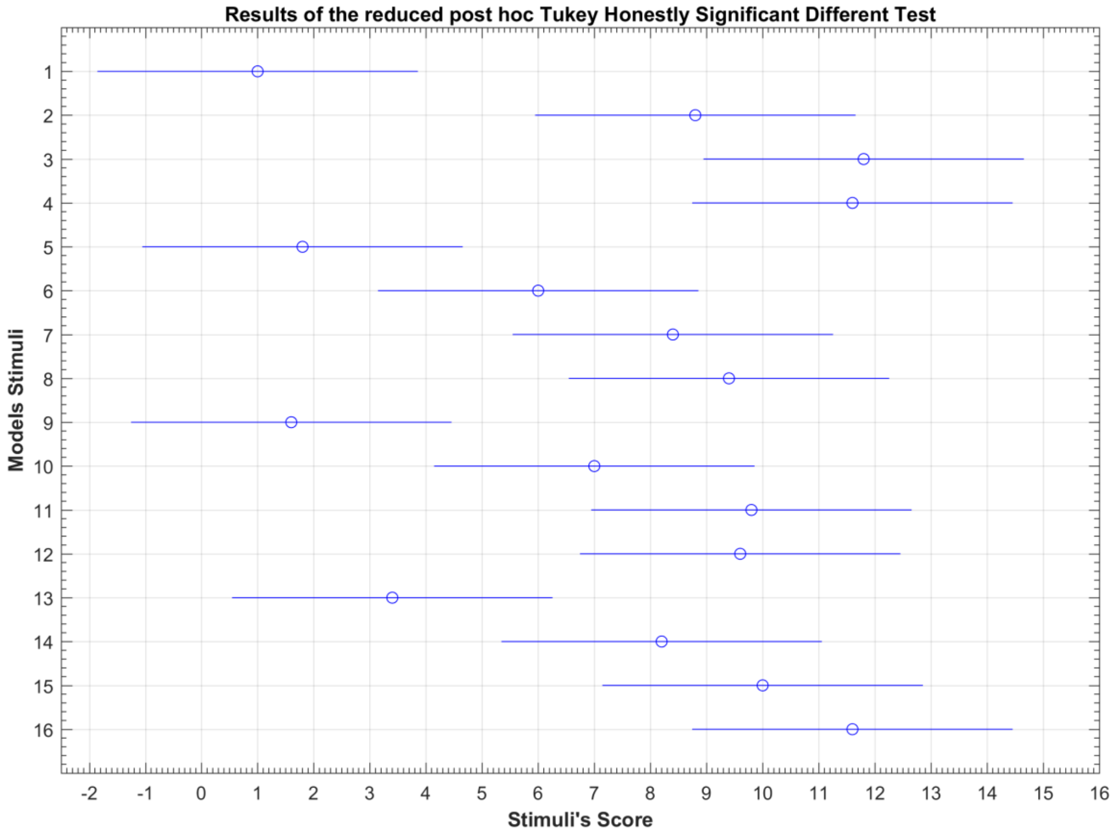

Figure 28: Results of the full post hoc Tukey Honestly Significant Different Test

5.6 Subjective results

Through the paired forced paired comparison test, it is possible to create an ordered ranked ranking of the models and how well they were perceived against each other. However, it is not possible without more information to assess how well they are perceived next to the real artefact. In order to have a more complete evaluation of the 3D digital representations and their perceived quality, a subjective experiment comparing a digital stimulus to the real-world artefact was undertaken. The participant was asked *"How does this 3D model and texture compare to the real object on a scale of 1 to 10? 1 being the worst and 10 being the best."* Other open ended questions were also asked to assess how accurate and well the 3D stimulus was perceived. However as the 3D model is not the only element within the scene presented to the participant, the HDR-VDP2 image metric [249], was used in conjunction with the stimuli generated for the forced comparison test, to create the stimulus for the subjective test. The procedure for creating the stimulus and information is presented earlier in this chapter. The resulting stimulus for each object was 55% of the 100% resolution stimulus from the forced paired comparison experiment, with a 1024x1024px texture.

This subjective experiment was conducted after the forced paired comparison, allowing the participant to interact in full with the object before answering the questions. The full results of the subjective test, for each of the 4 objects can be found in the appendices. A summary of the subjective results will be found below and discussed in more detail in the discussion section of this chapter.

5.6.1 Question 1, How does this 3D model and texture compare to the real object on a scale of 1 to 10? 1 being the worst and 10 being the best.

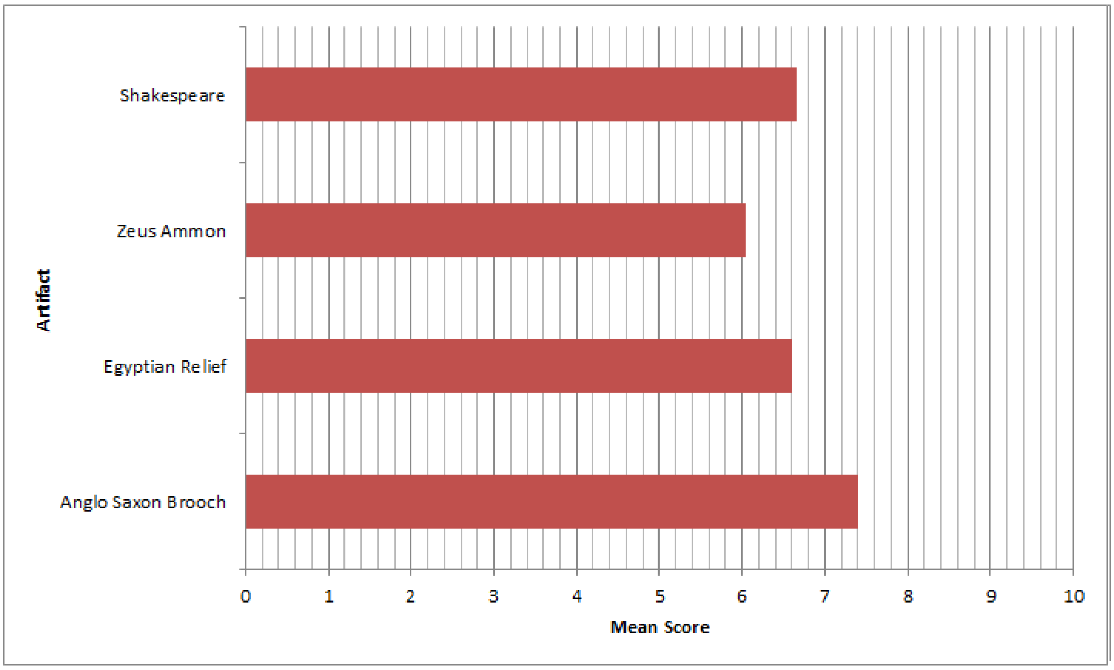

Figure 29: Mean scores of how the 3D digital replica compares against the real life artefact

This is the key question about the quality of the 3D stimulus and how well it is perceived against the artefact. The stimulus was created by using the stimuli used for the forced pairwise comparison experiment, and the HDR-VDP2 [237] image metric. Images taken of the stimuli from within the scene were evaluated using the HDR-VDP2 metric [249], to create a stimulus, which was decimated further than 100% stimuli yet contain very little errors or distortion within the scene according to the HDR-VDP2 image metric [237]. The user was then asked to compare this against the artefact, to evaluate if this stimulus provided a good perceptual experience comparative to the real-world artefact, not the full resolution 3D dataset. The user was asked to rate the stimulus from 1 to 10, on how they perceived the stimulus compared to the artefact. The mean scores and standard deviation can be found in table 13 and figure 31. The stimuli were rated very similarly amongst the four objects, with the Anglo Saxon Brooch, being rated the highest as 7.4 with the lowest being the Zeus Ammon bust with 6.05. The provided stimulus performed very well against the artefact, especially for heavily decimated versions of the original 3D dataset. It should be noted that in some cases, this stimulus for these objects was as low as 8% and as high as 25% of the original resolution of the original 3D dataset. There is also the fact that these objects were decimated using a simple decimation technique (Quadratic Edge Collapse Decimation [242]), which does not focus on preserving details within a mesh. The scores will be discussed in more details with the scores from the forced comparison test, in the discussion section.

Object	Mean Score (between 1:10)	Standard Deviation	Resolution of the original 3D dataset
Anglo Saxon Brooch	7.4000	1.8439	16.7365 %
Egyptian Relief	6.600	0.985611	26.664 %
Zeus Ammon Bust	6.0500	1.700619	8.371%

Shakespeare Bust	6.6500	0.8751	14.289 %

Table 13: Mean Scores and standard deviation for how well the stimulus compared to the real world artefact

5.6.2 Question 2, What do you think this 3D model is made out of?

Participants were asked this question if the stimulus was textured accurately, and if the scene in which it was rendered allowed the user to accurately guess the material of the object. Most participants were able to guess roughly what material objects and its real-world artefact was made from, this is displayed in figure 32. The Anglo Saxon Brooch had a majority of participants guessed that the object was made from either Bronze or Gold. For the relief, most participants generalised their answer to stone, with a few hazarding a guess at plaster or sandstone. The same can be seen in the Shakespeare bust, where participants generalised they are choices to clay or stone though that is in the same vein of materials that these artefacts are made from. For the Zeus Ammon bust, the majority of people generalised their choice to metal, though some were able to deduce that it was meant to represent bronze, or contained copper due to the blue patina of the texture.

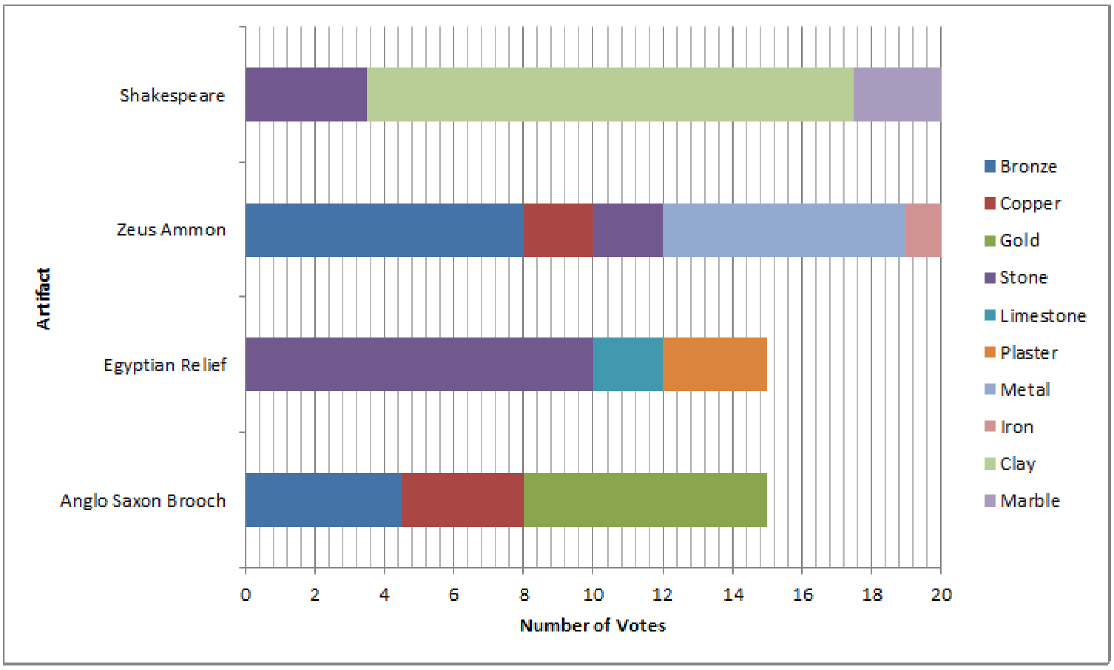

Figure 30: What material each user guessed the 3D replica was made from

5.6.3 Question 3, How important is the texture for you when interacting with this 3D model?

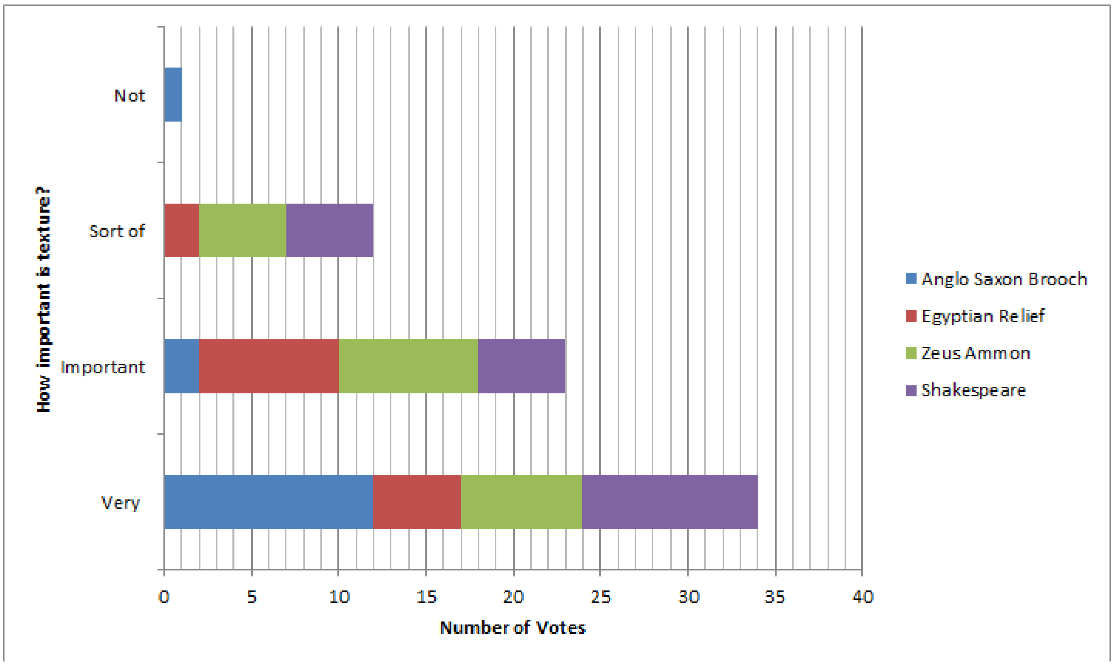

Figure 31: How users answered the question about the importance of texture for interacting with the 3D cultural heritage artefacts.

Another key question regarding the quality of the 3D stimulus is how important is the surface material of the 3D digital artefact, and is it important in the interaction. There was a strong agreement among the 70 participants, across all the objects, that the texture was quite important in the interaction. Apart from two participants, that took part in the Anglo Saxon Brooch, all of the other participants agree the texture was either “sort of”, “very important” or “important”. This can be seen above in figure 33.

5.6.4 Question 4, Would you like the option to choose to display and remove the texture from the 3D model? In conjunction with the above question, it was put to the observer if they would like the option to see the 3D stimulus without the texture like that presented in the forced paired comparison experiment. Would participants, like the ability to see the 3D stimulus, in full detail without a texture obscuring and masking surface details on the 3D model. Similar to the above there was a strong agreement among the 70 participants, agreeing that they would like that option. 7 out of the 70 observers, said they would not like that option or that they did not see it as that important. This is shown below in figure 34.

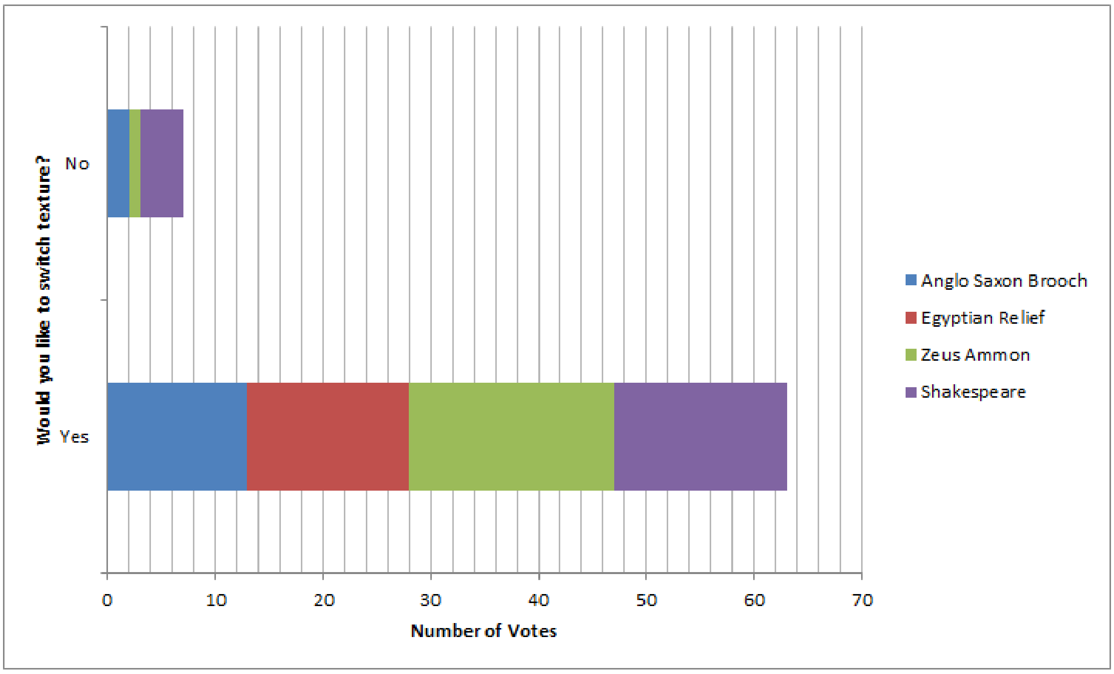

Figure 32: How users answered the question regarding if they would like to change switch between the texture and a non-texture.

5.6.5 Question 5, What would you prefer interacting with: the original/replication or the 3D model?

The 3D models that were created by Conservative Technologies for NML, were created from artefacts that were either too large to handle, extremely delicate and could not be handled by the general public. These 3D models can enhance visitor engagement, allowing visitors to interact with cultural heritage artefacts in new and exciting ways. However, the participants were asked if they would prefer to primarily interact with; the original artefact, a replica (3D printed, bronze cast, CNC milled) or the 3D digital artefact. The answers can be found to question to which they preferred to interact within figure 35, table 14 and with the appendices in section G Subjective Questionnaire Results. The answers were mixed across each of the objects. The answers were mixed across each of the objects. For the Anglo Saxon brooch, 7 out of the 15 participants would rather interact with the original artefact or replica, with 2 purely wishing to interact with the 3D digital replica. The other 6 participants wished to interact with both the original and digital artefact. For the Egyptian Relief, it followed a similar trend, with 7 out of 15 wishing to interact with the original artefact or replica. The other 8 wished to purely engage with the 3D digital replica. Participants for the Zeus Ammon Bust were mixed again, with 12 out of the 20 observers preferring either a mixture of both the original or the digital artefact. This was also the case for the Shakespeare bust, with 11 out of the 20 participants, preferring a mixture of the two.

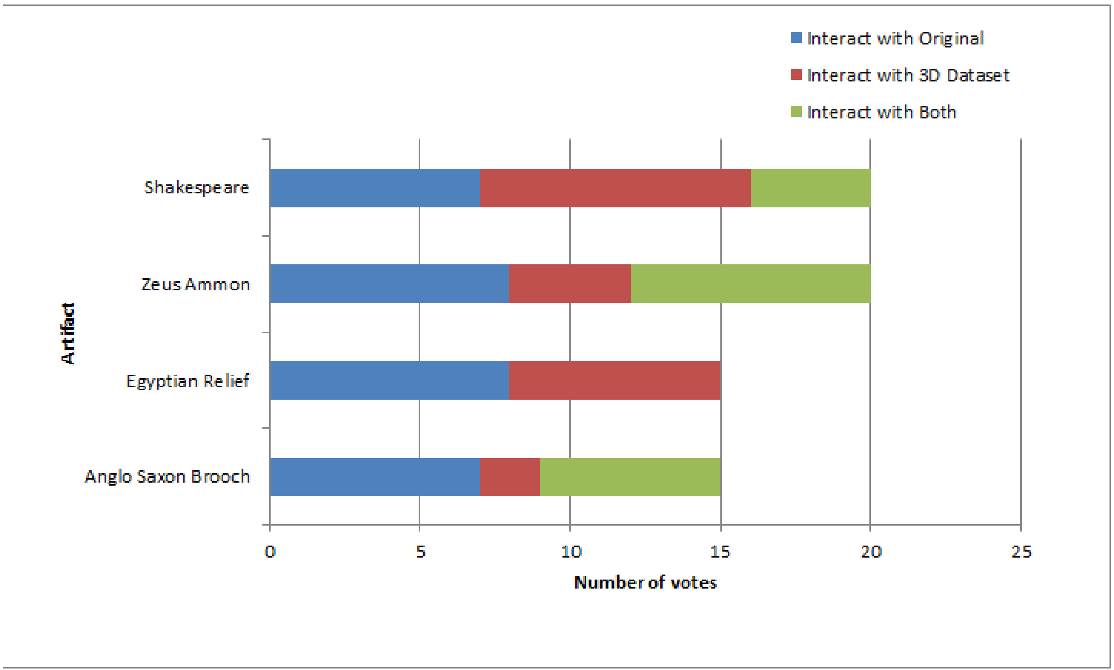

Figure 33: How users answered the question regarding interacting with either the digital replica or the real world artefact

Object	Would rather interact with original/replica?	Would rather interact with digital representation?	Would like to interact with both
Anglo Saxon Brooch	7 out of 15	2 out of 15	6 out of 15
Egyptian Relief	7 out of 15	8 out of 15	0 out of 15
Zeus Ammon Bust	8 out of 20	4 out of 20	8 out of 20
Shakespeare	8 out of 20	9 out of 20	3 out of 20

Table 14: Answers to which they would prefer to interact with

5.6.6 Question 6, After this experiment, would you like to learn more about the collections, or the 3D models that the National Museums have?

This question was asked to garner whether or not participants would be interested in knowing about the collections within the NML and if 3D models would help to generate more interest. How the users chose to answer the question can be seen in figure 36 below. Overall the answers are very mixed across all four objects. The majority of participants for the experiment are involving the Anglo Saxon Brooch and Egyptian relief were interested in learning more, with 55 % of participants in the Zeus Ammon and Shakespeare bust wished to learn more, whereas 87% of users from the Anglo Saxon Brooch and 60% of users for the Egyptian relief wished to learn more about the collections. However, as a whole, 61% of the users in the study chose that they would like to know more about the collections. While the 3D stimuli may have been perceived as good representations of the real-world artefacts, and with participants may wish to interact and look at the original artefact or a replica. The use of 3D objects may be of limited use to generate interest in collections and would be best used alongside additional media. However, this would need to be tested, especially against the likes of 3D printed cultural heritage artefacts.

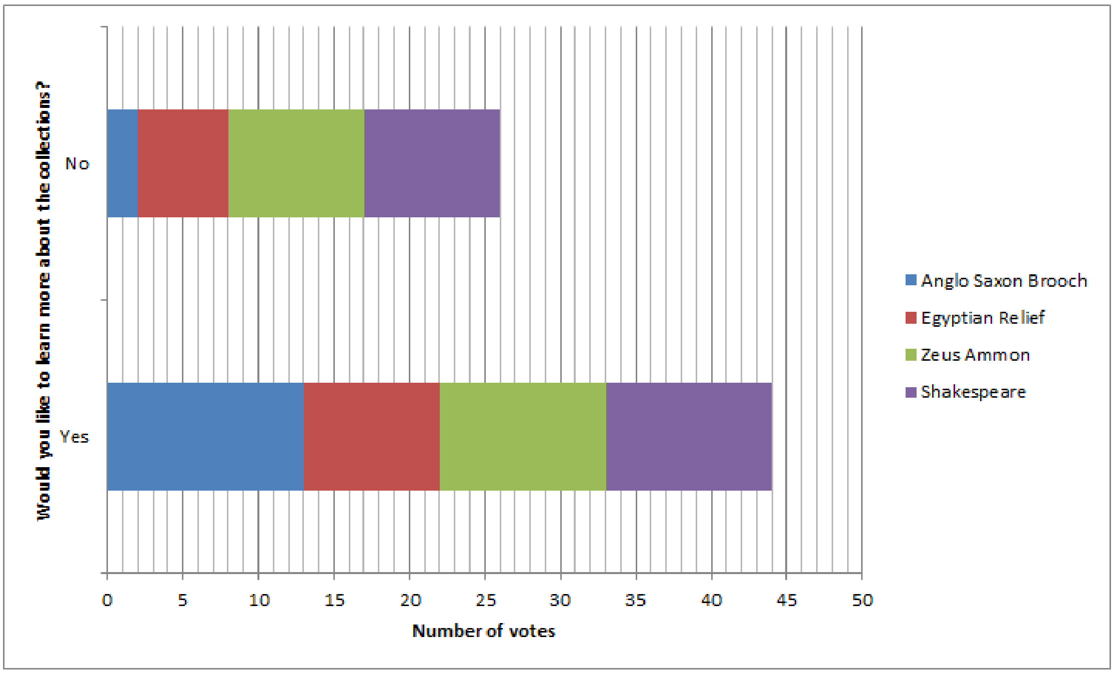

Figure 34: Asked if users would like to know more about the collections after interacting with the digital replica.

5.7 Discussion

This chapter has presented a study, investigating how humans perceive the quality of 3D digital datasets of real-world cultural artefacts through the use of a forced pair wise comparison study and a subjective questionnaire. The study has implications for cultural heritage institutions to help find the acceptable border between polygonal and texture resolution to offer the best perceptual experience. The study asked users to assess the quality of four digital 3D objects created by Conservation Technologies: 2 busts, a relief and a brooch. This study wished to answer the overarching question of this chapter; at what cost to texture and polygonal resolution for a 3D digital cultural heritage artefacts offer the best perceptual effect?

The first experiment aimed to answer this question, explored whether perceived quality is linked with the texture and polygonal resolution of a 3D mesh using differing levels of texture and polygonal resolution. The results of this study supported studies and their claims that texture is important to the perception of quality [221, 226, 243]. However, the results of this chapter did offer interesting new insights, where it was thought that an increase in texture resolution increased quality linearly [226]. This appears to not be the case, although this varies across each of the four objects. This study is also thought to be the first to compare non-textured models versus textured models. Low resolution textures harm the perceived quality of the 3D meshes regardless of the polygonal resolution [224], and no studies using textures have compared against meshes without a texture. This is important and significant, as the results are the first that show that in some cases a non-textured mesh can be perceived as highly as a textured mesh with a high polygon and texture resolution. Stimuli 2, 3, 4, (40 – 100% polygonal resolution, no texture) from the Zeus Ammon and Shakespeare bust both scored highly against the meshes with textures. In the case of the Shakespeare bust, stimuli 2 was the highest scored for perceived quality against the other stimuli. However, due to overlaps of the confidence intervals, there are no significant differences; there is not enough evidence to suggest they are perceived as better as or worse than textured stimuli. However, it does suggest that non-textured meshes can provide a good perceptual experience in some cases, whereas as it has been shown that low resolution textures can harm the perceived quality of the 3D objects regardless of the polygonal resolution [224]. The use of no texture may be a good alternative to display the 3D datasets that offers a "3D interaction". The reason for this trend could be due to a number of details being masked by textures, unsatisfactory texturing for these models, or simply participants preferred the model rendered without a texture within the scene. More research needs to be conducted to explore this further.

The results of the first experiment showed that the worst perceived stimulus was always the most extremely decimated mesh at 10% polygonal resolution. This was supported in the full comparison design table, which rated the stimuli with 10% polygonal resolution consistently as the worst with the exception of the Zeus Ammon which was the 10% stimuli with the 512x512px texture. Participants also tended to rate models with high polygonal resolution as better quality than those with lower. However, this is common across other studies [221, 226, 243]. What makes this study significant, for cultural heritage and other industries, is that there is no significant difference between meshes with polygonal resolutions greater than 40% at any of the texture resolutions in the reduced comparison experiment. There is no evidence to suggest that the texture resolutions differ in perceived quality. However, the full comparison control group produced similar results yet there was no significant difference between meshes with polygonal resolutions greater than 10% with texture resolutions greater than 512x512px. In the reduced comparison, meshes with 40% at points showed no significant differences from the higher resolution meshes except in certain conditions.

This is significant, and very important for both cultural institutions and other industries. As it has been suggested that with geometry alone, after a certain point perceived quality plateaus [226], this study has provided evidence that this is also the case for texture resolution. It also suggests that the trend for increasing both polygonal and texture resolutions for games, films and other industries to increase perceived quality will not have as much of an impact as some believe.

The perceived quality for each model was perceived differently across the four objects. It was observed that the Egyptian Relief, while having an overall high polygon count is a very simple shape, very flat with bold details. The results of its One Way ANOVA and post hoc Tukey HSD results were similar to what is described in Pan et al. study [226], where the texture seems more important in the perceived quality. The score for this object increases linearly with the worst perceived is the 10% polygonal resolution, with the score increasing linearly before it plateaus. The material also applied to the mesh, would also be sensitive to artefacts caused by lowering the resolution as it can be observed more easily. This results in the increased perceived quality by increasing the texture and polygonal resolution. The Zeus Ammon bust also follows a similar trend on the full comparison matrix, where the reduced comparison matrix, shows no difference between any of the texture or polygonal resolutions. This is not the case for the Anglo Saxon Brooch, which seems to be the opposite. While the highest rated meshes are those with the texture resolutions greater than 512x512px, the scores for the meshes are very similar at their own resolutions regardless of texture resolution. It appears as though both the polygonal and texture resolution plateau after the 512x512px texture resolution. However, there is no significant difference between meshes with a polygonal resolution greater than 10% regardless of texture resolution. This is the same for the Shakespeare, where scores apart from the 10% polygonal resolution mesh; they all share very similar scores with no significant differences between themselves. There are very little differences between the polygonal and texture resolution. There is no evidence to suggest an increase in either a texture or polygonal resolution increases the perceived quality of the 3D object.

The results from this first experiment have allowed us to investigate the boundary between texture and polygonal resolution and what offers a good perceptual experience with 3D digital cultural heritage artefacts. The results can be interpreted that by using a method similar to that of Pan et al. [226], to appoint a new lower resolution would be a visually acceptable new resolution to display to the public. However, it is still possible to suggest it can be decimated further to 40% or greater and still be acceptable to the general public. The findings also suggest that a texture resolution of 1024x1024px would be visually appealing without the need to increase the texture resolution.

It should be noted at this point, that previous studies involving textures by Pan et al. [226], Guo [221] and to a certain degree Thorn et al. [243], have attempted to deduce an image metric that can predict perceived quality. This study has not. The results from this study alone show how difficult it is how to predict quality, where results for objects from the same group differ greatly. However, this would be useful for future work focusing on the creation of an image metric that both utilised 3D image metrics and an image metric that evaluates an entire scene such as the HDR-VDP2 [237].

The second experiment focused on analysing how users reacted to a created stimulus via the HDR-VDP2 image metric, and their evaluation of this versus the real-world artefact. Other questions were also introduced to the participant, to assess the importance of the material and if the use of 3D objects would generate further interest in the collections of the NML. The stimuli that were created for this experiment were 55% of the 100% stimuli with a 1024x1024px texture from the paired comparison test. Participants were asked to rate the stimuli rendered within the scene against the artefact presented to them from 1 to 10. The results showed that the stimuli performed well and were a good representation of the digital artefact, but not a perfect replica. There is still room for improvement and it should be noted that these are highly decimated versions of the original 3D dataset. Table 14 shows the actual percentage of the polygonal resolution against the original 3D dataset; the stimuli are for the most case roughly 15% of the original datasets resolution.

The results from the questionnaire were also similar with other studies for the use of textures and their importance in the perceived quality of 3D objects [118, 221, 226, 243]. There was a near unanimous agreement between participants for each object that texture was important for the display of 3D digital cultural artefacts. Participants also thought there was a need to have the ability to change between a textured and non-textured state [118].

The results also seem to support the theory, that there needs to be additional media alongside the 3D object to generate interest in the object itself, and the other collections within a cultural institution [118]. This is reflected in the results of the questionnaire. When asked if they would prefer to interact with the 3D object or original artefact/replica and if the 3D model would help create interest for the participant in the collections, the results were mixed. Participants for a majority would prefer to interact with the original artefact/replica or the digital artefact with the real-world artefact. This may be so that they have a context regarding the 3D digital replica, improving their experience from being novel, to something deeper [118]. They may also wish for more information surrounding the artefact, to help facilitate generating interest in the 3D model and in the collections.

The results of this study have been already been applied within the NML for the display of 3D digital artefacts. Using the evidence provided by this study, the Egyptian relief and Zeus Ammon are being displayed on the NML's website in conjunction with the opening of a new Egyptian Gallery at the World museum [256]. The 3D models are being displayed and rendered with WebGL, part of the reasons for this choice of platform are as a direct impact of the research outcomes achieved in this engineering doctorate [256]. The 3D models have been uploaded to SketchFab [95] and embedded within the NML website. The 3D models themselves are the 70% resolution stimuli from the paired comparison test, with a 1k texture applied to the Egyptian Relief. The Zeus Ammon bust has no texture applied after curators and the web team perceived it to be a realistic representation of the real-world artefact even without one. The parties involved were told about the decimated models and that they were not the full resolution, and were pleased with this. It allowed the museum to share and disseminate these 3D models freely, offering a good perceptual experience to visitors without the various issues of copyright, or IP getting involved. They believe the 3D models used were decimated and far removed enough from the original dataset that they would not need to worry about intellectual theft. This is also allowing NML, to test the water for sharing their 3D models online, and if they are popular enough to create a community around them.

5.8 Conclusion

The aim of this chapter was to offer the best perceptual experience to visitors when they interacted with 3D cultural heritage; investigating the relationship between texture and polygonal resolution and how it impacts on the perceptual experience. The results from the comparative and subjective experiments allowed for interesting conclusions to be drawn regarding the perceived quality of 3D cultural heritage artefacts.

The results from the comparative study revealed that for each of the objects there was no significant difference between meshes with polygonal resolutions greater than 40% regardless of texture resolution in the reduced comparative study. Similar results were found in the full comparison results which revealed there was no difference between meshes with polygonal resolutions greater than 10% and texture resolution greater than 512x512px. This would suggest that both polygonal [221, 226, 243], and texture resolutions plateau. The trend of increasing texture and polygonal resolution may only increase perceived quality slightly. However, with these results, there was an overlap between confidence intervals. This reveals there is no evidence to significantly decide which one is perceived as being of better quality. However, this does not mean that both produce equally good quality, just that there is no statistical difference in quality and more testing is needed.

The results also revealed that the two 3D interactive style 3D datasets when presented without a texture were perceived quite highly. They showed similar scores to those of stimuli with a high polygonal and texture resolution and in some cases perceived as the best way to display the cultural artefact.

These results are very significant for cultural institutions. It allows institutions to offer a good perceptual experience to visitors, by reducing the polygon count to 40% if following the methodology presented on page 93. The texture resolution does not need to be overly large and in some cases may not need a texture at all.

The second experiment aimed at quantifying how users would react to a digital replica and how it compared to the artefact/replica. The results showed that the stimuli performed well against the object. However, when participants were asked about the preferred engagement with the cultural artefact, participants favoured either the real-world object or a preference for both. This could possibly be due to the level of immersion that a monoscopic display can offer. This has been studied buy Thorn et al. [243], which suggested that the use of a VR system, helped to increase the immersion of visitors. There is also research by Di Franco et al. [118], which suggests that participants are willing to engage with replicas of the original artefact for a tactile experience and prefer this over a digital experience.

Chapter 6 Conclusions and Perspectives

6.1 Introduction

The research carried out in this thesis was motivated by the potential of the 3D laser scanned digital artefacts. It enabled visitors and specialists to engage, view and interact with cultural artefacts in new and exciting ways. Since the inception of this thesis, new avenues for dissemination and visualisation of artefacts have arisen such as WebGL (addressed in chapter 3), augmented and virtual reality and 3D printing. These allow for greater access to the 3D cultural heritage artefacts, offering either remote access or surrogate replicas of the original cultural artefact. However, as dissemination and visualisation of artefacts is a broad field for research, the direction of the research focused on the main issues for dissemination and visualising 3D cultural artefacts; IP issues, interacting with 3D objects via new mediums and creating a good perceptual experience.

6.2 Conclusions and contribution to knowledge

The first part of this thesis focused on IP issues surrounding 3D cultural heritage artefacts (see chapter 2). To address this issue, a literature review of the IP status of laser scanned cultural heritage artefacts, addressing case law, and government studies produced by the British Government were undertaken. Chapter 2 helped the NML to understand the rights and their position in regards to their 3D objects that they hold within their archives. The chapter quickly concluded that, in regards to IP it was most likely to be unaffected by Patent [26], Trademark [27], Design and Copyright [37] law, where copyright would have the largest impact. The complex nature of copyright, does not allow us allow us to draw a conclusive answer to questions raised in chapter 2 such as can an institution hold the copyright for a scanned work of a public domain object. It is complicated by wither the dataset was created from an artefact that is utilitarian in purpose or if it was to be enjoyed visually. Evidence provided in this chapter supports and contradicts but ultimately it needs to be settled in the courts to establish the copyright rights of scanned cultural artefacts. However under current law, it should be possible to acquire copyright on the digital file created from the scanning process, and derivative works would also be afflicted with this copyright. This would include the use of the digital file to 3D print objects or visualise the object. The chapter delivered a notable contribution of knowledge in the area of IP for the NML. It also provided additional information for dissemination solutions and to allow them to monetise their digital assets. It also addressed the issues and risks involved in sharing their assets, highlighting the benefits of sharing the digital artefacts. However, as the research was conducted without a legal expert, the research and conclusion can be only used for guidance and not provide actual legal advice.

The next area of focus for research was the dissemination, visualisation and interacting with 3D cultural heritage artefacts. As the use of the internet is becoming more prominent and ubiquitous in our lives [93, 94], the research focused on the display of 3D cultural artefacts within web pages. The research conducted in chapter 3 focused on the creation of a prototype viewer, to display 3D artefacts and an HCI study to investigate how users would interact and use the viewer itself. The study concluded that the preferred navigation and interaction style was very similar to styles already used within 3D applications: rotating the object around its center point and not inverting the axes. Users preferred the simple interaction style but wished for the additional tools such as a raking light to allow for more in-depth engagement for the user. Due to the instability of WebGL at the time NML chose to incorporate the research from this study into one of their projects, the Pre-Hispanic Caribbean Sculpture in Wood [139].

With the growth of 3D content and its use within various mediums, the look and feel are becoming more important in the display of 3D content [154]. However, for cultural heritage, there are many objects that have been digitally recreated without surface material information. Chapter 4 addressed these issues in the form of a literature review. It identified that one of the possible options is the use of surface parameterisation, which is widely adopted within CGI and is the industry standard. Its popular nature and wide implementation would allow for its implementation within cultural heritage. It can also be used in conjunction with 2D image synthesis software such as Substance Painter [240]. Solid texturing also provides a way to texture a 3D model without a UV map, but it requires a large amount of computational power and storage for the smallest of textures. It has improved upon since its first conception but the growth in this area is still lacking, yet ripe for potential future research.

The literature review also provided knowledge for assessing the quality of the 3D content, and how it may be perceived by the general public. This is relevant as the displayed 3D content must be an acceptable representation of the original artefact in accordance with the London Charter [13], and offer a good perceptual experience for visitors. There are many 3D metrics available to assess the mesh, but they are narrow in focus. Parameters such as shading, lighting, texture and others can have a major impact on the perceived quality by observers but are not considered by automatic image metrics. The perceived quality may be assessed using human observers; however, the image metrics need to be expanded to incorporate these parameters for the evaluation of the 3D model. Thus an evolution of the 3D metric needs to incorporate both a metric that evaluates the mesh itself for quality, but then considers the scene, the lighting and materials and environment in which the mesh is being rendered via a 2D image metric such as the HDR-VDP2 [238].

The realisations brought about by this literature review led to the research focus shifting. It moved towards investigating the balance between acceptable visual quality of a 3D digital cultural artefact and how the relationship between polygonal and texture resolution affects this. Chapter 5 addressed this with a large study involving visitors from one of NML venues, where research was conducted to investigate this relationship through a comparative and subjective study. Visitors were asked to assess the quality of the created stimuli for the experiments and compare them against the artefacts. The study also compared non-textured datasets against textured datasets for the first time. The results from this study were similar to other studies [221, 226, 243], which reported that texture was important in perceived quality. Yet, the results themselves differed across each object, where objects from the same class of interaction and group differed greatly. The results would suggest that after a point texture resolution no longer increases the perceived quality of a mesh, similar to polygonal resolution [221, 226, 243]. Interestingly, the two objects that offered 3D interactions, a non-textured model was scored as highly as the highest texture and polygonal resolutions. These results are significant as it reveals that cultural heritage datasets can be heavily decimated, and potentially do not need to be textured. The cost to the polygonal resolution can be up to a cost of 40% of the polygonal resolution following the Pan et al methodology [226] and texture resolution can be 1024x1024px and above to offer a good perceptual experience.

The second experiment of this study focused on quantifying the actual perceived quality of a digital replica versus the real-world artefact. The stimuli were created via the use of the HDR-VDP2 [238] image metric, which assessed the decimated mesh within the scene to evaluate and create a stimulus that offered the best perceived quality. The results of the experiment showed that the created stimulus performed well against the presented artefact, but was not a perfect replica eliciting the same response as the artefact. There was further evidence provided via the subjective experiment that the 3D digital replica could not function as a surrogate replacement for the original artefact. Participants' answers to questions regarding their preferred interaction pointed towards having the artefact/replica present in some capacity. This is also present in a Study by Di Franco [118], where participants preferred a tactile experience over a digital experience, suggesting that a monoscopic experience with 3D digital artefacts can only immerse a participant to a certain degree. Though this may not be the case for mediums such as virtual reality or 3D stereoscopic experiences that increase immersion when engaging with objects [243].

This body of work documents the research and investigations into the possibility of creating a user-appropriate viewer for high resolution interactive engagement with 3D digital cultural heritage artefacts. This thesis has contributed:

- In depth knowledge of how IP interacts with laser scanned cultural heritage
- Provided knowledge regarding the impact and resolutions for sharing and disseminating 3D cultural heritage artefacts
- A methodology for the creation of an interactive viewer
- Results from a small HCI study regarding preferred interaction with 3D content within a webpage
- In depth knowledge of surface parameterisation, solid texturing and quality assessment 3D metrics via a literature review
- Results that indicate an acceptable cost to texture and polygonal resolution of a 3D digital cultural artefact to offer the best perceptual experience.
- The first study that includes compares non textured meshes versus textured meshes. The study provides results demonstrating that a non-textured mesh may be as highly received as a high resolution textured mesh.
- Results indicating the effectiveness of a texture when compared to the high level detail that is captured via laser scanning.

This body of work has asked and answered many important questions that may affect the interaction with high resolution digital 3D cultural heritage artefacts. To spread this knowledge and allow other institutions to make use of this information; the research within will be published in the form of papers in journals and at conferences and will be made publicly available to other institutions via networking with NML.

6.3 Limitations

Due to the research being completed for NML, there are requirements that had to be met that influenced the direction of the research. The research needed to be able to accommodate the needs of the museum, while considering the growing uncertainties over possible future budget cuts. The thesis considers three key areas that feed into each other and present a feasible option for displaying and interacting with 3D cultural heritage.

The projects that were conducted as part of this thesis are limited to NML; limited to their resources and interests. The limitations of their resources has resulted in the research been conducted to run on their hardware and integrate with their practices, which will not change for the foreseeable future, introducing limitations on what can be displayed, and limited the testing of new technology such as the Unreal Engine 4 [257] to display 3D datasets, using tools such as substance painter [240] to easily texture 3D datasets realistically or exploring the use of VR or AR within NML.

The projects were conducted within NML venues, in real-world environments interacting with museum visitors; they are subject to change outside the influence of the researcher. The experiment participants as explained on page 103, did not truly reflect the diverse age range of NML visitors and 74% of the 70 participants were aged between 18-33 years old, limiting and possibly biasing the results of this thesis. More research would need to be conducted with more users within NML to discover if this is a limitation or a true reflection of how users wish to interact with 3D cultural heritage.

In hindsight, certain actions or aspects of these projects may or may not have introduced limitations to the results presented in this thesis. However, the resource limitation for NML and myself as the sole developer and researcher working with the 3D cultural artefacts limited what actually could be done in practice.

6.4 Further Work

There are a number of areas highlighted by the work in this thesis that would benefit from further exploration and provide positive directions for future work.

- There is a need to further explore IP rights of scanned cultural heritage artefacts. This is an important area of research that is likely to become very important in the coming years especially with the emergence of 3D printing. The research in this thesis was not able to provide a conclusive answer to the questions raised in chapter 2. It is still ambiguous as to the rights of the 3D datasets and their digital file that have been created from objects within a museum, and if they are entitled to IP protection. This needs to be further explored with 3D printing becoming more mainstream and cultural institutions expressing interest in its application [118]. Would these same rights if they are on the digital file still hold for the 3D printed copy? Would the rights be affected if you have a copyright on a sculpture or artistic work, but when 3D printed is a near exact copy of a public domain work? What of design or trademark rights that may also come into play? There is currently not a lot of research being conducted into the use of 3D printing within cultural heritage but its use is likely to increase in the coming years. There is also another need to continue this research, due in part to the decision of Britain to leave the European Union. Current IP laws are likely to change, and it is unknown if the laws and decisions that have been passed in recent years will still hold true with regards to cultural heritage.

- Another important area is the intangible nature of cultural heritage artefacts with regards to IP laws. This is currently a grey area especially with regards to artefacts that may be a part of traditional expressions of a culture or artefacts that may have been appropriated without a cultures permission, and wish for their return. The rights of these cultures need to be considered when 3D artefacts are created from their appropriated artefacts.

- There is also a further need to investigate how IP may impact upon other methods of creating 3D cultural datasets, especially Photogrammetry and Structure from Motion (SfM) techniques. These technologies are being adopted rapidly within cultural heritage [258, 259, 260], as they are capable of creating 3D datasets with high resolution textures and accurate surface details rivaling non-contact laser scanning [260]. The 3D datasets are created via photos instead of the use of a laser scanner [259, 260]; therefore add a further complication to the intellectual status of a created 3D dataset. Artistry is recognized in the taking of photographs [17, 45], yet it is unknown if this artistry is applicable to the creation of a 3D dataset from images. If the artistry was recognised it could lead to confusion about the copyright status of datasets created from the images. This needs to be review especially in the light of the rapid adoption of SfM for the creation of 3D cultural artefacts [259, 260].

- There is considerable scope to expand the results of chapter 5. The objects used in chapter 5, offered different shapes of flat objects and of busts, that were small and easily viewable within the viewer. To improve upon the results of this chapter, objects such as sculptures where details are not focused in one particular area may be used instead. There is also the potential to use objects where it is hard to discern the main focus point.

- Decimation techniques may be able to benefit from the work provided in this thesis, providing data that can be used in future work to drive perception based mesh decimation research. Potential research would be able to generate a method that can decimate a mesh, preserving areas of interest, and replicating a good perceptual replica of the high resolution dataset. There is also the potential of using the data provided in this thesis to drive a deep learning algorithm that would provide a metric to rate a decimated mesh. This could be accomplished by either training the dataset on 2D images, or the 3D meshes. This would be extremely useful when identifying decimated textures and meshes that could provide a good perceptual experience for users, both within cultural heritage and other industries. The deep learning algorithm could also be used to create a generative/adversarial network to drive a perceptual based decimation technique by being trained on the meshes generated for the studies of chapter 5. This would be as far as I know, never been attempted in cultural heritage or other industries.

- An area of development is the growing use of 3D printing (see Chapter 5). There is a growing interest in the use of 3D printing by cultural heritage institutions [118], to offer a tactile experience with cultural heritage artefacts. There is potential research focusing on the use of 3D printing to compare various resolutions of 3D printed artefacts, and to study how they compare against each other. This would help to understand whether the results from chapter 5 are transferrable to other mediums such as 3D printing. There is also the possibility of exploring the greater role that material may play in the interaction with3D printed objects that have been created from cultural heritage artefacts. Does the material that the 3D object has been printed in effect the overall interaction? Would the use of a plastic polymer be a good material for interaction or would it cheapen the over interaction? Does the cost of the 3D printed object also come into the thought process of users as they interact with the 3D printed object? This research would be able to push the boundaries of interacting within cultural heritage, providing information on if a 3D printed object would be able to replace interacting with the original. This research could be very beneficial for cultural institutions, as it would allow them to be able to print exhibits for galleries via digital files instead of transporting the object from one institution to another.

- Further work might also be conducted by expanding the work from chapter 5 into both AR and VR space. These fields of research are currently undergoing a renaissance of sorts as the technology and software is becoming more accessible and easy to use, allowing creators and researchers to explore and expand upon current and new research. The results presented in chapter 5, were displayed on a 2D touch screen object, which limits both the results that were obtained but it also limits the immersion with the cultural heritage artefact. A further study could be conducted with an AR or a VR device, with the same set up as the studies conducted in chapter 5, making it possible to explore if the results obtained in that chapter still hold true. It would also provide results regarding if the polygonal and texture resolutions still play a major impact on the perceptual experience with cultural heritage artefacts. This would be very important, especially when VR devices normally present a lower quality render versus a monitor, but still achieves a more immersive experience.

6.5 Research impact for National Museums Liverpool

The research and work conducted as part of this thesis has been extremely beneficial for NML, offering them in depth knowledge and practical applications [139, 256]. It has evolved naturally with the increased penetration of 3D content within cultural institutes, and with the emergence of WebGL, AR and VR. This thesis has addressed the complexities of IP law, and the impact it plays on the dissemination of 3D cultural heritage artefacts for NML. It has also provided tools to display and explore how visitors prefer to interact with 3D cultural heritage. The research also provided a methodology that would allow NML to reduce the polygonal resolution to 80% of the original dataset and provide a texture without reducing the perceived perceptual quality of the cultural heritage datasets. With the knowledge generated from this thesis and the increased penetration of 3D content within other industries and cultural institutions; it has provided an opportunity for NML to be at the forefront for the display and dissemination of high resolution 3D cultural artefacts.

References

1. Google, 2011. *Google Arts and Culture* [Online] https://artsandculture.google.com/ [Accessed June 2012]
2. Whitfield, C.W./ S., 2014.. *Digitisation at the British Library* [online]. http://www.bl.uk/aboutus/stratpolprog/digi/digitisation/ [Accessed December 2014]
3. Hughes, L.M. ed., 2012. *Evaluating and measuring the value, use and impact of digital collections*. Facet Publishing.
4. La Pensée, A., 2008 3D Laser Scanning in 3D Documentation and Digital Reconstruction of Cultural Heritage, *JISC 3D Visualisation in the Arts Network - 3DVisA Bulletin Issue 4*
5. UNESCO, 2015. *Launch of global Unite for Heritage Coalition in Bonn* [online] Available at: http://whc.unesco.org/en/news/1303/ [Accessed October 2016]
6. UNESCO, 2015. *Recommendation concerning the preservation of, and access to, documentary heritage including in digital form* [online] Available at: http://portal.unesco.org/en/ev.php-URL_ID=49358&URL_DO=DO_TOPIC&URL_SECTION=201.html [Accessed October 2016]
7. Barton, J. 2007. *Preservation by Documentation: The Cyark Commitment* [online] Available at: http://www.cyark.org/news/preservation-by-documentation-the-cyark-commitment [Accessed March 2013]
8. Cooperstock, J.R., 2011. Multimodal telepresence systems. *IEEE Signal Processing Magazine*, *28*(1), pp.77-86.
9. Occulus. *Occulus Rift* [online] Available at: https://www.oculus.com/ [Accessed September 2016]
10. Vive, *Vive* [online] Available at: https://www.vive.com/uk/ [Accessed September 2016]
11. European Patents Office and Office for Internal Harmonization in the Internal Market, 2013. *Intellectual property rights intensive industries: Contribution to economic performance and employment in the European Union Industry-Level Analysis Report* [online] Available at: http://documents.epo.org/projects/babylon/eponet.nsf/0/8E1E34349D4546C3C1257BF300343D8B/$File/ip_intensive_industries_en.pdf [Accessed June 2014]
12. Rainie, L., 2010, February. Networked creators: how users of social media have changed the ecology of information. In *VALA Libraries 2010 Conference, Melbourne*.
13. Denard, H., The London Charter for the computer-based Visualisation of cultural heritage. Draft 2.1, 7 February 2009.
14. Oxford Dictionary, *The definition of Accuracy in English* [online] Available at: https://en.oxforddictionaries.com/definition/accuracy [Accessed December 2017]
15. Shao, G., 2009. Understanding the appeal of user-generated media: a uses and gratification perspective. *Internet Research*, *19*(1), pp.7-25.
16. National Museums Liverpool, 2013. *Strategic Plan 2011 – 2015* [online] Available at: http://liverpoolmuseums.org.uk/about/corporate/strategic-plan/Strategic-Plan-2011-2015- Year-3-Interim-Annual-Plan-2013--2014.pdf [Accessed June 2015]
17. Antiquesportfolio.com plc v Rodney Fitch & Co Ltd. 2001. FSR 345
18. Bridgeman Art library, Ltd. V. Corel Corp. 2000, 36 F. Supp. 2d 191 (S.D.N.Y. 1999).
19. New Media Consortium, 2012. *Horizon Report –2012 Museum Edition* [online] Available at: http://www.nmc.org/pdf/2012-horizon-report-museum.pdf [Accessed February 2013]
20. British Library, 2013. *Digitized Manuscript*s [online] http://www.bl.uk/manuscripts/About.aspx [Accessed December 2013]
21. Creative Commons, 2013. *About The Licenses* [online] http://creativecommons.org/licenses/ [Accessed Mar 2014]
22. Copyright, Designs and Patents Act, 1988. Section 2.
23. Copyright, Designs and Patents Act, 1988. Section 11
24. Torsen, M. and Anderson, J., 2010. *Intellectual property and the safeguarding of traditional cultures: Legal issues and practical options for museums, libraries and archives* (p. 51). WIPO.
25. Intellectual Property Office, 2013. *The rights granted by copyright* [online] Available at: https://www.gov.uk/guidance/the-rights-granted-by-copyright [Accessed June 2015]
26. Patents Act 1977
27. Trade Marks Act 1994
28. Berne convention for the protection of literary and artistic works 1886, *completed at Paris on May 4, 1896, revised at Berlin on November 13, 1908, completed at Berne on March 20, 1914, revised at Rome on June 2, 1928, revised at Brussels on June 26, 1948, and revised at Stockholm on July 14, 1967*. Geneva, United International Bureaux for the Protection of Intellectual Property.
29. The Gillette Company v. LA-Laboratories Ltd. 2005. Oy Case C-228/03 All ER (EC) 940 at para. 42.
30. Intellectual Property Office, 2017. *Design Rights* [online] Available at: https://www.gov.uk/design-right [Accessed December 2017]
31. Mendis, D., 2013. The Clone Wars: Episode 1–The Rise of 3D Printing and its implications for Intellectual Property Law–Learning Lessons from the Past?. *EIPR*, *35*(3), pp.155-169.
32. Korn, N. 2005. *Guide to Intellectual Property Rights and Other Legal Issues.* [online] Available at: http://www.minervaeurope.org/publications/guideipr1_0.pdf [Accessed July 2014]
33. Intellectual Property office, 2014. *The Copyright and Rights in Performances (Quotation and Parody) Regulations 2014* [online] Available at: http://www.legislation.gov.uk/ukdsi/2014/9780111112717 [Accessed June 2014]
34. Dan Anthony, 2016 *National Treasures* [online] Available at: https://ipo.blog.gov.uk/2016/07/07/national-treasures/ [Accessed September 2016]
35. Kelly, K., 2013. Images of works of art in museum collections: the experience of open access; a study of 11 museums.
36. Pantalony, R,E. 2013. A WIPO Guide on Managing Intellectual Property for Museums.
37. Copyright, Designs and Patents Act, 1988
38. Feist Publications, Inc., v. Rural Telephone Service Co. 1991, 499 U.S. 340
39. Museum Copyright Group, 2013. *Copyright in Photographs of Works of Art* [online] Available at: http://museumscopyright.org.uk/resources/articles/bridgeman/ [Accessed July 2014]
40. Target Wire, 2009. *Conference on copyright law in the visual arts puts Bridgeman v Corel back in the spotlight* [online] Available at: http://www.targetwire.com/targetwire/2007/05/17/po180/po180_uk.html [Accessed July 2014]
41. Temple Island Collections Ltd v New English Teas Ltd. 2012 EWPCC 1

42. UK Intellectual Property Office, *Copyright Notice: digital images, photographs and the internet* [online] Avaliable at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/481194/c-notice-201401.pdf [Accessed November 2015]
43. Express Newspapers v Liverpool Daily Post & Echo Plc. 1985. FSR 306
44. Ong, B., 2010. Originality from copying: fitting recreative works into the copyright universe. *Intellectual Property Quarterly*, *2*, pp.165-191.
45. Painer v. Standard Verlags GmbH, Axel Springer AG, Süddeutsche Zeitung GmbH, Spiegel-Verlag Rudolf Augstein GmbH & Co. KG, Verlag M. DuMont Schauberg Expedition der Kölnischen Zeitung GmbH & Co. KG. 2012. ECDR 6
46. Intellectual Property Office comment. 2016, National Treasures [online] Available at: https://ipo.blog.gov.uk/2016/07/07/national-treasures/#comment-34489 [Accessed September 2016]
47. Meshwerks, Inc. v. Toyota Motor Sales USA Inc. 2008. 528 F.3D 1258 (10th Cir.)
48. Weinberg, M, 2013. *What's the Deal with Copyright and 3D printing?* [online] Available at: http://www.publicknowledge.org/files/What's%20the%20Deal%20with%20Copyright_%20Final%20version2.pdf [Accessed August 2014]
49. Burrow-Giles Lithographic Co. v. Sarony. 1985. 111 US 53
50. Bradshaw S., Bowyer A., & Haufe P. 2010, The Intellectual Property Implications of Low Cost 3DPrinting *Vol. 7, Issue 1 Script-ed pp. 1-31*
51. Copyright, Designs and Patents Act, 1988. Section 4.
52. Lucasfilm Ltd. & Ors v Ainsworth Anor. 2011. UKSC 39
53. Intellectual Property Office. 2015 *Consultation on New Transitional Provisions for the repeal of section 52 of the Copyright, Designs and Patents Act 1988: Government Response and Summary of Responses*, [online] Available at: https://www.gov.uk/government/consultations/transitional-arrangements-for-the-repeal-of-section-52-cdpa [accessed September 2016]
54. Moody, G, 2016 *UK copyright extension on designed objects is "direct assault" on 3D printing* [online] Available at: http://arstechnica.co.uk/tech-policy/2016/08/uk-copyright-extension-designed-objects-3d-printing/ [Accessed September 2016]
55. Koller, D., Turitzin, M., Levoy, M., Tarini, M., Croccia, G., Cignoni, P. and Scopigno, R., 2004, August. Protected interactive 3D graphics via remote rendering. In *ACM Transactions on Graphics (TOG)* (Vol. 23, No. 3, pp. 695-703). ACM
56. Digital Millennium Copyright Act, 1998
57. Copyright, Designs and Patents Act, 1988. Section 11
58. Freedom of Information Act, 2000.
59. Wohlers, T., 2013. Wohlers Report 2013: Additive Manufacturing and 3D Printing State of the Industry–Annual Worldwide Progress Report, Wohlers Associates. *Fort Collins, CO*.
60. Koller, D. and Levoy, M., 2005. Protecting 3d graphics content. *Communications of the ACM*, *48*(6), pp.74-80.
61. Koller, D., Frischer, B. and Humphreys, G., 2009. Research challenges for digital archives of 3D cultural heritage models. *Journal on Computing and Cultural Heritage (JOCCH)*, *2*(3), p.7.
62. Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L., Ginzton, M., Anderson, S., Davis, J., Ginsberg, J. and Shade, J., 2000, July. The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*(pp. 131-144). ACM Press/Addison-Wesley Publishing Co.
63. Zhu, J., Bakdash, J.Z., Koller, D., Banton, T., Proffitt, D.R. and Humphreys, G., 2008, August. Quantifying usability in secure graphics: assessing the user costs of protecting 3D content. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization* (pp. 91-96). ACM.
64. Google, 2013. *Chrome Experiments* [online] Available at: Http://www.chromeexperiments.com/ [Accessed November 2013]
65. Behr, J., Jung, Y., Keil, J., Drevensek, T., Zoellner, M., Eschler, P. and Fellner, D., 2010, July. A scalable architecture for the HTML5/X3D integration model X3DOM. In *Proceedings of the 15th International Conference on Web 3D Technology* (pp. 185-194). ACM.
66. Deveci, H.A., 2012. Can hyperlinks and digital rights management secure affordable access to information?. *Computer Law & Security Review*, *28*(6), pp.651-661.
67. Jung, E.K., Levien, R.A., Lord, R.W., Malamud, M.A., Rinaldo Jr, J.D., Tegreene, C.T. and Wood, L.L., Invention Science Fund I LLC, 2012. *Manufacturing control system*. U.S. Patent 8,286,236.
68. Kain, E. 2013. *EA Admits SimCity Could Have Run Offline* [online] Available at: http://www.forbes.com/sites/erikkain/2013/03/17/ea-admits-simcity-could-have-run-offline-went-against-developers-vision/ [Accessed September 2014]
69. Walker, J. 2012. *Ubisoft Scrapping Always-on DRM for PC Games* [online] Available at: http://www.rockpapershotgun.com/2012/09/05/ubisoft-scrapping-always-on-drm-for-pc-games/ [Accessed August 2014]
70. Fogliatto, F.S. and Da Silveira, G.J. eds., 2010. *Mass customization: engineering and managing global operations*. Springer Science & Business Media.
71. FigurePrints, *FigurePrints Frequently Asked Questions* [online] Available at: http://www.figureprints.com/wow/Help/FAQ.aspx [Accessed September 2014]
72. The British Museum, 2014. *Night at the museum: Fact vs Fiction* [online] Available at: http://www.britishmuseum.org/visiting/family_visits/night_at_the_museum/fact_vs_fiction.aspx [Accessed November 2014]
73. Terrassa, J. 2012. *Met - 3D: The Museums First 3D Scanning and Printing Hackathon* [online] Available at: Http://www.metmuseum.org/about-the-museum/now-at-the-met/features/2012/hackathon [Accessed November 2014]
74. Vernik, D.A., Purohit, D. and Desai, P.S., 2011. Music downloads and the flip side of digital rights management. *Marketing Science*, *30*(6), pp.1011-1027.
75. Sudler, H., 2013. Effectiveness of anti-piracy technology: Finding appropriate solutions for evolving online piracy. *Business Horizons*, *56*(2), pp.149-157.

76. Crisp, J. 2013. *One Year Later, the Results of Tor Books UK Going DRM-Free* [online] Available at: http://www.tor.com/blogs/2013/04/tor-books-uk-drm-free-one-year-later [Accessed December 2014]
77. Zhang, L., 2016. Intellectual property strategy and the long tail: Evidence from the recorded music industry. *Management Science*.
78. Motion Pictures Association of America, 2011. *Cost of Content Theft By the Numbers* [online] Available at: http://www.scribd.com/doc/62848402/mpaa-infographic [Accessed December 2014]
79. Moses, A. 2011. *Piracy, are we being conned?* [online] Available at: http://www.smh.com.au/technology/technology-news/piracy-are-we-being-conned-20110322-1c4cs.html [Accessed January 2015]
80. Wilson, S. L. 2011. *Box Office Records and Whining Twatwaffles of the MPAA* [online] Available at: http://www.pajiba.com/trade_news/box-office-records-and-whining-twatwaffles-of-the-mpaa.php [Accessed Jan 2015]
81. Scott K., 2012 *The Pirate Bay Adds ''Physibles' 3D Printing Category* [online] Available at: http://www.wired.co.uk/news/archive/2012-01/24/pirate-bay-introduces-physibles [Accessed June 2014]
82. Vatu, G., 2014. *The PirateBay reaches 10 millionth torrent milestone* [online] Available at: http://news.softpedia.com/news/The-PirateBay-Reaches-10-Millionth-Torrent-Milestone-438660.shtml [Accessed January 2015]
83. Thingiverse, 2015 *Makerbot Thingiverse* [online] Available at: http://www.thingiverse.com/about [Accessed Jan 2015]
84. Shapeways, *Shapeways Content Policy and Notice Takedown Procedure*,[online] Available at: http://www.shapeways.com/legal/content_policy [Accessed Jan 2015]
85. Lerch, T., MacGillivray, M. and Domina, T., 2007. 3D Laser Scanning: A Model of multidisciplinary research. *Journal of Textile and Apparel, Technology and Management*, *5*(4), pp.1-8.
86. Bernardini, F. and Rushmeier, H., 2002, June. The 3D model acquisition pipeline. In *Computer graphics forum* (Vol. 21, No. 2, pp. 149-172). Blackwell Publishers Ltd.
87. Larson, J., 1994. The use of laser scanners for recording and reproducing sculpture. *Conservation news (United Kingdom Institute for Conservation of Historic and Artistic Works)*, (54), pp.10-11.
88. Levoy, M., 1999. The Digital Michelangelo project: creating a 3D archive of his sculpture using laser scanning. *Proceedings of Electronic Imaging and the Visual Arts (EVA)*, *99*, pp.22-26.
89. Smithsonian Museum of National Museum of Natural History, *What does it mean to be human?* [online] Available at: http://humanorigins.si.edu/evidence/3d-collection [Accessed September 2013]
90. Europeana, *3D Icons*, [online] Available at: https://pro.europeana.eu/project/3d-icons, [Accessed November 2013]
91. Jung, Y., Behr, J. and Graf, H., 2011, March. X3DOM as carrier of the virtual heritage. In *International Workshop 3D-ARCH* (Vol. 4, p. 2011).
92. Schifferstein, H.N. and Zwartkruis-Pelgrim, E.P., 2008. Consumer-product attachment: Measurement and design implications. *International journal of design*, *2*(3).
93. G. Sverdlov, R Reitsma, and S. Jaddou, Understanding the changing needs of US Online Consumers 2012. North American Consumer Technographics, 2012
94. Piazza, J. and Bering, J.M., 2009. Evolutionary cyber-psychology: Applying an evolutionary framework to Internet behavior. *Computers in Human Behavior*, *25*(6), pp.1258-1269.
95. Sketchfab, *Enter New Dimensions*, [Online] Available at: https://sketchfab.com [Accessed September 2016]
96. Cairns, P. and Cox, A.L. eds., 2008. *Research methods for human-computer interaction* (Vol. 12). Cambridge: Cambridge University Press.
97. Khronos, *WebGL- OpenGL ES for the Web*, [online] Available at: http://www.khronos.org/webgl/ [Accessed March 2012]
98. Levoy, M. and Hanrahan, P., 1996, August. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (pp. 31-42). ACM
99. Wood, D.N., Azuma, D.I., Aldinger, K., Curless, B., Duchamp, T., Salesin, D.H. and Stuetzle, W., 2000, July. Surface light fields for 3D photography. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (pp. 287-296). ACM Press/Addison-Wesley Publishing Co.
100. Mudge, M., Schroer, C., Earl, G., Martinez, K., Pagi, H., Toler-Franklin, C., Rusinkiewicz, S., Palma, G., Wachowiak, M., Ashey, M. and Mathews, N., 2010. Principles and practices of robust, photography-based digital imaging techniques for museums. *In, 11th VAST International Symposium on Virtual Reality, Archaeology and Cultural Heritage*
101. Adobe. *Adobe Flash Player 11 / Features* [online] Available at: http://www.adobe.com/products/flashplayer/features.html [Accessed march 2012]
102. Adobe, 2009.*AIR Passes 100 Million Installs* [online] Available at: http://blogs.adobe.com/flashplayer/2009/01/air_passes_100_million_install.html?sdid=EENCL [Accessed March 2012]
103. World Wide Web Consortium (W3C) 2012.*HTML 5* [online] Available at: http://www.w3.org/TR/html5/the-canvas-element.html#security-with-canvas-elements [Accessed March 2012]
104. TIME Techland, 2010. *Steve Jobs:* 'Flash is no longer Necessary' and other Musings [online] Available at: http://dev.w3.org/html5/spec/introduction.html#history-1 [Accessed March 2012]
105. Deep Blue Sky, 2012. *HTML5 and CSS3 Support* [online] Available at: http://www.findmebyip.com/litmus [Accessed March 2012]
106. Adobe, 2011. *Flash to Focus on PC Browsing and Mobile Apps; Adobe to more aggressively contribute to HTML5* [online] (last update : 9 November 2011) Available at: http://blogs.adobe.com/conversations/2011/11/flash-focus.html[Accessed February 2012]
107. Adobe, 2012. *Adobe and Google Partnering for Flash Player on Linux* [online] Available at: http://blogs.adobe.com/flashplayer/2012/02/adobe-and-google-partnering-for-flash-player-on-linux.html [Accessed March 2012]
108. Adobe, 2017, *Flash and the Future of interactive content* [online] Available at: https://theblog.adobe.com/adobe-flash-update/ [accessed August 2017]
109. Scabia, M., 2011. *How Stage3D works* [online] Available at: http://www.adobe.com/devnet/flashplayer/articles/how-stage3d-works.html [Accessed March 2012]

110. Context, 2011. *WebGL – A New Dimension for Browser Exploitation* [online] Available at: http://www.contextis.com/research/blog/webgl [Accessed February 2012]
111. WebGL-loader. 2012. *Happy Buddha: Loading 1 million Triangles* [online] Available at: http://code.google.com/p/webgl-loader/wiki/HappyBuddha [Accessed March 2012]
112. Chromium, 2010. *Introducing the ANGLE Project* [online] Available at: http://blog.chromium.org/2010/03/introducing-angle-project.html [Accessed February 2012]
113. Airtight Interactive, 2011. *Stage3D vs WebGL Performance* [online] Available at: http://www.airtightinteractive.com/2011/10/stage3d-vs-webgl-performance [Accessed February 2012]
114. Jung, Y., Behr, J. and Graf, H., 2011, March. X3DOM as carrier of the virtual heritage. In *International Workshop 3D-ARCH* (Vol. 4, p. 2011).
115. Schwartz, C., Ruiters, R., Weinmann, M. and Klein, R., 2013. WebGL-based streaming and presentation of objects with bidirectional texture functions. *Journal on Computing and Cultural Heritage (JOCCH)*, *6*(3), p.11.
116. Rodriguez, M.B., Agus, M., Bettio, F., Marton, F. and Gobbetti, E., 2015, September. Digital Mont'e Prama: 3D Cultural Heritage presentations in museums and anywhere. In *Digital Heritage, 2015* (Vol. 2, pp. 557-564). IEEE.
117. Barsanti, S.G., Caruso, G., Micoli, L.L., Rodriguez, M.C. and Guidi, G., 2015. 3D visualization of cultural heritage artefacts with virtual reality devices. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, *40*(5), p.165.
118. Di Franco, P.D.G., Camporesi, C., Galeazzi, F. and Kallmann, M., 2015. 3D printing and immersive visualization for improved perception of ancient artifacts. *Presence: Teleoperators and Virtual Environments*, *24*(3), pp.243-264.
119. Klimke, J., Hagedorn, B., Trapp, M. and Döllner, J., 2014. Web-based and mobile provisioning of virtual 3d reconstructions. In *Tagungsband der 12. Konferenz Kultur und Informatik: Reality and Virtuality* (pp. 17-28).
120. WebGL Stats, *WebGL stats* [online] Available at: Http://webglstats.com/ [Accessed June 2014]
121. WebGL Stats, 2016 WebGL Stats [online] Available at: Http://webglstats.com/ [Accessed November 2016]
122. Koan Interactive, 2012. *WebGL in enterprise apps? Uh, no*. [online] Available at: Http://blog.kaon.com/2012/05/11/webgl-in-the-enterprise/ [Accessed June 2014]
123. Mozilla. *Blacklisting/Blocked Graphic Drivers [*online] Available at: https://wiki.mozilla.org/Blocklisting/Blocked_Graphics_Drivers [Accessed April 2014]
124. Khronos. *WebGL Security* [online] Available at: Http://www.khronos.org/webgl/security [Accessed March 2013]
125. 3D Coform, *3D content,* [online] Available at: http://www.3d-coform.eu/index.php/3d-content [Accessed April 2014]
126. 3D Coform, *X3Dom 3D data catalog*, [online] Available at: http://3d-coform.eu/index.php/3d-content/47-3d-content/3d-content/99-x3dom-3d-data-catalog6 [Accessed April 2014]
127. Prahalad, C.K. and Ramaswamy, V., 2004. Co-creation experiences: The next practice in value creation. *Journal of interactive marketing*, *18*(3), pp.5-14.
128. Arora, H., 2009. A conceptual study of brand communities. *The Icfai University Journal of Brand Management*, *6*(2), pp.7-21
129. Turner, P., 2010, August. The anatomy of engagement. In *Proceedings of the 28th Annual European Conference on Cognitive Ergonomics* (pp. 59-66). ACM.
130. Marcum, J.W., 2000. Out with motivation, in with engagement. *Global Business and Organizational Excellence*, *19*(4), pp.57-60.
131. Zayas, B. and Bustos, J.V., 2011, May. Designing a 3D desktop virtual environment for teaching. In *Proceedings of the 2011 International Conference on Computers and Computing, Bangkok* (pp. 125-130).
132. Usability.Gov, 2017. *Prototyping* [online] Available at: https://www.usability.gov/how-to-and-tools/methods/prototyping.html [Accessed December 2017]
133. Nielsen, J., 2003. Paper prototyping: Getting user data before you code. *Last Reviewed on September*, *22*, p.2007.
134. Tan, O.S., 2011. Problem-based learning approach to human computer interaction. *World Academy of Science, Engineering and Technology, International Journal of Social, Behavioral, Educational, Economic, Business and Industrial Engineering*, *5*(4), pp.356-359
135. Gentner, D. and Nielsen, J., 1996. The anti-mac interface. *Communications of the ACM*, *39*(8), pp.70-82.
136. M. Geelnard, *OpenCTM* [online] Available at: http://openctm.sourceforge.net/ [Accessed March 2013]
137. Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F. and Ranzuglia, G., 2008, July. Meshlab: an open-source mesh processing tool. In *Eurographics Italian Chapter Conference* (Vol. 2008, pp. 129-136).
138. Pixologic, *Zbrush Features*, [online] Available at: http://pixologic.com/features/ [Accessed September 2016]
139. National Museums Liverpool. 2014. *Pre-Hispanic Caribbean sculptural arts project*, [online] Available at: http://www.liverpoolmuseums.org.uk/collections/research/caribbean-sculpture/ [Accessed Jan 2015]
140. National Museums Liverpool, *3D Interactive 3D model of the Turin Cemi* [Online] Available at: http://www.liverpoolmuseums.org.uk/collections/research/caribbean-sculpture/turin-cemi/3d-viewer.aspx [Accessed September 2013]
141. 3D Hubs, *Local Manufacturing*, [online] Available at: https://www.3dhubs.com/ [Accessed September 2016]
142. Marmoset, *Marmoset Viewer*, [online] Available at: http://www.marmoset.co/viewer?utm_source=inapp&utm_medium=menu&utm_campaign=viewer [Accessed September 2016]
143. Maker Bot, 2011, 3D preview in your browser on thingieverse [online] available at: https://www.makerbot.com/media-center/2011/02/17/thingiview-js-on-thingiverse [Accessed December 2017]
144. P3D.in, 2017 Welcome to p3d.in, your 3D online [online] available at: https://p3d.in/ [Accessed December 2017]
145. National Museums Liverpool, *3D Models* [online] Available at: http://www.liverpoolmuseums.org.uk/wml/visit/floor-plans/ancient-egypt/3d-models.aspx [Accessed April 2017]
146. Sketchfab Admin, 2013. *50k Strong!* [Online] Available at: https://blog.sketchfab.com/50k-strong [Accessed September 2016]
147. Veldhuzen, Bart, 2016 *Half a million Sketchfab members!* [online] Available at: https://blog.sketchfab.com/half-million-sketchfab-members/# [Accessed September 2016]
148. Sheffer, A., Praun, E. and Rose, K., 2007. Mesh parameterization methods and their applications. *Foundations and Trends® in Computer Graphics and Vision*, *2*(2), pp.105-171.
149. Heeger, D.J. and Bergen, J.R., 1995, September. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (pp. 229-238). ACM.

150. Jagnow, R., Dorsey, J. and Rushmeier, H., 2004, August. Stereological techniques for solid textures. In *ACM Transactions on Graphics (TOG)* (Vol. 23, No. 3, pp. 329-335). ACM.
151. Kopf, J., Fu, C.W., Cohen-Or, D., Deussen, O., Lischinski, D. and Wong, T.T., 2007, August. Solid texture synthesis from 2d exemplars. In *ACM Transactions on Graphics (TOG)* (Vol. 26, No. 3, p. 2). ACM.
152. Daly, L. and Brutzman, D., 2007. X3D: Extensible 3D graphics standard [standards in a nutshell]. *IEEE Signal Processing Magazine*, *24*(6), pp.130-135.
153. Daly, S.J., 1992, August. Visible differences predictor: an algorithm for the assessment of image fidelity. In *Human Vision, Visual Processing, and Digital Display III* (Vol. 1666, pp. 2-16). International Society for Optics and Photonics.
154. Kundu, D. and Evans, B.L., 2015, September. Full-reference visual quality assessment for synthetic images: A subjective study. In *Image Processing (ICIP), 2015 IEEE International Conference on* (pp. 2374-2378). IEEE.
155. Floater, M.S. and Hormann, K., 2005. Surface parameterization: a tutorial and survey. In *Advances in multiresolution for geometric modelling* (pp. 157-186). Springer, Berlin, Heidelberg.
156. Bennis, C., Vézien, J.M. and Iglésias, G., 1991, July. Piecewise surface flattening for non-distorted texture mapping. In *ACM SIGGRAPH computer graphics* (Vol. 25, No. 4, pp. 237-246). ACM.
157. Maillot, J., Yahia, H. and Verroust, A., 1993, September. Interactive texture mapping. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*(pp. 27-34). ACM.
158. Pinkall, U. and Polthier, K., 1993. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, *2*(1), pp.15-36.
159. Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M. and Stuetzle, W., 1995, September. Multiresolution analysis of arbitrary meshes. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*(pp. 173-182). ACM.
160. Lévy, B. and Mallet, J.L., 1998, July. Non-distorted texture mapping for sheared triangulated meshes. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (pp. 343-352). ACM.
161. Sander, P.V., Snyder, J., Gortler, S.J. and Hoppe, H., 2001, August. Texture mapping progressive meshes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 409-416). ACM.
162. Floater, M.S., 1997. Parametrization and smooth approximation of surface triangulations. *Computer aided geometric design*, *14*(3), pp.231-250.
163. Floater, M.S., 2003. Mean value coordinates. *Computer aided geometric design*, *20*(1), pp.19-27
164. Lévy, B., Petitjean, S., Ray, N. and Maillot, J., 2002, July. Least squares conformal maps for automatic texture atlas generation. In *ACM transactions on graphics (TOG)* (Vol. 21, No. 3, pp. 362-371). ACM.
165. Desbrun, M., Meyer, M. and Alliez, P., 2002, September. Intrinsic parameterizations of surface meshes. In *Computer graphics forum* (Vol. 21, No. 3, pp. 209-218). Blackwell Publishing, Inc.
166. Sheffer, A. and de Sturler, E., 2001. Parameterization of faceted surfaces for meshing using angle-based flattening. *Engineering with computers*, *17*(3), pp.326-337.
167. Sander, P.V., Hoppe, H., Gortler, S. and Snyder, J., 2002. Signal-specialized parameterization.
168. Lee, Y., Kim, H.S. and Lee, S., 2002. Mesh parameterization with a virtual boundary. *Computers & Graphics*, *26*(5), pp.677-686.
169. Sheffer, A., Lévy, B., Mogilnitsky, M. and Bogomyakov, A., 2005. ABF++: fast and robust angle based flattening. *ACM Transactions on Graphics (TOG)*, *24*(2), pp.311-330.
170. Tutte, W.T., 1963. How to draw a graph. *Proceedings of the London Mathematical Society*, *3*(1), pp.743-767.
171. Ray, N. and Levy, B., 2003, Hierarchical least squares conformal map. In *Computer Graphics and Applications, 2003. Proceedings. 11th Pacific Conference on*(pp. 263-270). IEEE.
172. Hormann, K. and Greiner, G., 2000. *MIPS: An efficient global parameterization method*. Erlangen – Nuernberg Univ (Germany) Computer Graphics Group
173. Degener, P., Meseth, J. and Klein, R., 2003. An Adaptable Surface Parameterization Method. *IMR*, *3*, pp.201-213.
174. Zayer, R., Lévy, B. and Seidel, H.P., 2007, July. Linear angle based parameterization. In *Fifth Eurographics Symposium on Geometry Processing-SGP 2007* (pp. 135-141). Eurographics Association.
175. Zigelman, G., Kimmel, R. and Kiryati, N., 2002. Texture mapping using surface flattening via multidimensional scaling. *IEEE Transactions on Visualization and Computer Graphics*, *8*(2), pp.198-207.
176. Zhou, K., Wang, X., Tong, Y., Desbrun, M., Guo, B. and Shum, H.Y., 2005. TextureMontage. *ACM Transactions on Graphics (TOG)*, *24*(3), pp.1148-1155.
177. Kraevoy, V., Sheffer, A. and Gotsman, C., 2003. *Matchmaker: constructing constrained texture maps* (Vol. 22, No. 3, pp. 326-333). ACM.
178. Lee, T.Y., Yen, S.W. and Yeh, I.C., 2008. Texture mapping with hard constraints using warping scheme. *IEEE Transactions on Visualization and Computer Graphics*, *14*(2), pp.382-395.
179. Pietroni, N., Tarini, M. and Cignoni, P., 2010. Almost isometric mesh parameterization through abstract domains. *IEEE Transactions on Visualization and Computer Graphics*, *16*(4), pp.621-635.
180. Zhang, E., Mischaikow, K. and Turk, G., 2005. Feature-based surface parameterization and texture mapping. *ACM Transactions on Graphics (TOG)*, *24*(1), pp.1-27.
181. Kharevych, L., Springborn, B. and Schröder, P., 2006. Discrete conformal mappings via circle patterns. *ACM Transactions on Graphics (TOG)*, *25*(2), pp.412-438.
182. Ben-Chen, M., Gotsman, C. and Bunin, G., 2008, April. Conformal flattening by curvature prescription and metric scaling. In *Computer Graphics Forum* (Vol. 27, No. 2, pp. 449-458). Blackwell Publishing Ltd.
183. Yang, Y.L., Kim, J., Luo, F., Hu, S.M. and Gu, X., 2008. Optimal surface parameterization using inverse curvature map. *IEEE Transactions on Visualization and Computer Graphics*, *14*(5), pp.1054-1066.
184. Lee, A.W., Sweldens, W., Schröder, P., Cowsar, L. and Dobkin, D., 1998, July. MAPS: Multiresolution adaptive parameterization of surfaces. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (pp. 95-104). ACM.
185. Praun, E. and Hoppe, H., 2003, July. Spherical parametrization and remeshing. In *ACM Transactions on Graphics (TOG)* (Vol. 22, No. 3, pp. 340-349). ACM.
186. Khodakovsky, A., Litke, N. and Schröder, P., 2003, July. Globally smooth parameterizations with low distortion. In *ACM Transactions on Graphics (TOG)* (Vol. 22, No. 3, pp. 350-357). ACM.
187. Tarini, M., Hormann, K., Cignoni, P. and Montani, C., 2004, August. Polycube-maps. In *ACM transactions on graphics (TOG)* (Vol. 23, No. 3, pp. 853-860). ACM.

188. Eckstein, I., Surazhsky, V. and Gotsman, C., 2001, September. Texture mapping with hard constraints. In *Computer Graphics Forum* (Vol. 20, No. 3, pp. 95-104). Blackwell Publishers Ltd.
189. Yu, H., Lee, T.Y., Yeh, I.C., Yang, X., Li, W. and Zhang, J.J., 2012. An rbf-based reparameterization method for constrained texture mapping. *IEEE Transactions on Visualization and Computer Graphics*, *18*(7), pp.1115-1124.
190. Tang, Y., Wang, J., Bao, H. and Peng, Q., 2003. RBF-based constrained texture mapping. *Computers & Graphics*, *27*(3), pp.415-422.
191. Lee, T.Y. and Huang, P.H., 2003. Fast and intuitive metamorphosis of 3d polyhedral models using smcc mesh merging scheme. *IEEE Transactions on Visualization and Computer Graphics*, *9*(1), pp.85-98.
192. Perlin, K., 1985. An image synthesizer. *ACM Siggraph Computer Graphics*, *19*(3), pp.287-296.
193. Ghazanfarpour, D. and Dischler, J.M., 1995. Spectral analysis for automatic 3-d texture generation. *Computers & Graphics*, *19*(3), pp.413-422.
194. Ghazanfarpour, D. and Dischler, J.M., 1996, August. Generation of 3D texture using multiple 2D models analysis. In *Computer Graphics Forum* (Vol. 15, No. 3, pp. 311-323). Blackwell Science Ltd.
195. Pietroni, N., Otaduy, M.A., Bickel, B., Ganovelli, F. and Gross, M., 2007, September. Texturing internal surfaces from a few cross sections. In *Computer Graphics Forum* (Vol. 26, No. 3, pp. 637-644). Blackwell Publishing Ltd.
196. Dischler, J.M., Ghazanfarpour, D. and Freydier, R., 1998, August. Anisotropic solid texture synthesis using orthogonal 2D views. In *Computer Graphics Forum* (Vol. 17, No. 3, pp. 87-95). Blackwell Publishers Ltd.
197. Jagnow, R., Dorsey, J. and Rushmeier, H., 2008. Evaluation of methods for approximating shapes used to synthesize 3d solid textures. *ACM Transactions on Applied Perception (TAP)*, *4*(4), p.5.
198. Qin, X. and Yang, Y.H., 2005, October. Basic gray level aura matrices: theory and its application to texture synthesis. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (Vol. 1, pp. 128-135). IEEE.
199. Elfadel, I.M. and Picard, R.W., 1994. Gibbs random fields, cooccurrences, and texture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *16*(1), pp.24-37.
200. Qin, X. and Yang, Y.H., 2007. Aura 3D textures. *IEEE transactions on visualization and computer graphics*, *13*(2), pp.379-389.
201. Wei, L.Y. and Levoy, M., 2000, July. Fast texture synthesis using tree-structured vector quantization. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (pp. 479-488). ACM Press/Addison-Wesley Publishing Co..
202. Wei, L.Y., 2003, July. Texture synthesis from multiple sources. In *ACM Siggraph 2003 Sketches & Applications* (pp. 1-1). ACM.
203. Kwatra, V., Essa, I., Bobick, A. and Kwatra, N., 2005, July. Texture optimization for example-based synthesis. In *ACM Transactions on Graphics (ToG)* (Vol. 24, No. 3, pp. 795-802). ACM.
204. Chen, J. and Wang, B., 2010. High quality solid texture synthesis using position and index histogram matching. *The Visual Computer*, *26*(4), pp.253-262.
205. Dong, Y., Lefebvre, S., Tong, X. and Drettakis, G., 2008, June. Lazy solid texture synthesis. In *Computer Graphics Forum* (Vol. 27, No. 4, pp. 1165-1174). Blackwell Publishing Ltd.
206. Ashikhmin, M., 2001, March. Synthesizing natural textures. In *Proceedings of the 2001 symposium on Interactive 3D graphics* (pp. 217-226). ACM.
207. Seo, M.K., Kim, H.M. and Lee, K.H., 2014. Solid texture synthesis for heterogeneous translucent materials. *The Visual Computer*, *30*(3), pp.271-283.
208. Zhang, G.X., Du, S.P., Lai, Y.K., Ni, T. and Hu, S.M., 2011. Sketch guided solid texturing. *Graphical Models*, *73*(3), pp.59-73.
209. Wang, L., Zhou, K., Yu, Y. and Guo, B., 2010. Vector solid textures. *ACM Transactions on Graphics (TOG)*, *29*(4), p.86.
210. Zhou, K., Ren, Z., Lin, S., Bao, H., Guo, B. and Shum, H.Y., 2008, August. Real-time smoke rendering using compensated ray marching. In *ACM Transactions on Graphics (TOG)* (Vol. 27, No. 3, p. 36). ACM.
211. Welsh, D.J. and Powell, M.B., 1967. An upper bound for the chromatic number of a graph and its application to timetabling problems. *The Computer Journal*, *10*(1), pp.85-86.
212. Zhang, G.X., Lai, Y.K. and Hu, S.M., 2013. Efficient synthesis of gradient solid textures. *Graphical Models*, *75*(3), pp.104-117.
213. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B. and Salesin, D.H., 2001, August. Image analogies. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 327-340). ACM.
214. Rademacher, P., Lengyel, J., Cutrell, E. and Whitted, T., 2001. Measuring the perception of visual realism in images. In *Rendering Techniques 2001* (pp. 235-247). Springer, Vienna.
215. BerjóN, D., MoráN, F. and Manjunatha, S., 2013. Objective and subjective evaluation of static 3D mesh compression. *Signal Processing: Image Communication*, *28*(2), pp.181-195.
216. Lavoué, G. and Mantiuk, R., 2015. Quality assessment in computer graphics. In *Visual Signal Quality Assessment* (pp. 243-286). Springer, Cham.
217. Banterle, F., Ledda, P., Debattista, K., Bloj, M., Artusi, A. and Chalmers, A., 2009, March. A psychophysical evaluation of inverse tone mapping techniques. In *Computer Graphics Forum* (Vol. 28, No. 1, pp. 13-25). Blackwell Publishing Ltd.
218. Masia, B., Agustin, S., Fleming, R.W., Sorkine, O. and Gutierrez, D., 2009. Evaluation of reverse tone mapping through varying exposure conditions. *ACM transactions on graphics (TOG)*, *28*(5), p.160.
219. Corsini, M., Larabi, M.C., Lavoué, G., Petřík, O., Váša, L. and Wang, K., 2013, February. Perceptual metrics for static and dynamic triangle meshes. In *Computer Graphics Forum* (Vol. 32, No. 1, pp. 101-125). Blackwell Publishing Ltd.
220. Watson, B., Friedman, A. and McGaffey, A., 2001, Measuring and predicting visual fidelity. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 213-220). ACM.
221. Guo, J., Vidal, V., Cheng, I., Basu, A., Baskurt, A. and Lavoue, G., 2017. Subjective and objective visual quality assessment of textured 3D meshes. *ACM Transactions on Applied Perception (TAP)*, *14*(2), p.11.
222. ITU-R BT.500-13: 2012, Methodology for the Subjective assessment of the quality of television picture.
223. Mantiuk, R.K., Tomaszewska, A. and Mantiuk, R., 2012, December. Comparison of four subjective methods for image quality assessment. In *Computer Graphics Forum* (Vol. 31, No. 8, pp. 2478-2491). Blackwell Publishing Ltd.
224. Rushmeier, H.E., Rogowitz, B.E. and Piatko, C., 2000, June. Perceptual issues in substituting texture for geometry. In *Human Vision and Electronic Imaging V* (Vol. 3959, pp. 372-384). International Society for Optics and Photonics.

225. Rogowitz, B.E. and Rushmeier, H.E., 2001, Are image quality metrics adequate to evaluate the quality of geometric objects?. In *Human Vision and Electronic Imaging VI* (Vol. 4299, pp. 340-349). International Society for Optics and Photonics.
226. Pan, Y., Cheng, I. and Basu, A., 2005. Quality metric for approximating subjective evaluation of 3-D objects. *IEEE Transactions on Multimedia*, *7*(2), pp.269-279.
227. Cleju, I. and Saupe, D., 2006, Evaluation of supra-threshold perceptual metrics for 3D models. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization* (pp. 41-44). ACM.
228. Lavoué, G., Gelasca, E.D., Dupont, F., Baskurt, A. and Ebrahimi, T., 2006, Perceptually driven 3D distance metrics with application to watermarking. In *Applications of Digital Image Processing XXIX* (Vol. 6312, p. 63120L). International Society for Optics and Photonics.
229. Lavoué, G., 2009. A local roughness measure for 3D meshes and its application to visual masking. *ACM Transactions on Applied perception (TAP)*, *5*(4), p.21.
230. Silva, S., Madeira, J., Ferreira, C. and Santos, B.S., 2007, Comparison of methods for the simplification of mesh models using quality indices and an observer study. In *Human Vision and Electronic Imaging XII* (Vol. 6492, p. 64921L). International Society for Optics and Photonics.
231. Váša, L. and Rus, J., 2012,. Dihedral angle mesh error: a fast perception correlated distortion measure for fixed connectivity triangle meshes. In *Computer Graphics Forum*(Vol. 31, No. 5, pp. 1715-1724). Blackwell Publishing Ltd.
232. Silva, S., Santos, B.S., Ferreira, C. and Madeira, J., 2009, A perceptual data repository for polygonal meshes. In *Visualisation, 2009. VIZ'09. Second International Conference in* (pp. 207-212). IEEE.
233. Subjective Quality Assessment of 3D models 2016, [online] Available at: http://liris.cnrs.fr/glavoue/data/datasets.html [Accessed March 2017]
234. Corsini, M., Gelasca, E.D., Ebrahimi, T. and Barni, M., 2007. Watermarked 3-D mesh quality assessment. *IEEE Transactions on Multimedia*, *9*(2), pp.247-256.
235. Lavoué, G., 2011, A multiscale metric for 3D mesh visual quality assessment. In *Computer Graphics Forum* (Vol. 30, No. 5, pp. 1427-1437). *Blackwell Publishing Ltd.*
236. Institute of Electronics and Informatics Engineering of Aveiro, *Overview* [online] Available at: http://wiki.ieeta.pt/wiki/index.php/Main_Page [Accessed September 2016]
237. Mesh Compression. *Data* [online] Available at: compression.kiv.zcu.cz [Accessed September 2016]
238. Mantiuk, R., Kim, K.J., Rempel, A.G. and Heidrich, W., 2011, HDR-VDP-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. In *ACM Transactions on graphics (TOG)* (Vol. 30, No. 4, p. 40). ACM.
239. O'Shea, J.P., Banks, M.S. and Agrawala, M., 2008, The assumed light direction for perceiving shape from shading. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization* (pp. 135-142). ACM.
240. Substance Painter, *3D painting Software*, [online] Available at: https://www.allegorithmic.com/products/substance-painter [Accessed September 2016]
241. Pixologic*, UV Master: UV creation has never been so easy!* [online] Available at: http://pixologic.com/zbrush/features/UV-Master/ [Accessed September 2016]
242. Garland, M. and Heckbert, P.S., 1998, October. Simplifying surfaces with color and texture using quadric error metrics. In *Visualization'98. Proceedings* (pp. 263-269). IEEE.
243. Thorn, J., Pizarro, R., Spanlang, B., Bermell-Garcia, P. and Gonzalez-Franco, M., 2016, October. Assessing 3D Scan Quality Through Paired-comparisons Psychophysics. In *Proceedings of the 2016 ACM on Multimedia Conference* (pp. 147-151). ACM.
244. Lavoué, G. and Corsini, M., 2010. A comparison of perceptually-based metrics for objective evaluation of geometry processing. *IEEE Transactions on Multimedia*, *12*(7), pp.636-649.
245. Silva, S., Santos, B.S., Madeira, J. and Ferreira, C., 2008, February. Perceived quality assessment of polygonal meshes using observer studies: A new extended protocol. In *Human Vision and Electronic Imaging XIII* (Vol. 6806, p. 68060D). International Society for Optics and Photonics.
246. Garrick, *Sculpture S0029* [online] Available at: http://garrick.ssl.co.uk/object-s0029 [Accessed December 2017]
247. W3C, *VRML Virtual Reality Modelling Language*, [online] Available at: https://www.w3.org/MarkUp/VRML/ [Accessed February 2017]
248. Unfold3D, *UV Mapping Solution* [online] Available at: http://www.polygonal-design.fr/ [accessed September 2016]
249. Mantiuk, R., Daly, S.J., Myszkowski, K. and Seidel, H.P., 2005, March. Predicting visible differences in high dynamic range images: model and its calibration. In *Human Vision and Electronic Imaging X* (Vol. 5666, pp. 204-215). International Society for Optics and Photonics.
250. MathWorks, M.A.T.L.A.B., 1992. the Mathworks. *Inc., Natick, MA*.
251. Silverstein, D.A. and Farrell, J.E., 2001. Efficient method for paired comparison. *Journal of Electronic Imaging*, *10*(2), pp.394-399.
252. Pallud, J., 2009. The application of a phenomenological framework to assess user experience with museum technologies.
253. Li Liew, C., 2005. Online cultural heritage exhibitions: a survey of information retrieval features. *Program*, *39*(1), pp.4-24.
254. Poria, Y., Reichel, A. and Biran, A., 2006. Heritage site perceptions and motivations to visit. *Journal of Travel Research*, *44*(3), pp.318-326.
255. Thurstone, L.L., 1927. A law of comparative judgment. *Psychological review*, *34*(4), p.273.
256. National Museums Liverpool, *3D Models* [online] Available at: http://www.liverpoolmuseums.org.uk/wml/visit/floor-plans/ancient-egypt/3d-models.aspx [Accessed April 2017]
257. Unreal, 2017. *Make something Unreal with the most powerful creation engine,* [online] Available at: https://www.unrealengine.com/en-US/what-is-unreal-engine-4 [Accessed March 2017]
258. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M. and Szeliski, R., 2011. Building Rome in a day. *Communications of the ACM*, *54*(10), pp.105-112.
259. De Reu, J., Plets, G., Verhoeven, G., De Smedt, P., Bats, M., Cherretté, B., De Maeyer, W., Deconynck, J., Herremans, D., Laloo, P. and Van Meirvenne, M., 2013. Towards a three-dimensional cost-effective registration of the archaeological heritage. *Journal of Archaeological Science*, *40*(2), pp.1108-1121.

260. Roncella, R., Re, C. and Forlani, G., 2011. Performance evaluation of a structure and motion strategy in architecture and cultural heritage. *Int. Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, *38*(5/W16), pp.285-292.

Appendices

A Visitor statistics generated by NML

National Museum Liverpool Visitor Profiling Survey 2016 – 17 Quarter 3 Report 2016/17

Visitor Total
NML -3,033,228
World Museum – 670,826

Visitors for November 2016:
NML – 192,945
World Museum – 38,295

All respondents were asked for either their postcode or home town. This was then used to code responses up into the categories below. (Note: the annual report is used to code each venue's visitors by more detailed geographies – in general, the smaller quarterly samples having lower reliability for such detail).

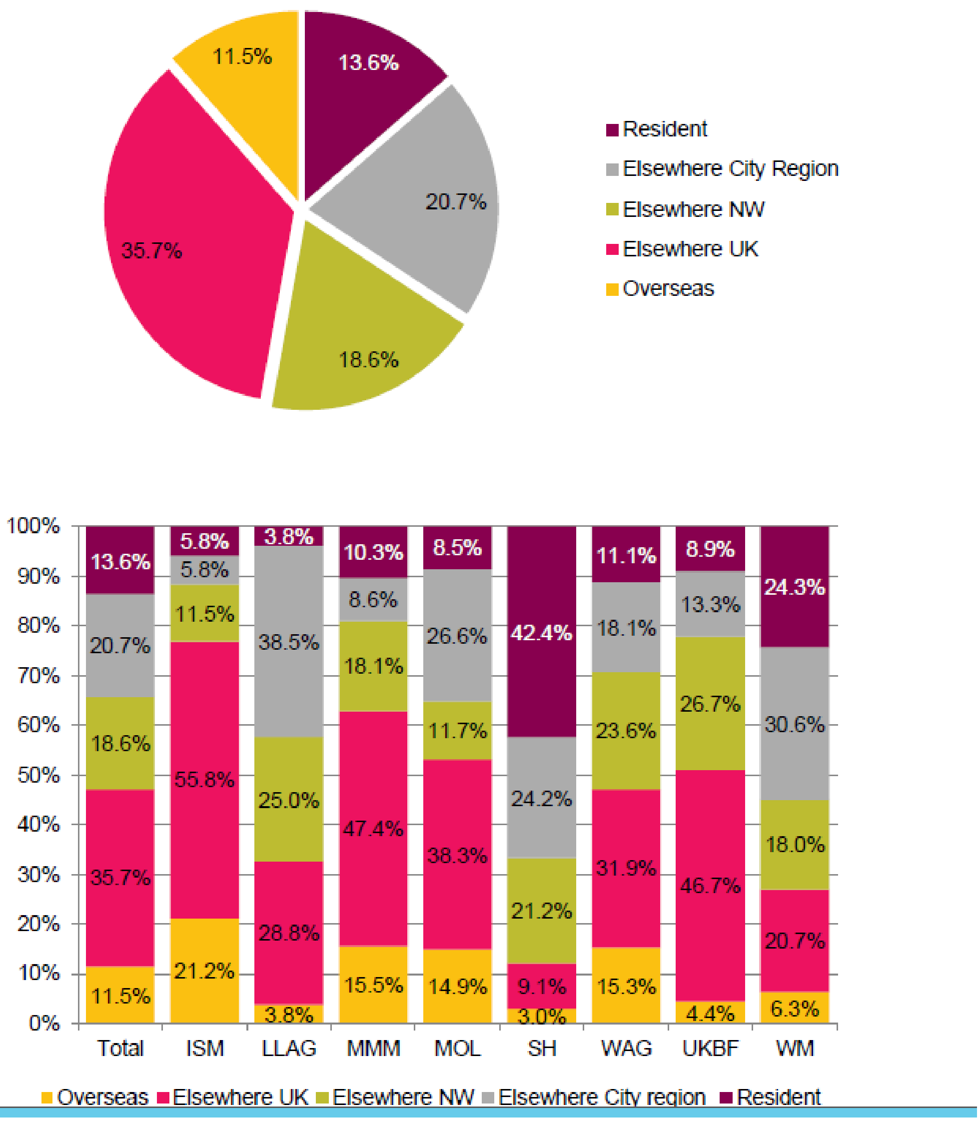

Figure 35: Visitor Origins

2.2: TYPE OF VISITOR

Respondents were asked what the nature of their trip to Liverpool was.

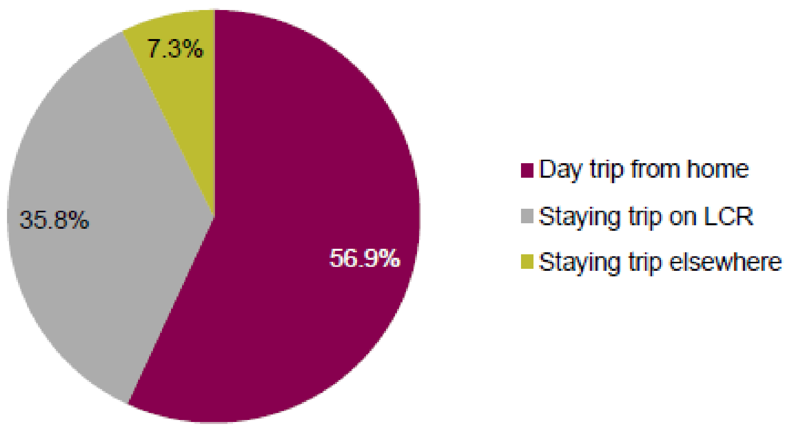

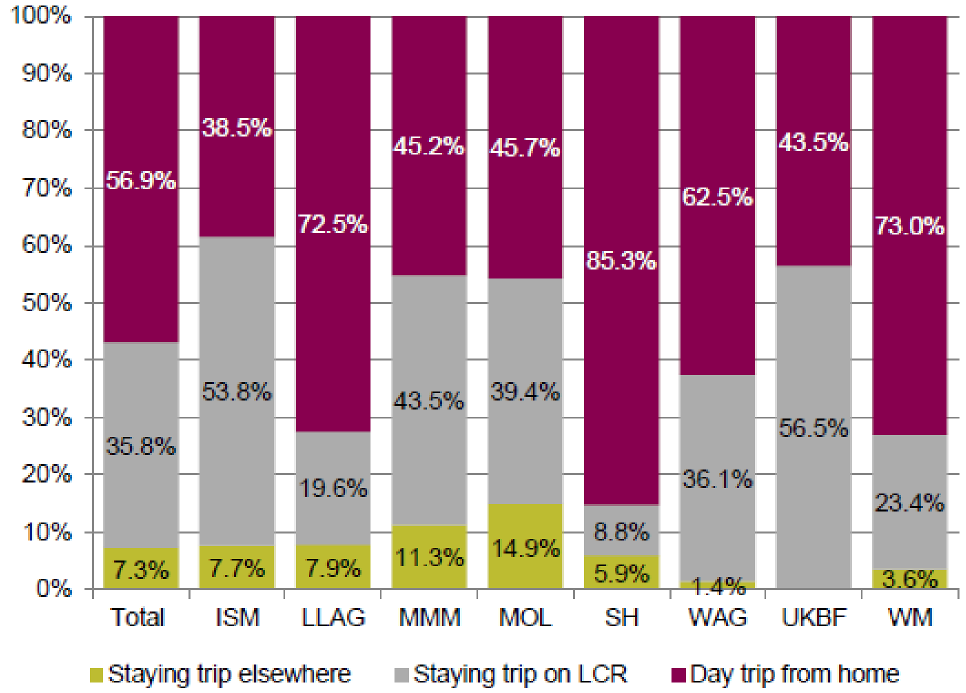

Figure 36: Visitors were asked about the nature of their trip

2.3: RESPONDENT DEMOGRAPHICS

Within this section we display the age and gender of those who were surveyed; although this is important from one perspective, as this sho those who were giving their opinions, in order to gain an understandin who the entire audience was at a venue, it is necessary to see the demographics of each *party* visiting – see section 3.

Both gender and age are broadly comparable with last year's results.

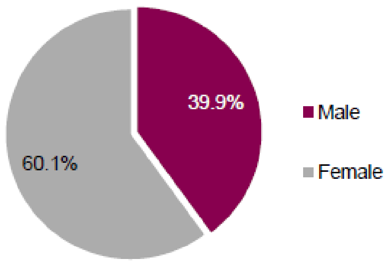

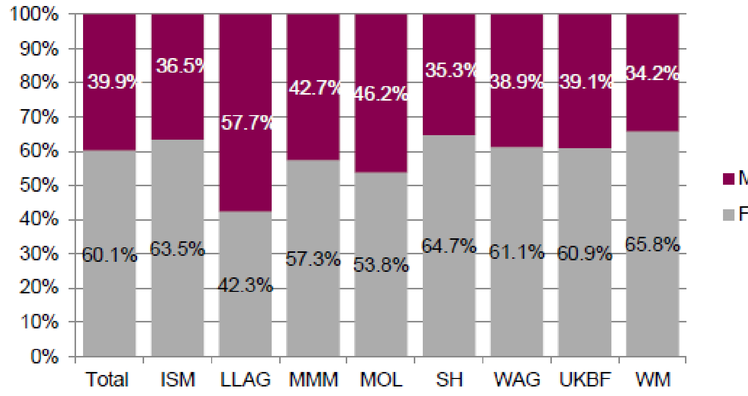

Figure 37: Visitor Demographics

	Total	ISM	LLAG	MMM	MOL	SH	WAG	UKBF	WM
On my own	12.6%	10.3%	17.2%	11.8%	19.0%	14.0%	25.0%	13.5%	8.5%
With family (inc children)	18.7%	11.5%	5.2%	23.3%	24.1%	8.8%	8.6%	10.6%	46.6%
With family (exc children)	10.7%	8.0%	13.8%	6.6%	5.6%	12.3%	9.4%	8.1%	4.2%
With spouse / partner	34.4%	48.3%	46.6%	39.9%	39.8%	31.6%	35.9%	45.9%	18.5%
With friends	24.0%	24.1%	24.1%	18.4%	12.0%	31.6%	20.3%	17.1%	24.3%
Coach party	2.2%	1.1%	0.0%	1.3%	0.5%	1.8%	0.8%	2.7%	0.0%
School / college group	0.9%	3.4%	0.0%	0.9%	0.9%	1.8%	1.6%	3.6%	0.5%
Club / community / organised group	0.3%	0.0%	1.7%	0.0%	8.5%	1.8%	0.0%	0.0%	0.0%

.

Figure 38: Visitor Group profile

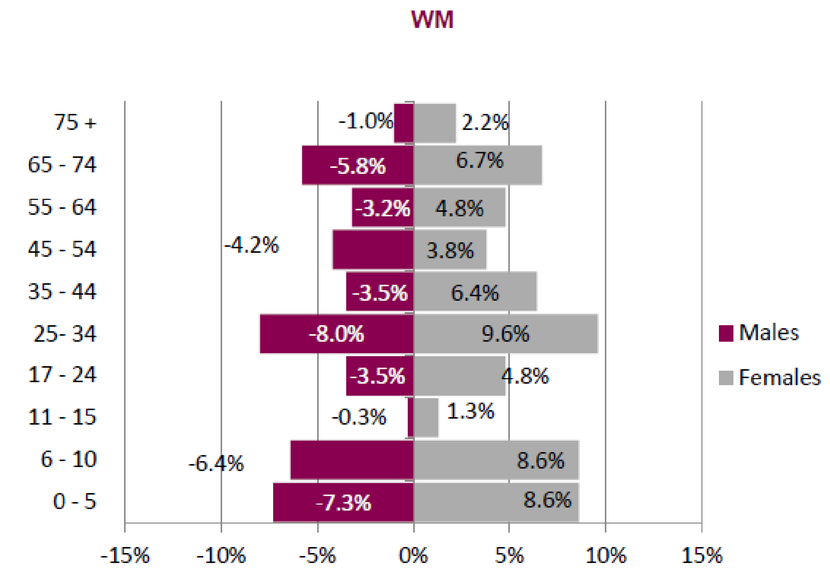

Figure 39: Visitor Age group statistics

3.3: MODE OF TRANSPORT

Overall, travelling by their own car or travelling on foot accounts for over a quarter of all visits, with almost a quarter travelling by train. Travel by car accounted for a far higher proportion of visits to Lady Lever Art Gallery (88%) and Sudley House (56%).

Respondents were asked to detail the main mode of transport they used to reach each venue on the day of their visit.

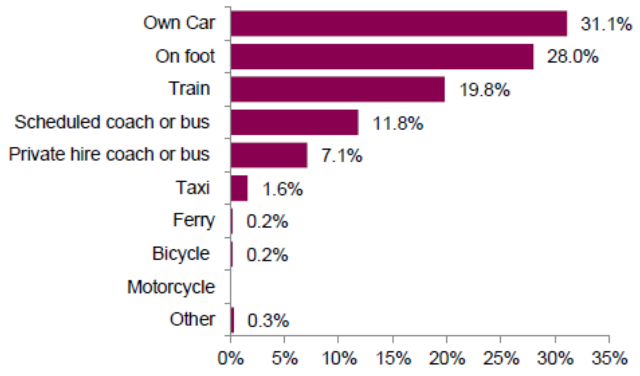

	Total	ISM	LLAG	MMM	MOL	SH	WAG	UKBF	WM
Car	31.1%	21.2%	88.2%	17.1%	15.1%	55.8%	12.5%	39.1%	38.7%
Taxi	1.6%	0.0%	0.0%	0.0%	4.3%	0.0%	1.4%	2.2%	2.7%
Scheduled bus / coach	11.8%	9.6%	2.0%	11.1%	11.8%	14.7%	9.7%	13.0%	18.0%
Private hire bus / coach	7.1%	7.7%	0.0%	19.7%	5.4%	0.0%	11.1%	0.0%	0.9%
Train	19.8%	13.5%	3.9%	17.1%	26.9%	11.8%	33.3%	13.0%	23.4%
Ferry	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.9%
Bicycle	0.2%	0.0%	0.0%	0.0%	0.0%	2.9%	0.0%	0.0%	0.0%
On foot	28.0%	46.2%	5.9%	34.2%	36.6%	14.7%	31.9%	32.6%	15.3%
Other	0.3%	1.9%	0.0%	0.9%	0.0%	0.0%	0.0%	0.0%	0.0%

Figure 40: Statistics on how visitors travelled to venues

B Example of icons used in the paper prototype and the prototype generated for the experiments in chapter 3

These were some of the Icons used in the creation of the low level paper prototype from chapter 3. It shows icons to represent the 3D object, and buttons to rotate and control the object. Other Icons were included to represent potential tools that may be of use in the interactive viewer.

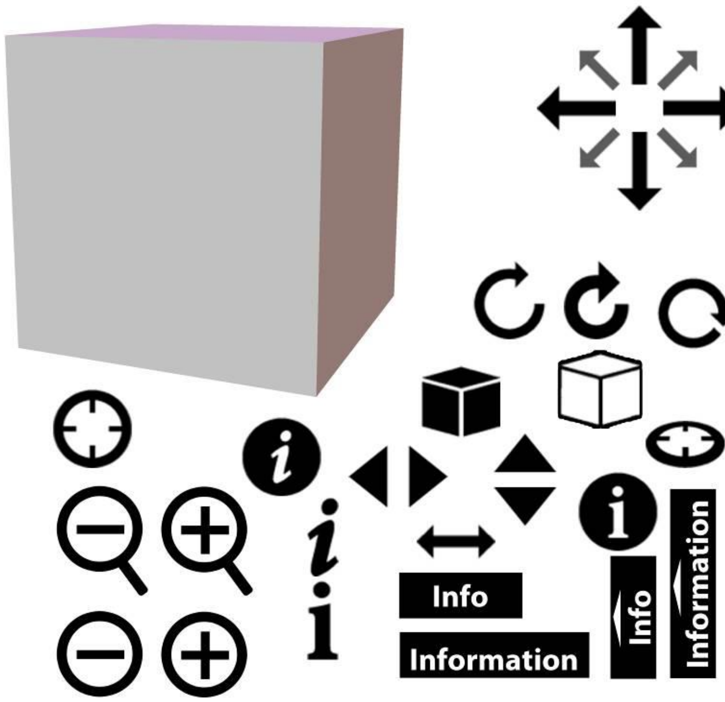

Sed ut perspiciatis unde omnis iste natus error sit voluptatem accusantium doloremque laudantium, totam rem aperiam, eaque ipsa quae ab illo inventore veritatis et quasi architecto beatae vitae dicta sunt explicabo. Nemo enim ipsam voluptatem quia voluptas sit aspernatur aut odit aut fugit, sed quia consequuntur magni dolores eos qui ratione voluptatem sequi nesciunt. Neque porro quisquam est, qui dolorem ipsum quia dolor sit amet, consectetur, adipisci velit, sed quia non numquam eius modi tempora incidunt ut labore et dolore magnam aliquam quaerat voluptatem. Ut enim ad minima veniam, quis nostrum exercitationem ullam corporis suscipit laboriosam, nisi ut aliquid ex ea commodi consequatur? Quis autem vel eum iure reprehenderit qui in ea voluptate velit esse quam nihil molestiae consequatur, vel illum qui dolorem eum fugiat quo voluptas nulla par-

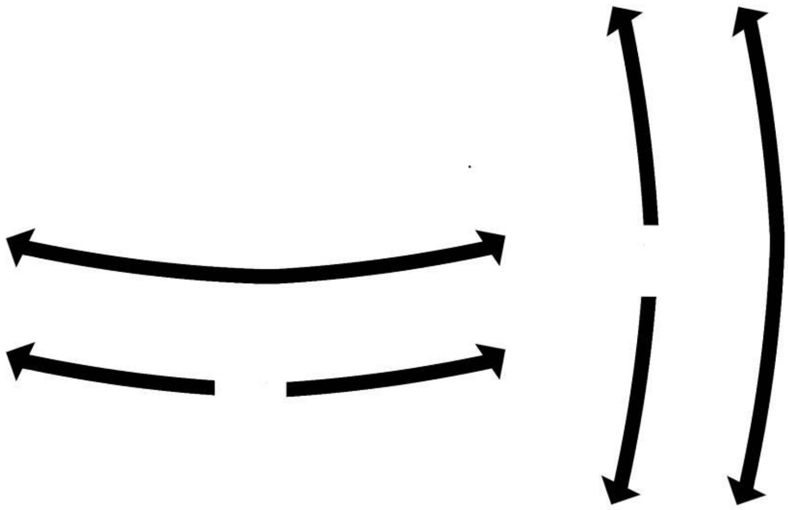

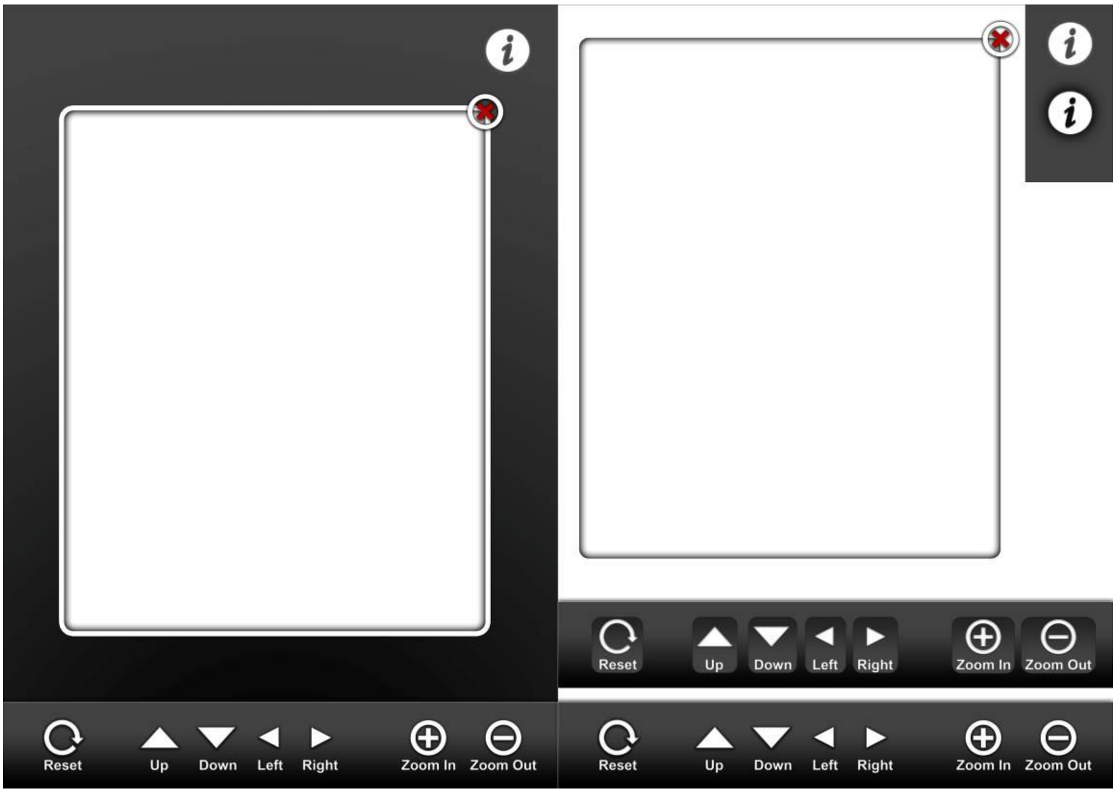
Reset
Up
Down
Left
Right
Zoom In
Zoom Out
Reset
Up
Down
Left
Right
Zoom In
Zoom Out
Reset
Up
Down
Left
Right
Zoom In
Zoom Out

C Information Sheet provided to participants for the usability study

Information Sheet for evaluation and testing of a 3D prototype interactive viewer National Museums Liverpool

Name of project:

Usability Testing for Interacting with 3D digital cultural heritage

You are invited to participate as a subject in this study investigating the preferred interaction style with 3D cultural heritage within a webpage and a mobile device, and the evaluation of the interactive prototype itself. The study itself should take between 30 minutes to an hour to complete.

You will be asked to first familiarise yourself with the interactive viewer, with the help of the researcher next to you before being asked to carry out 4 tasks. This study is not about testing how good you are, but to evaluate the preferred method to interact with 3D digital cultural heritage and if the viewer is easy to use. You do not have to undertake or complete any of the tasks asked of you if you do not wish too. You may also stop at any point.

The researcher will observe and make notes during your use and completion of the tasks, to be used for further analysis. You will also be asked to fill out a questionnaire at the end of the session regarding your experience and evaluating the interactive viewer.

The results of this study maybe published, but the data and information you provide will be completely confidential and anonymous. To assure anonymity and confidentiality, participants will not be identifiable in the results or publications of this data. An ID will be generated only so that the data can be withdrawn if you ask us too.

You may withdraw your participation at any time, including the withdrawal of any information you have provided. However, by signing the consent form attached, it is understood that you have consented to participate in this experiment and to publication of the results, with the understanding that anonymity will be preserved. The project is being carried out by:

Name of researcher: David Gillespie

Email Address: david.gillespie@liverpoolmuseums.org.uk

He will be pleased to discuss any concerns you have about participation in the project.

D Scenarios for usability study concerning preferred interaction style with 3D cultural content

This was the scenarios and instructions given to the participants in the interaction study in chapter 3. The instructions include the information on what is expected from the participant and the areas to navigate too.

Usability Testing for Interacting with 3D digital cultural heritage

Instructions for the participants:

This study will involve completing four scenarios, using different interaction styles to engage with a 3D cultural heritage model. Three of the scenarios will be conducted on a laptop and 1 will be on a mobile device. You will be asked to navigate to points on the 3D model as shown below for each model. When you are done, you will be asked to complete a questionnaire and evaluate the interactive viewer.

This task is to be completed following the think out method, where you will need to vocalise your thoughts for any actions you are performing, any difficulties, confusion or any questions you would like to ask. If you are silent and forget to vocalise your thoughts, you will be given a reminder.

This is a study focusing on the usability of this interactive prototype and engaging with 3D cultural content. The observer cannot answer questions regarding the interaction or the specifics of the controls.

This session will take approximately 30 minutes to an hour to complete.

Please navigate to these points on the 3D model for scenario 1

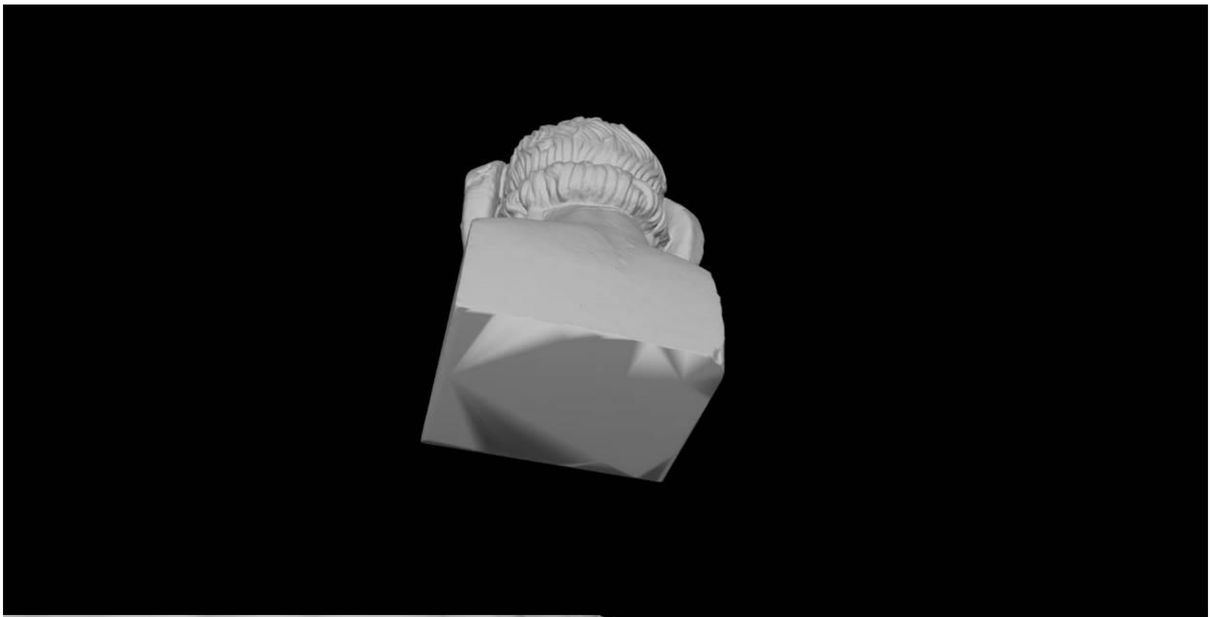

Please navigate to these points on the 3D model for scenario 2

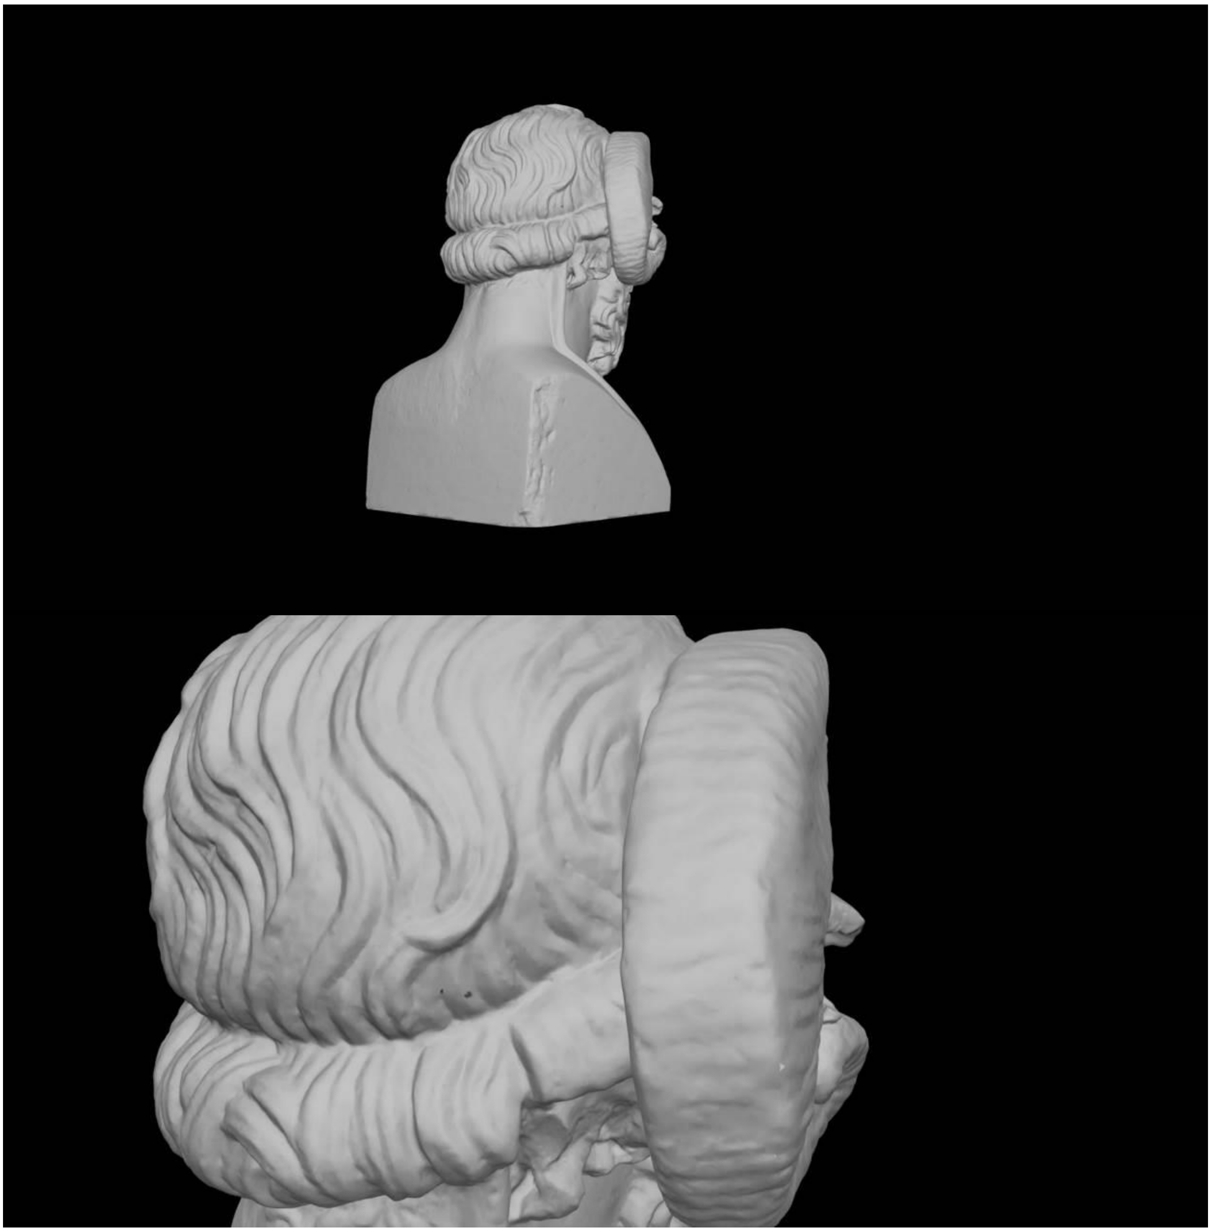

Please navigate to these points on the 3D model for scenario 3

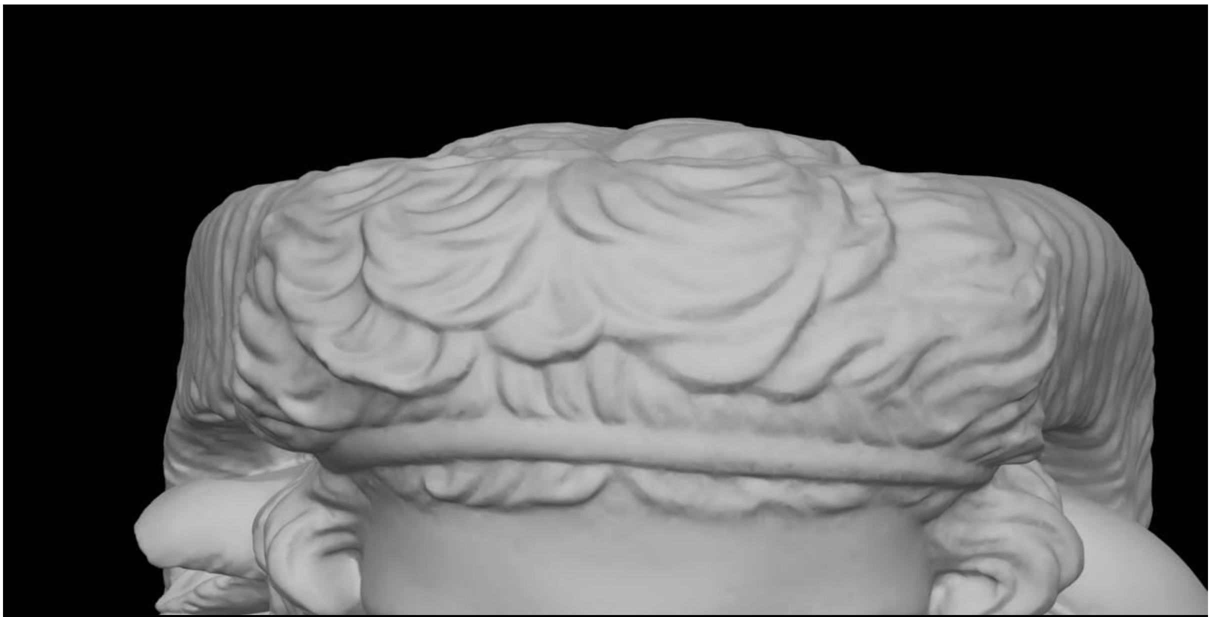

Please navigate to these points on the 3D model for scenario 4

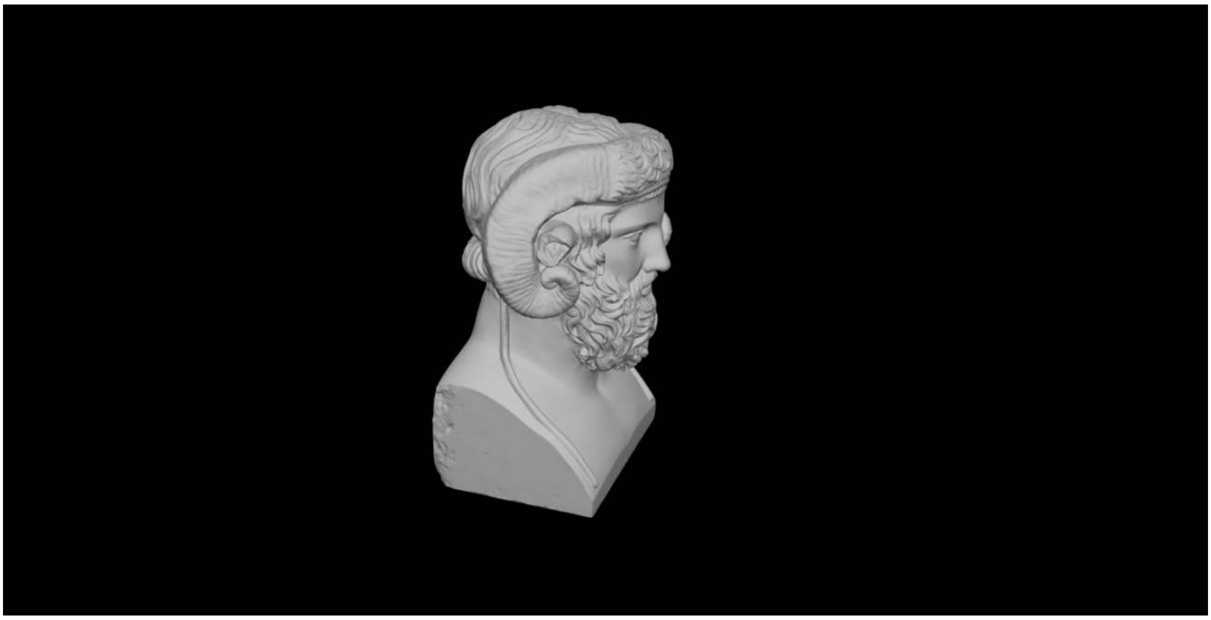

E Feedback generated from the experiments conducted in chapter 3

The experiment conducted in chapter 3 asked for feedback on 4 different interactive styles with the generated prototype. The styles included normal and an inverted y axis, changing where the centre of the screen rotates and the use of it on a mobile device. The feedback generated asked the users to evaluate the overall design of the interactive prototype to measure how well it was perceived. Eight participants took part in this experiment, 2 male, 6 females aged between the ages of 22 – 38 with normal or corrected vision took part in the user testing. Participants ranged from curators, conservators and volunteers within NML. The participants had various degrees of experiences with 3D graphic applications with most having very little to no experience. Their feedback is presented below.

Participant A

Age: 22
Gender: female
Experience with 3D: not a lot

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Commented that it was easy to navigate – most controls as expected i.e zoom with the mouse scroll, rotate by left click movement
- They found the rotation overly sensitive at first but once used they got used to the sensation and it becomes easier to control
- They didn't seem to understand the icon for panning, and accidentally discovered the right mouse button to pan.
- They thought this made sense though (left click to rotate, right click to move object around).
- Commented the panning makes sense, not overly sensitive – doesn't move too fast or slow
- They commented that they were confused why there were holes in the 3D model
- Found the background to be a bit clinical
- Thought it was user friendly for all experience levels
- They tried the icons, but they preferred to use the mouse for interacting
- They thought they could zoom in a little too close before the model disappeared behind the camera

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- They commented the found the controls easy to use/understand again
- Thought the rotation still a bit sensitive but easy to use once accustomed to it
- Found the inversion to be confusing, pulling the mouse down rotated the model the opposite way from what they expected and vice versa.
- Commented that it would be good to see the object holes filled on the object rather than a blank space

- Found the back ground to be clinical.
- Again didn't use the icons

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Commented it was simple and similar to the other interactions
- Commented that moving and rotating the object are simple when the object is in its original position.
- However, they commented that once the object moved to outside of the centre, rotation feels uncomfortable. They described it as viewing an object in person you would be face on to examine it with the object directly in your eye line instead of in your peripheral vision. So why wouldn't you expect something similar when viewing it interactively
- Commented that they would like a reset button.

Observation on Interaction style 4 (interaction on mobile device)

- Found it easy to use.
- Commented they could use this easily as they are used to using a smart phone.
- Had no troubles with gestures or that
- Used the icons more, though this could have been due to the lack of gestures to pan a model.
- Commented the icons look quite good on the mobile device.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

Yes, but it took a bit of time to find how to navigate.

Is there any confusion over any of the icons or controls?

Was a bit confused by the pan icons, and didn't understand why it was there. Otherwise it was ok, didn't really use them though when I started to use the mouse to move the object. Didn't like it when I clicked the info icon, it floated over the object, could have done without it, or had it on the side.

Does the interactive viewer work like you thought it would?

Discovering the navigation took a bit of time but it worked pretty well such as the zoom and panning the object.

Did the interactive viewer allow you to explore the object as much as you wanted to?

I found it to be a bit clinical, but I could get to all of the points quite easily. Would have liked a few more tools. Would have liked to be able to change the lighting too. Guessing you would need more direction needed for the average user though? Maybe restrict the amount you can zoom in, was a bit disconcerting when it suddenly disappeared when you got too close and couldn't find it again.

Would you like to interact with 3D cultural heritage models using this?

Yeah

Do you have any other comments?

Information of the model could easily have been put on the side.

Participant B

Age: 27
Gender: female
Experience with 3D: not a lot

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Seemed to be able to use the interactive intuitively without any instructions
- Mentioned it was simple and easy to use, liked the mouse controls and it was how they expected it to work
- They mentioned they thought the icons and interface were designed quite nicely
- They commented its good for an accessibility standpoint, but they would probably use the mouse controls
- They were able to navigate easily with the controls
- They didn't use the icons to control the 3D model , except to test their functuality
- Liked the options to zoom and pan, mentioned if there would be any other tools to see the surface in more detail at all?
- Thought the icon for panning was a bit confusing, especially with the arrow controls, which they thought was to move the 3D model.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- They once again were able to use it intuitively without instructions
- Mentioned it was very similar to the other interaction, but they could easily use it with the inverted axis
- Completed the task quite easily.
- Mentioned again that it was easy to use,
- Commented that the interface seemed a bit small for a webpage, maybe an option to go full screen would be nice.

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Commented that it was easy to use, very similar to the other interactions
- Thought that moving the object and rotating around it in the centre was good,
- They commented they didn't like it when it went off centre; it wasn't how they thought it would be when interacting with the 3D model.
- Was able to complete the tasks, but mentioned again they didn't like the pan and rotate interaction. They thought it was off putting and would like something like a reset button to make it slightly better.

Observation on Interaction style 4 (interaction on mobile device)

- They found it intuitively to use
- Commented that the design looked better on the phone, than the small app within the webpage.
- Found the use of gestures easy to use on a smartphone, liked that the gesture to zoom, was the same to zoom into the 3D model
- Completed tasks easily with minimal use of the icons to control the 3D object.

- They thought the use of information on the phone app would need to be redesigned. It covered the model and didn't seem designed for the phone.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

Yes.

Is there any confusion over any of the icons or controls?

No

Does the interactive viewer work like you thought it would?

yes

Did the interactive viewer allow you to explore the object as much as you wanted to?

Yes.

Would you like to interact with 3D cultural heritage models using this?

If it is related to the collections on the gallery then yes.

Do you have any other comments?

No

Participant C

Age: 34
Gender: female
Experience with 3D: not a lot

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Commented that the design of the GUI looked nice
- Mentioned they liked the 3D model a lot.
- Commented that it was nice to be able to interact with such a large object within the viewer, mentioning how hard it would be to interact with something like this
- They seemed to have some difficulty getting used to the controls at first mentioning a lack of experience with 3D models.
- Was able to navigate to points on the 3D model, using the icons and controls
- Mentioned the controls were easy to use after having a couple of minutes to use it.
- Thought it was quite easy to use, and user friendly for people with a low level of experience with 3D models
- Used a mixture of the icons and mouse to navigate.
- Liked being able to zoom in and out but asked about other tools such as a raking light.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- They mentioned that it was very similar to the other interaction
- Took a couple of seconds to get used to the inverted axis, commented that they would easily get accustomed either of the interaction styles, depending on which was presented first.
- Mentioned about sensitivity of the model when zoomed out, thought it moved too fast to be practical.
- Though the interaction was quite easy to use, and the icons helped out a lot
- Used a mixture of the icons and mouse to move about

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Commented it was similar to the other interactions rotation and zoom in
- Was able to navigate around the object quite well by this stage
- Panned the model and then rotated and seemed putting off by the interaction, especially navigating to points on the 3D model.
- They commented that they found it a strange way to interact with an object; It is not how they would move the object in the gallery or within conservation.
- They mentioned it is not how they would like visitors to interact with the object.
- They used a combination of icons and mouse to move around the object.

Observation on Interaction style 4 (interaction on mobile device)

- They mentioned that it looked like an app on the mobile and seemed to fit there
- They mentioned they found the interaction very easy to use on the mobile device
- Things were clear and laid out, making more sense on the small device.

- They commented they would like to see something like on an ipad or something for the galleries.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

It was ok, I struggled to get used to it to begin with, but it was easy to use after a couple of tries.

Is there any confusion over any of the icons or controls?

Apart from the pan button choice, not really. I would have preferred it if the information for the object was separate from the interaction. It distracts from the interaction.

Does the interactive viewer work like you thought it would?

I don't know how it was supposed to work. It however was easy to use after getting used to it.

Did the interactive viewer allow you to explore the object as much as you wanted to?

Yes, even with the slight difficulty using the controls to begin with.

Would you like to interact with 3D cultural heritage models using this?

Yes ,I think it would be great as an app on the iPhone or iPad for visitors to use on the museum galleries.

Do you have any other comments?

Participant D

Age: 28
Gender: male
Experience with 3D: some from computer games

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Commented that it was easy to use
- Commented that it look good
- Didn’t seem to struggle to navigate to points on the object
- Used the icons when they had zoomed in closer
- Understood the pan and moved between that and rotate a lot.
- Mentioned it was a bit like Skyrim, when interacting with an object.
- Commented it was a bit stuttery and felt like the frame rate dropped a bit
- Seemed to be able to use this quite easily.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- Found the interaction similar to the first interaction
- Commented that it was more like a flight game mode for looking at things with the inverted axis.
- They also didn’t seem to enjoy it that much
- However, they mentioned they could get used to it.
- Once again mentioned the dropped frame rate at points.

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Commented on frame rate and how it was easy to use
- Worked as he was expecting it with respect to rotation and zoom
- Panning he seemed to find it a bit difficult when panning then rotating. He would always drag it back to the centre of the screen.
- Commented it’s probably not the best way to interact with one object
- Suggested maybe using a first person like style if there were more objects in the viewer.

Observation on Interaction style 4 (interaction on mobile device)

- Commented that it was easy to use,
- Used the icons a lot more
- Commented that it ok, and not mind blowing for him.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

Yes

Is there any confusion over any of the icons or controls?

No

Does the interactive viewer work like you thought it would?

Yes

Did the interactive viewer allow you to explore the object as much as you wanted to?

Yes

Would you like to interact with 3D cultural heritage models using this?

Yeah, seems pretty good, just need to watch the FPS

Do you have any other comments?

no

Participant E

Age: 38
Gender: female
Experience with 3D: moderate

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Commented It didn't take long to get used
- Mentioned it was easier after they got going
- Liked the design, thought it was well laid out
- Mentioned it worked well, and would be great to interact with other 3D models
- Used the icons as much as the mouse, even after discovering what the mouse buttons did
- They commented they liked being able to zoom into the model and being able to see the model in detail.
- Managed to navigate and complete the task
- They seemed to be enjoying using it when interacting with the 3D model.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- Thought it was similar to the original interactive
- They were confused with the inverted axis but mentioned that they could get used to it, and that overall it was very similar to other interactive
- Mentioned it didn't feel as responsive as the other interactive

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Mentioned the responsiveness again, could be due to frames being dropped
- Found it similar to the other interactives for interacting with the 3D model
- Commented that they found it frustrating with the panning and zooming, finding it to be extra work to complete a simple navigation.
- They found it not to be natural to use with the panning as it was.

Feedback on Interaction style 4 (interaction on mobile device)

- They mentioned they found it good to use,
- However, they thought the screen was too small to see the details clearly
- Thought the interface was a bit too cramped with the additional icons
- They mentioned it was responsive, and easy to use.
- Expressed an interest that it would be ideal to be used on the galleries within the National Museums Liverpool.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

My experience so far has been good; I mainly use similar software for research and work. This seems targeted more towards people without a lot of interaction with 3D models or software.

Is there any confusion over any of the icons or controls?

Not really, it provides the basic controls, and I understood the icons easily. The text underneath also highlighted what it was meant to do.

Does the interactive viewer work like you thought it would?

Yes. I believe it’s a bit low level for the likes of research, but if you provide a couple more tools. It could be more impressive. Annotating and pointing out places of interest would be good but also being able to change the 3D models shader and how it’s viewed would be good.

Did the interactive viewer allow you to explore the object as much as you wanted to?

I believe it is a bit basic to fully explore the object. However, if you included something like lighting, animation or annotations it could be a lot better.

Would you like to interact with 3D cultural heritage models using this?

I can see this being good for the general public, and would be fantastic on galleries with a few changes. However I can’t see this being so great for research purposes.

Do you have any other comments?

No

Participant F

Age: 32
Gender: female
Experience with 3D: not a lot

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Commented that the icons looked clear and concise
- They seemd to take a couple of seconds to get used to controls,
- They used the icons to look around the object and mixed it with the mouse
- Never seemed to use any of the mouse buttons apart from the left one (right controlled the pan, scroll to zoom)
- Mentioned how they liked to be able to interact with it easily.
- Was confused by the pan button, had to ask what it was.
- Was able to navigate to the points after getting used to the system.
- They said the interaction seemed to be either really responsive or jittery depending on how close they zoomed in.
- Commented that there should be a limit on the amount they could zoom in.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- Made the same comments about interface looking clear and it explained the functions quite well
- Was a bit confused by the inverted axis but got used to it shortly afterwards
- Mentioned about the inverted axis being a bit strange, but was used to it after a bit.
- Style used a mixture of the mouse and icons to navigate within the viewer.

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Mentioned again, about how the initial interaction was very similar to the previous ones
- They commented that they didn’t notice a difference in the interaction for rotating and panning
- They were confused when they rotated after panning the model
- They commented this was not how they were expecting it to work, they expected the screen to rotate round the model, not the model rotating off screen.
- They would drag the pan to the centre to avoid this interaction
- Commented they were a bit uncomfortable with that interaction

Observation on Interaction style 4 (interaction on mobile device)

- They commented it looked clear on the mobile device, very similar to an iphone app/design
- Found it to be responsive, though could be sluggish at parts.
- Mentioned it was easy to use, and the controls helped out on the smaller device.
- Mentioned displaying the information on the phone would need to be reworked as it blocked the object.

- Found gesturing to be a bit unresponsive for them.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

I think it is, a few things here and there were a bit confusing but overall was quite good to use.

Is there any confusion over any of the icons or controls?

Well, there was some for when I tried to move the model, it wasn't clear that when I pressed the pan button, the arrow buttons would move the object instead of panning it. This seemed to lock the mouse left button to move the object too, would be nice if you could change between them a bit more naturally. Also the information button when it was pressed obscured the object. Could that at be placed outside the viewer at all, or at the side of the webpage.

Does the interactive viewer work like you thought it would?

It took a bit of time to get used to but overall I enjoyed using it, and found it easy to use. I could easily see a few changes being made here and there to make it a bit more usable for the general pulic.

Did the interactive viewer allow you to explore the object as much as you wanted to?

Yes, but I would have liked more tools to engage with it more, changing the lighting, or being able to mark areas of interest and allowing for annotations would be great too! Will there be tools for this when it's placed within the website?

Would you like to interact with 3D cultural heritage models using this?

Yes I can see the potential here to let visitors interact with this. Would allow them to see the 3D objects from new angles, and it would be extremely useful for research as well! Additional tools would be great!

Do you have any other comments?

Would be better if you could make it go full screen so you can see it in more detail. Seems a bit constrained. Would like more tools if possible to inspect and interact with the model, maybe even more information regarding the stages of the creation of the 3D model and how it was made?

Participants G

Age: 22
Gender: male
Experience with 3D: a lot

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Found it clear and well laid out.
- Commented it was pretty responsive and intuitive to use, similar to some 3D packages they had used.
- Commented it was easy for them to use, and does what they thought it would do.
- Found the button set up on the mouse, for rotate, pan and zoom to be good.
- They seemed to navigate to the point on the 3D model easily.
- Didn't seem to struggle to use the viewer at all
- Found the pan button to be valuable to navigate to points on the 3D model
- Found the buttons ok to use, but they preferred interacting with the mouse controls.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- Commented that it felt very similar to the first interaction
- Commented it only took a few minutes to get used to the controls with the inverted axis.
- Mentioned they are used to and prefer it the other way though.
- Didn't seem to struggle at all, when using the 3D viewer.
- Seemed to find the interaction enjoyable

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- Commented that it was similar to the other interactions
- Commented that once again, they liked the zoom to see details and that it was responsive
- Seemed to struggle slightly when they panned the model and moved it off center
- Commented they found it a bit confusing, when it was not rotating around the centre of the screen

Observation on Interaction style 4 (interaction on mobile device)

- They seemed to find it nice and responsive.
- Commented that it was still very intuitive and easy to use.
- Seemed to struggle with zooming in, without the use of a mouse. Seemed to struggle with mobile gestures.
- They commented that the icons used in the application were distracting as they seemed a bit squished on the mobile and a bit more design would be needed.
- However, they did comment that the icons were clear and communicated their functions.
- They preferred to use their finger to interact and not use the icons much

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

Yes, found it easy to use, but could be a bit cluttered on the phone. Otherwise looked great!

Is there any confusion over any of the icons or controls?

Each icon was labelled and easy to understand. The buttons for moving the 3D model could be made a bit more responsive, but I am guessing they are there for those with a disability or can't use a website well? Also didn't realise you could pan with the arrows when it was selected.

Does the interactive viewer work like you thought it would?

It did! Controlling the 3D model wasn't hard and pretty intuitive. Fast and responsive enough to use. Works like I was using 3ds max controls wise. Lacks some of the controls of the 3D software I use but can see something like controllable lighting being added easily.

Did the interactive viewer allow you to explore the object as much as you wanted to?

It does, it allowed me to zoom in far enough, and see all the details. Could see lots of details on the object. Must have been a high poly model!

Would you like to interact with 3D cultural heritage models using this?

Yes. Would be great to see some more 3D models,

Do you have any other comments?

Would be better if you could make it go full screen so you can see it in more detail than it being constrained so much. Was a bit on the small scale.

Participant H

Age: 35
Gender: female
Experience with 3D: ok amount

Experiment 1

Observation for Interaction style 1 (normal, axis not inverted, rotate round centre of object)

- Mentioned the interaction was similar to some software they had used for research.
- Commented on the design, complimenting it clear and concise look
- They mentioned that the 3D model was a good representation of the model, which they pass in the World museum
- They seem to find it easy to interact with the viewer, working out the mouse buttons quickly.
- Liked the responsiveness, but mentioned that there seems to be a frame rate drop when you zoom in too much.
- Completed the task easily
- They didn't really use the icons to navigate, after working out the mouse buttons and not needing the icons.
- Mentioned the interaction was quite basic, needed some tools common in research like a raking light or that to bring it to life.

Observation on interaction style 2 (inverted axis, rotated round the centre of the object)

- They commented it was very similar to the other interaction
- Found that even with the inversion, they could still easily use the viewer without discomfort.
- Completed the task quite easily.
- Once again mentioned the frame rate issue
- Did not use the icons present

Observation on Interaction Style 3 (normal, axis not inverted, rotate round centre of screen)

- They found the interaction quite easy
- Discovered the pan was different, mentioning it was similar to Meshlab, which they use regularly.
- Was able to navigate easily. They did mention that it could possibly be disconcerting for people not used to this when interacting with a 3D object.

Observation on Interaction style 4 (interaction on mobile device)

- They found the interaction to be easy to use on the mobile device.
- Though it was intuitive to use and very similar to other apps they had seen on mobile devices.
- They commented that maybe some extra tools might make it more interactive on the mobile.
- They said they could see the appeal of this and trying to promote it to visitors on the website or gallery, but didn't know how they would proceed with that.

- Was able to complete the task easily, using just fingers and switching between pan and rotate with the one button.

Experiment 2:

Do you find the interactive viewer easy to understand and accessible?

Yes.

Is there any confusion over any of the icons or controls?

Not at really, the language cues under the icons easily denoted what they meant.

Does the interactive viewer work like you thought it would?

It presented the 3D model quite clearly and it was very easy to access. The first sets of controls were easy to learn and made the interaction. I don't know if it would be fit for research, but would be usable by the general public I don't know how it was supposed to work. It however was easy to use after getting used to it.

Did the interactive viewer allow you to explore the object as much as you wanted to?

Yes, it presented a clear and precise representation of the Bust; I have seen it within the World Museum and think it seems to be a close match. If we included some more relevant research regarding the object within the webpage, or other tools like a raking light it would be a lot better.

Would you like to interact with 3D cultural heritage models using this?

If it was practical and affordable to deploy to the museum website and galleries, then yes. Though with our stretched budgets, I can't see that happening anytime soon.

Do you have any other comments?

There could be other fancier ways of possibly presenting the 3D model, maybe some colours or relate it to National Museums Liverpool but I don't really know.

F Images computed for experiments conducted in chapter 5

The experiment's conducted in chapter asked users to compare generated stimuli at different polygonal and texture resolutions, and a stimulus was generated with the use of the HDR-VDP2 [217] image metric. The images rendered for the comparison experiment and used generating the stimuli in the subjective test are presented here. They are presented as the following for each artefact:

- The generated stimuli used for the comparison test
- The images rated using the HDR-VDP2 image metric against a reference image in the RGB.BT 709 colour space
- The differences computed using the HDR-VDP2 image metric matlab code in the RGB.BT 709 colour space.
- The images rated using the HDR-VDP2 image metric against a reference image in the sRGB colour space
- The differences computed using the HDR-VDP2 image metric matlab code in the sRGB colour space.

Anglo Saxon Brooch –

Rendered images of the different stimuli at different texture and polygonal resolutions

Figure 41: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 42: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 43: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 44: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 45: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 46: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 47: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 48: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 49: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 50: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 51: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 52: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 53: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 54: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 55: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Figure 56: Reference model with stimuli at 100% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

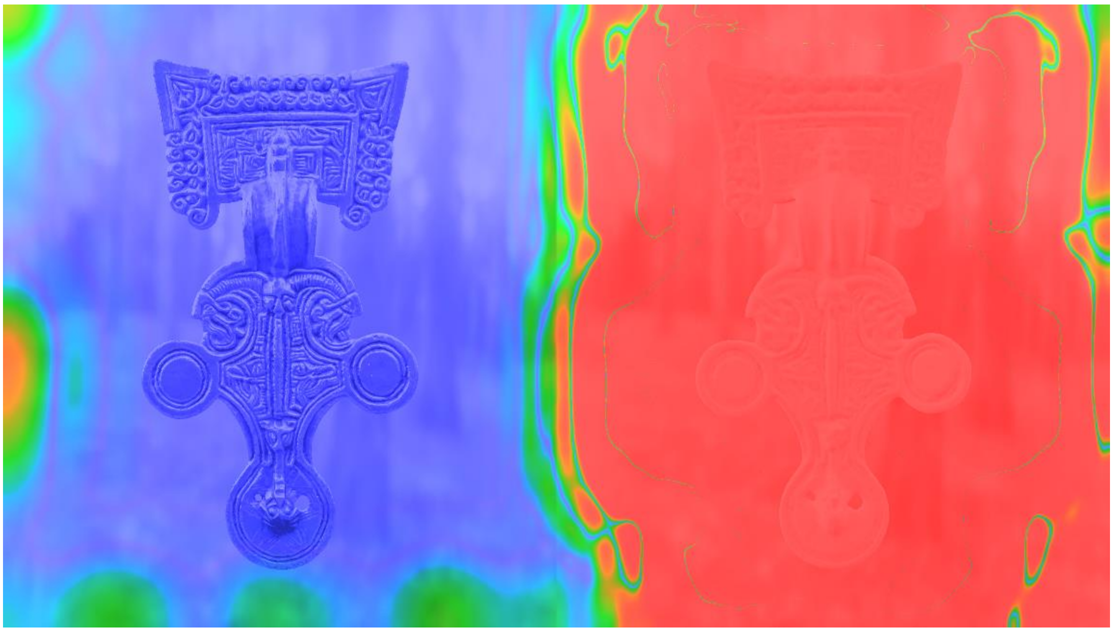

Figure 57: Reference model with stimuli at 10% polygonal resolution and no texture

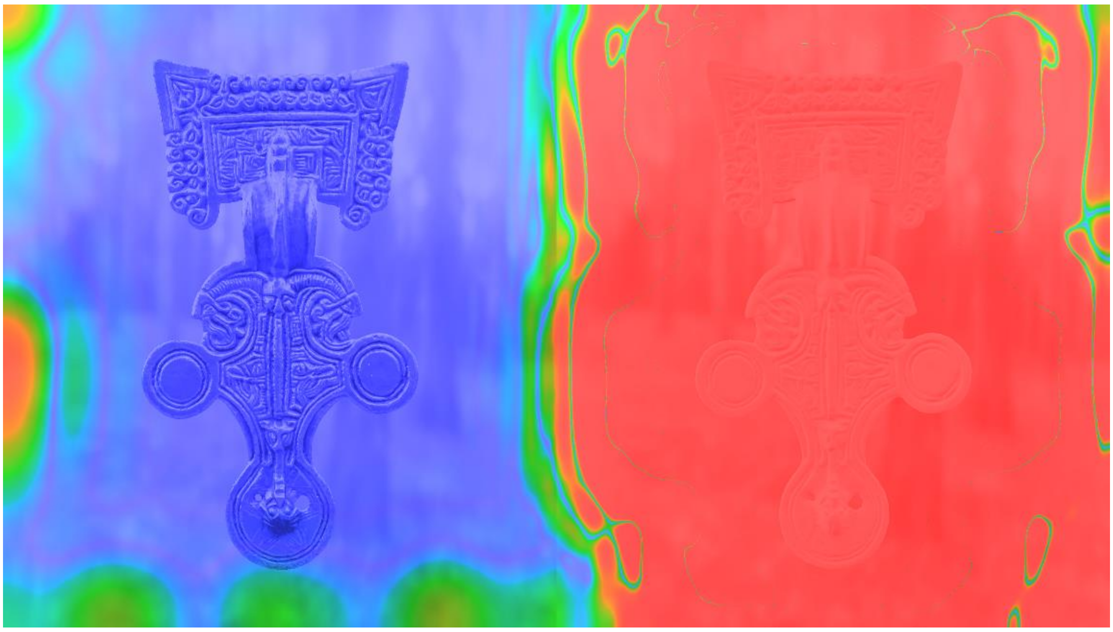

Figure 58: Reference model with stimuli at 40% polygonal resolution and no texture

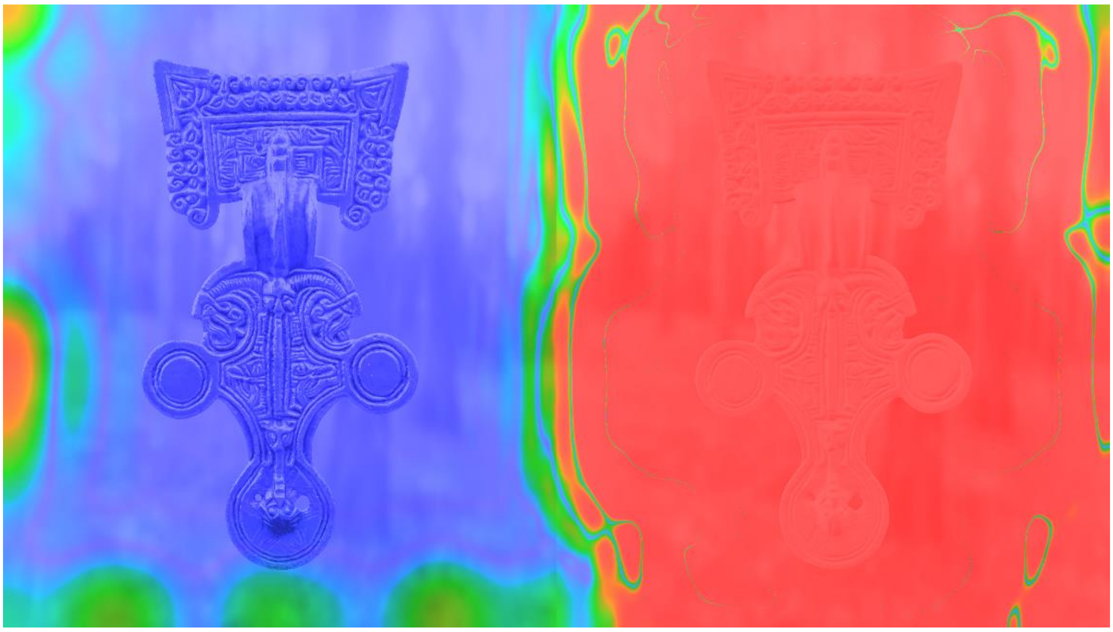

Figure 59: Reference model with stimuli at 70% polygonal resolution and no texture

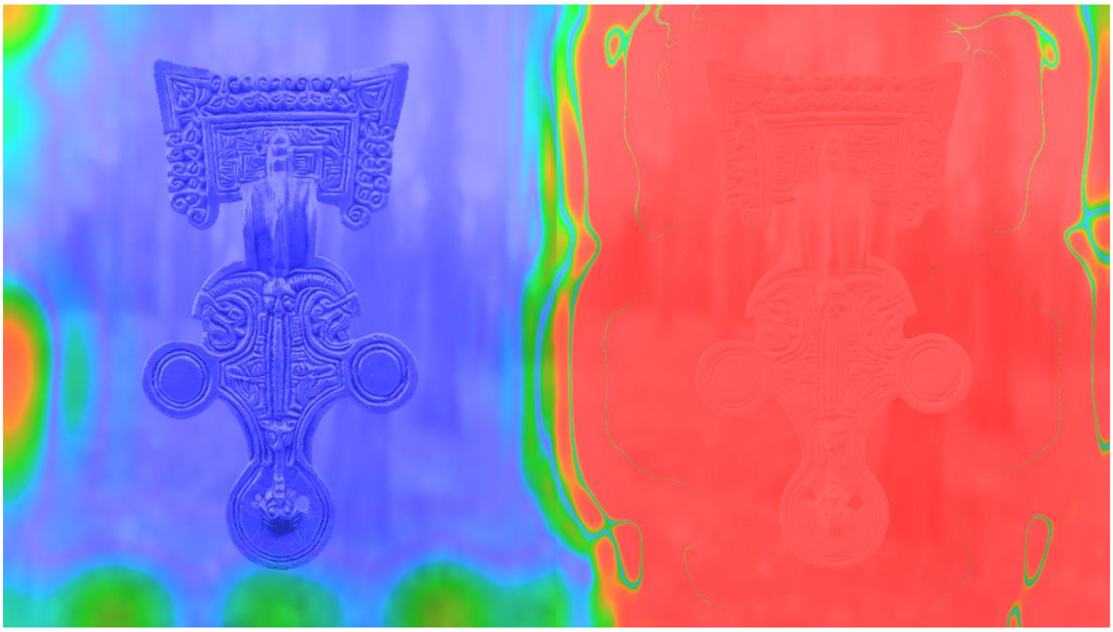

Figure 60: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 61: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 62: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 63: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 64: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 65: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 66: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 67: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

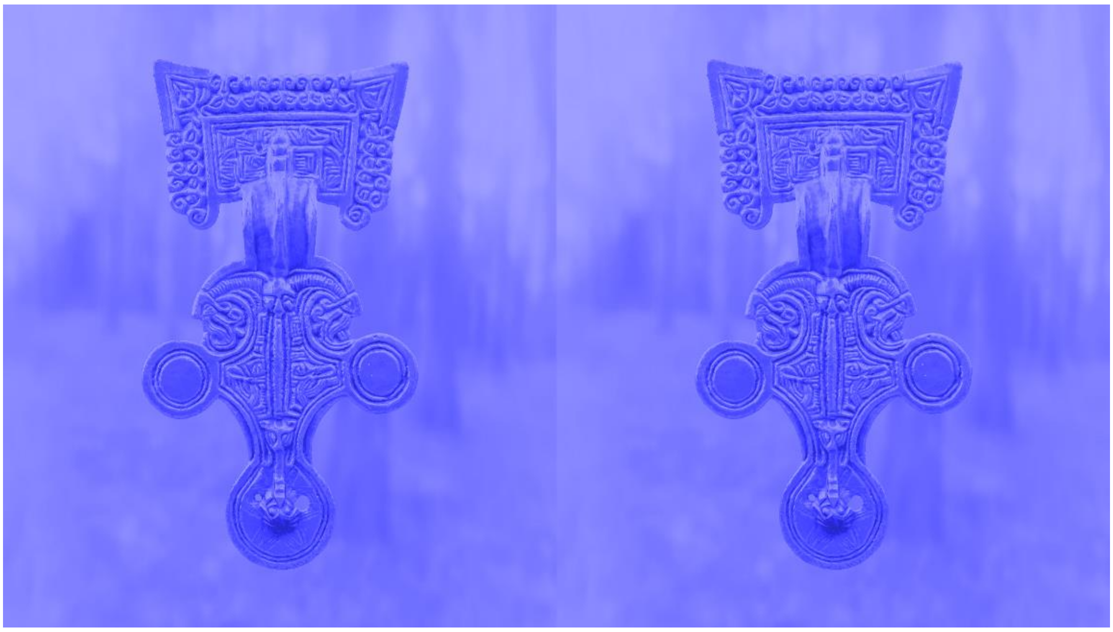

Figure 68: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 69: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 70: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

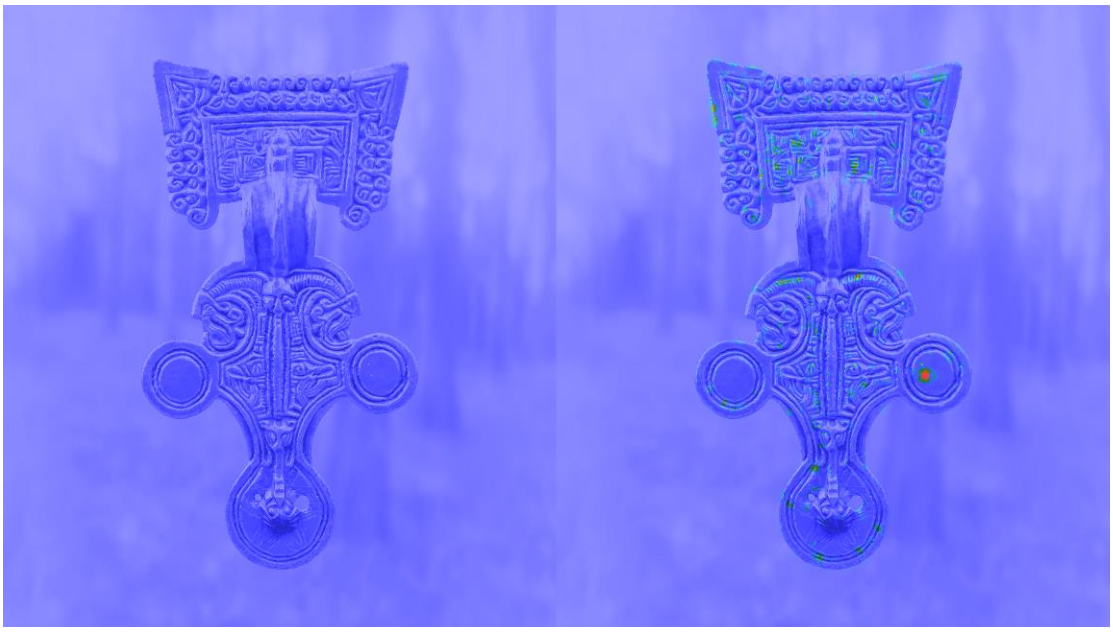

Figure 71: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space.

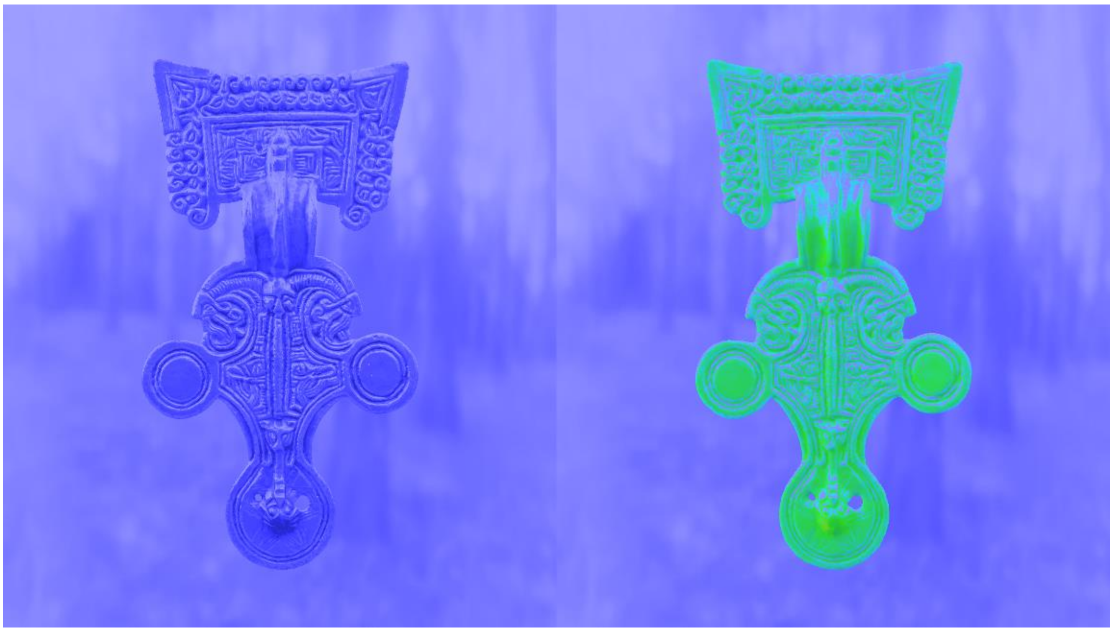

Figure 72: Reference model with stimuli at 10% polygonal resolution and no texture

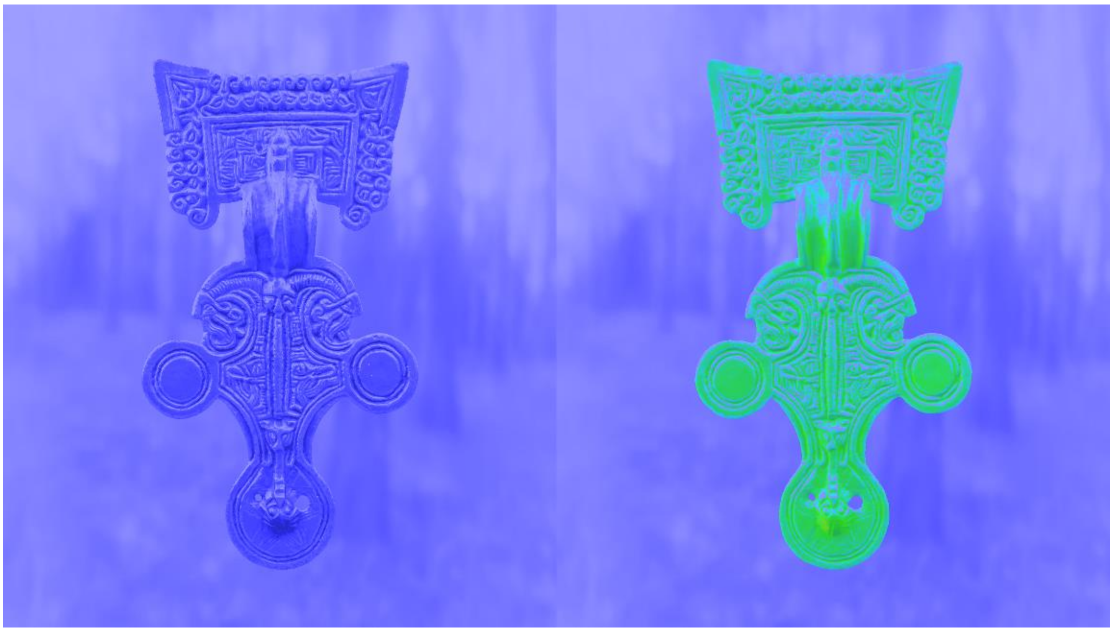

Figure 73: Reference model with stimuli at 40% polygonal resolution and no texture

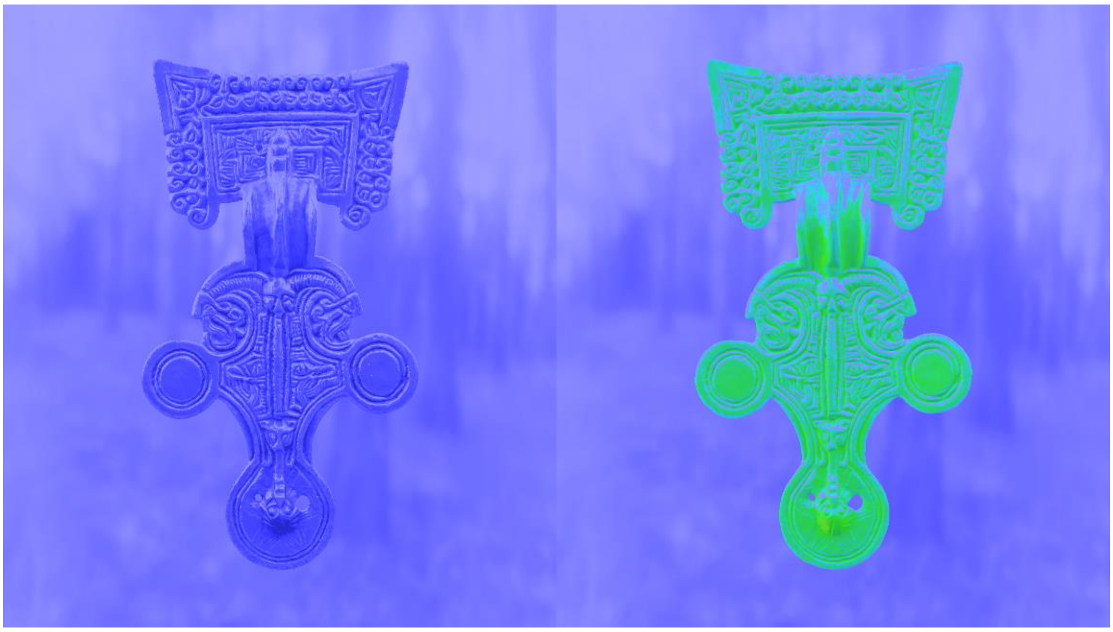

Figure 74: Reference model with stimuli at 70% polygonal resolution and no texture

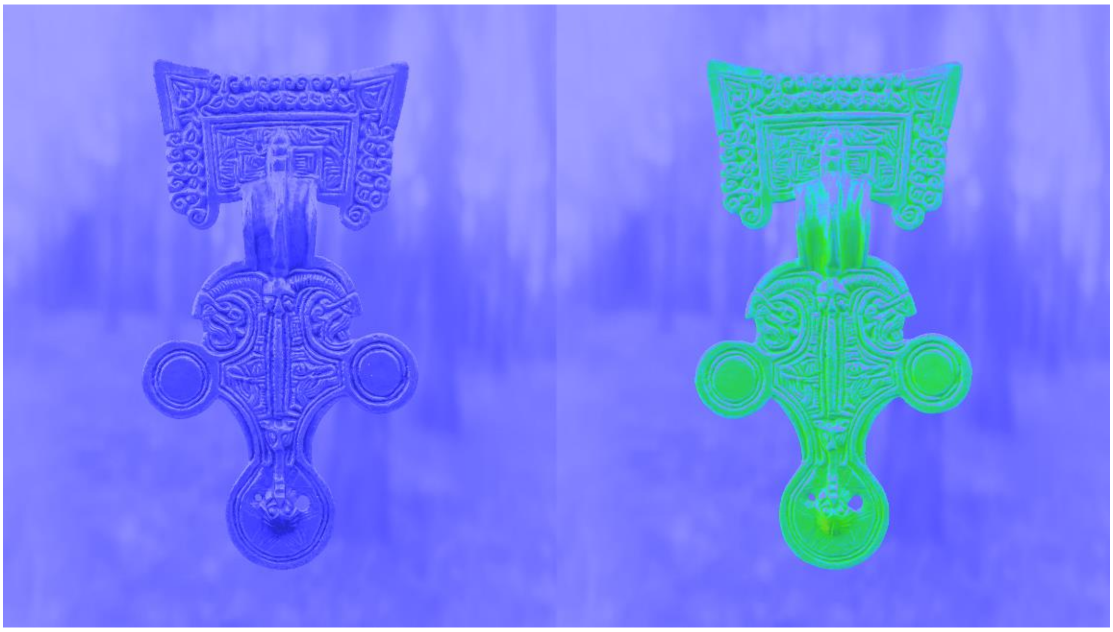

Figure 75: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 76: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

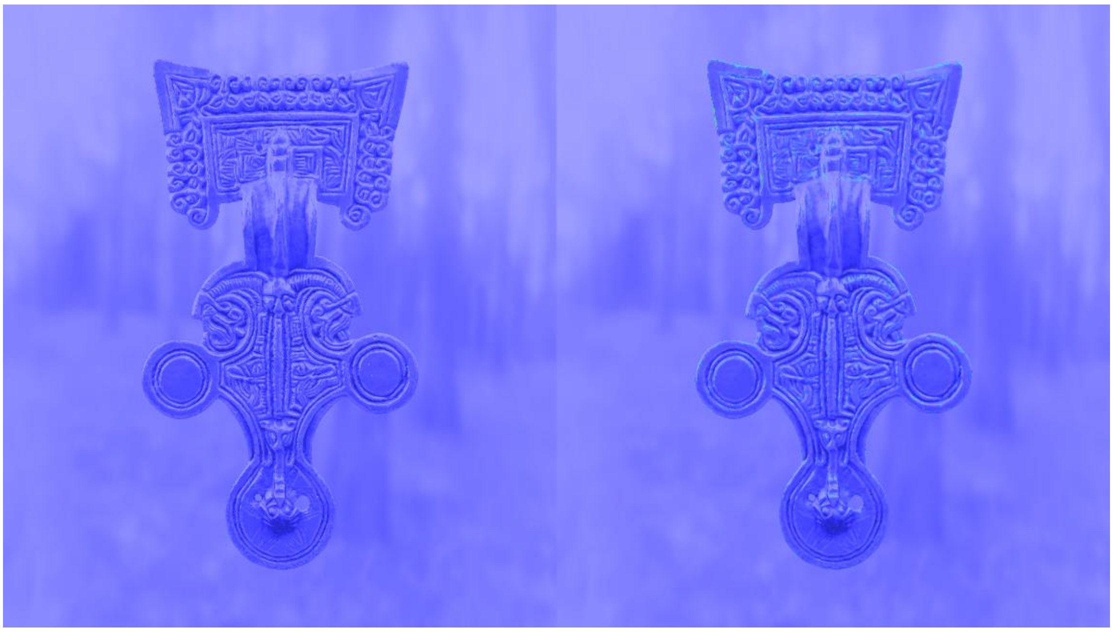

Figure 77: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

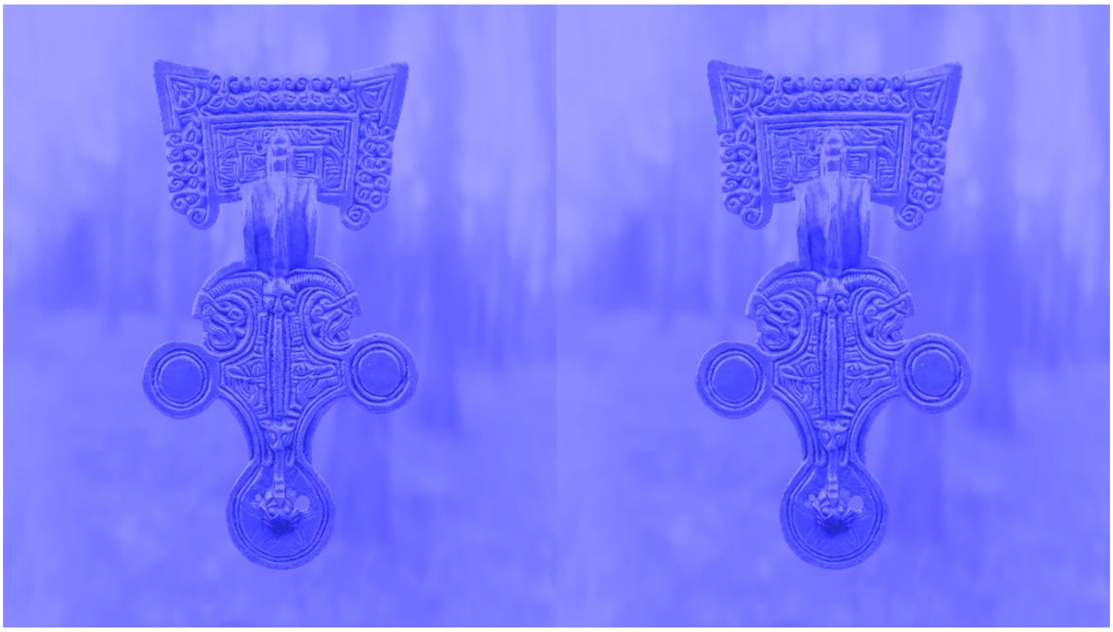

Figure 78: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

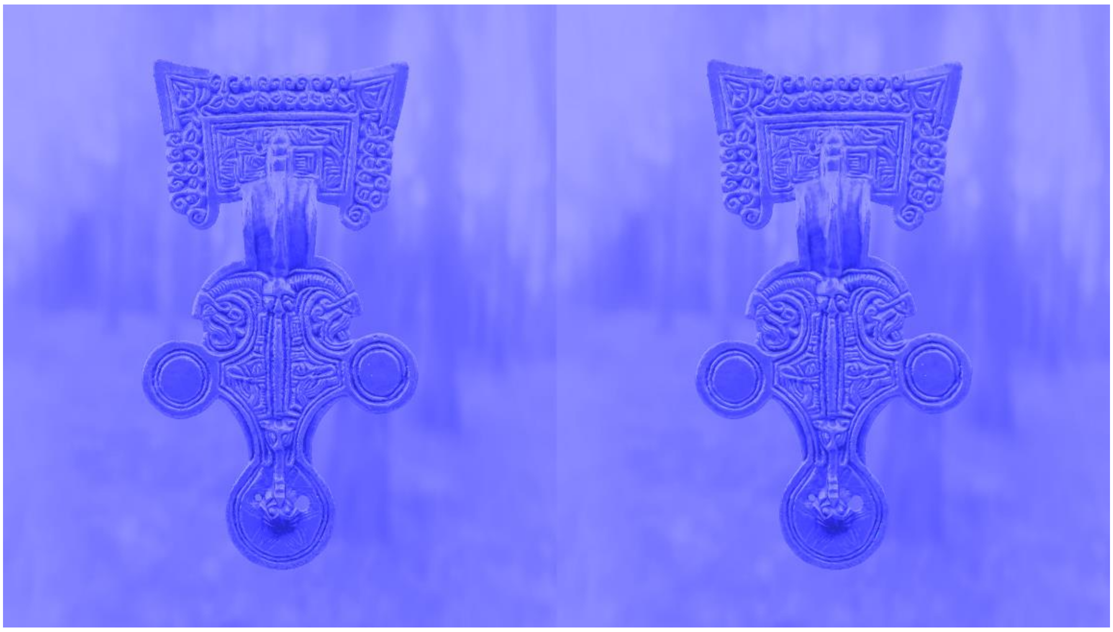

Figure 79: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 80: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

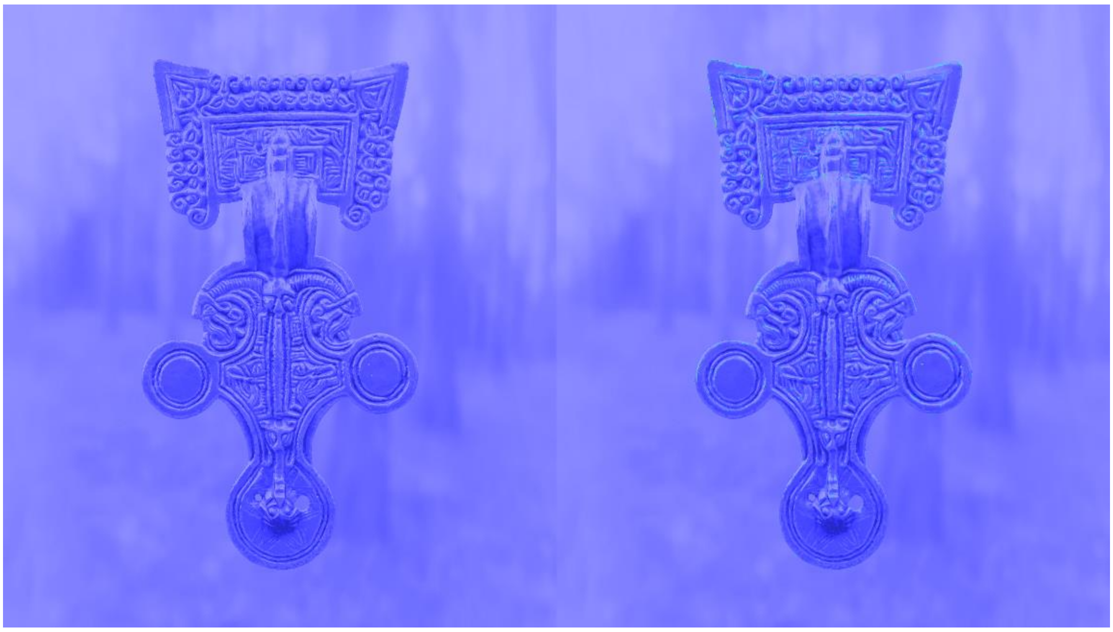

Figure 81: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

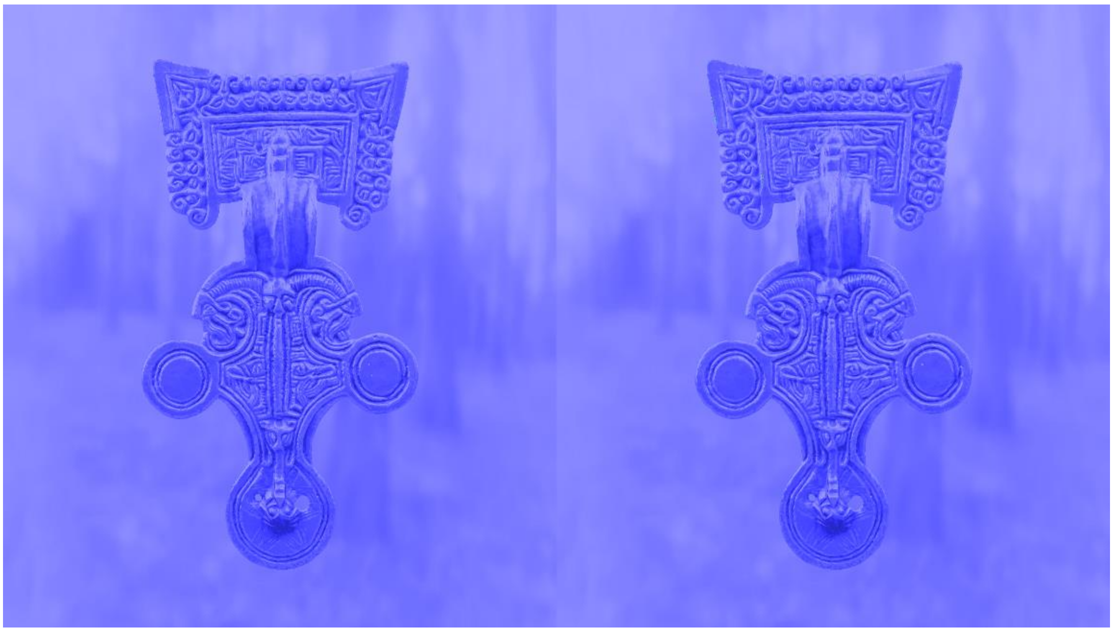

Figure 82: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

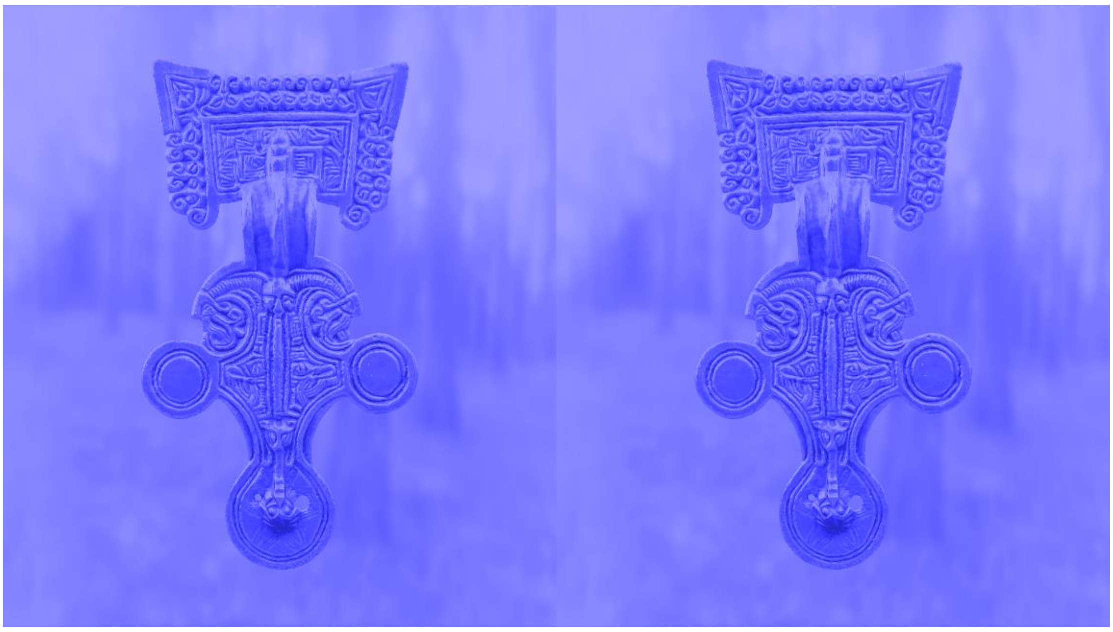

Figure 83: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 84: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

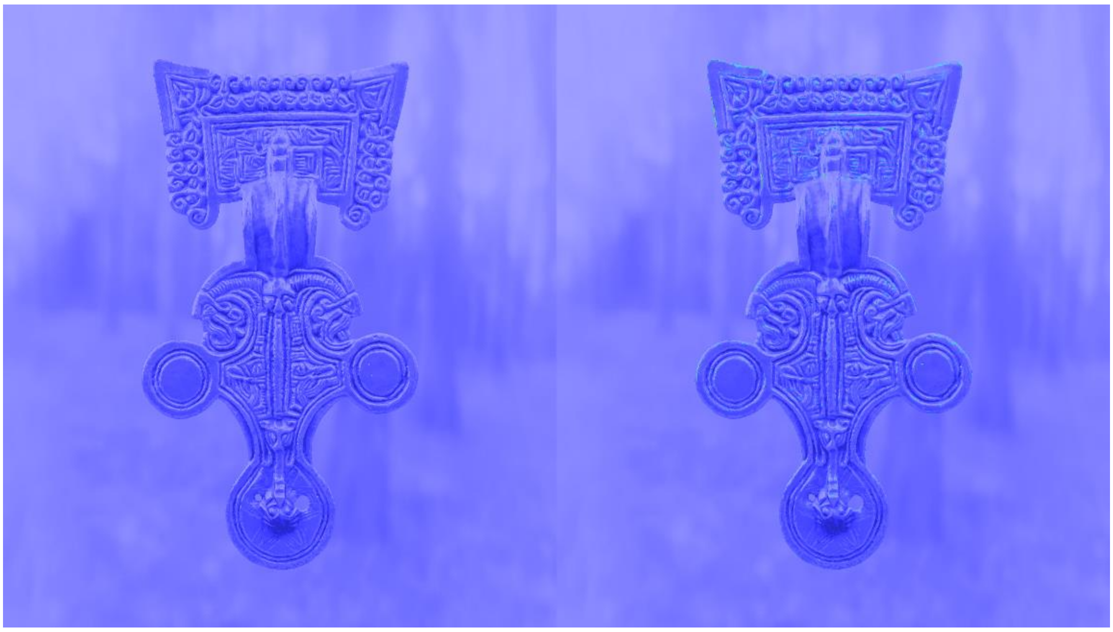

Figure 85: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

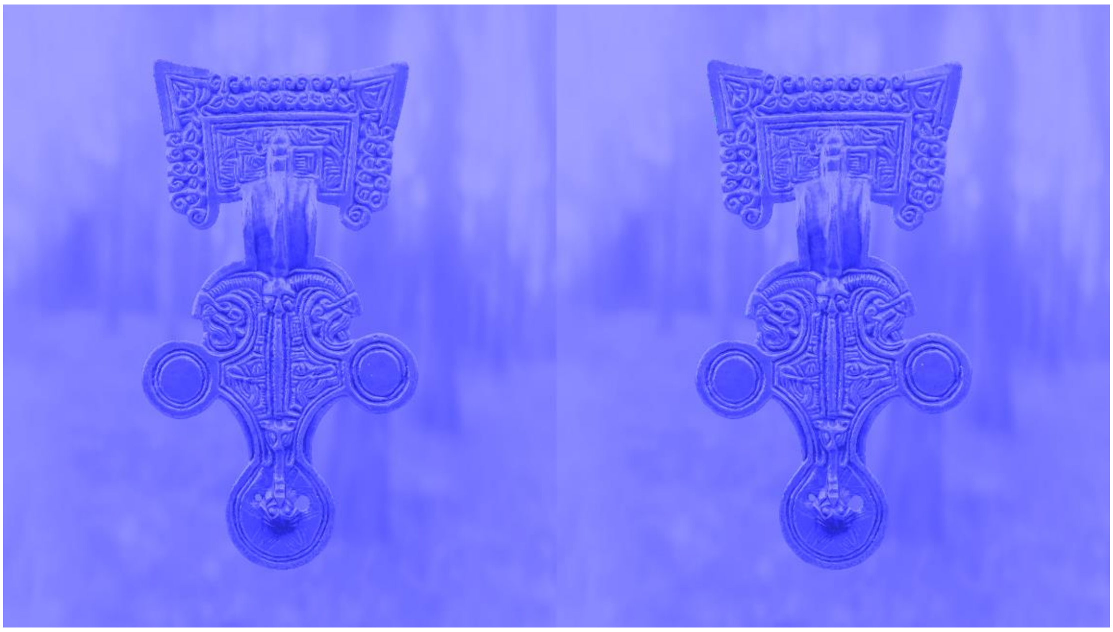

Figure 86: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 87: Reference model with stimuli at 10% polygonal resolution and no texture

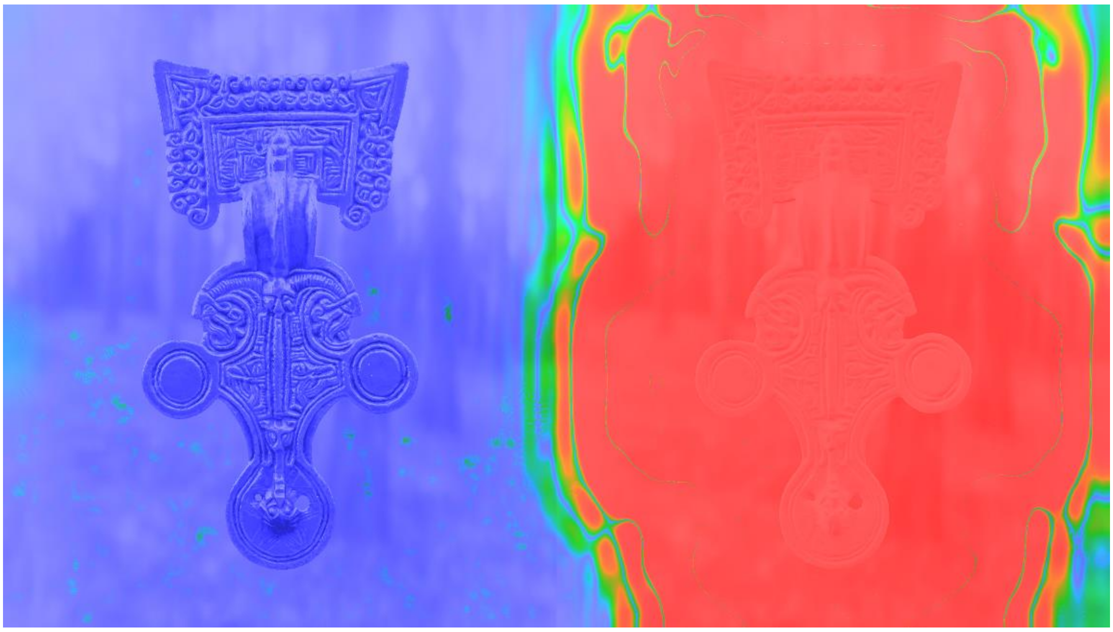

Figure 88: Reference model with stimuli at 40% polygonal resolution and no texture

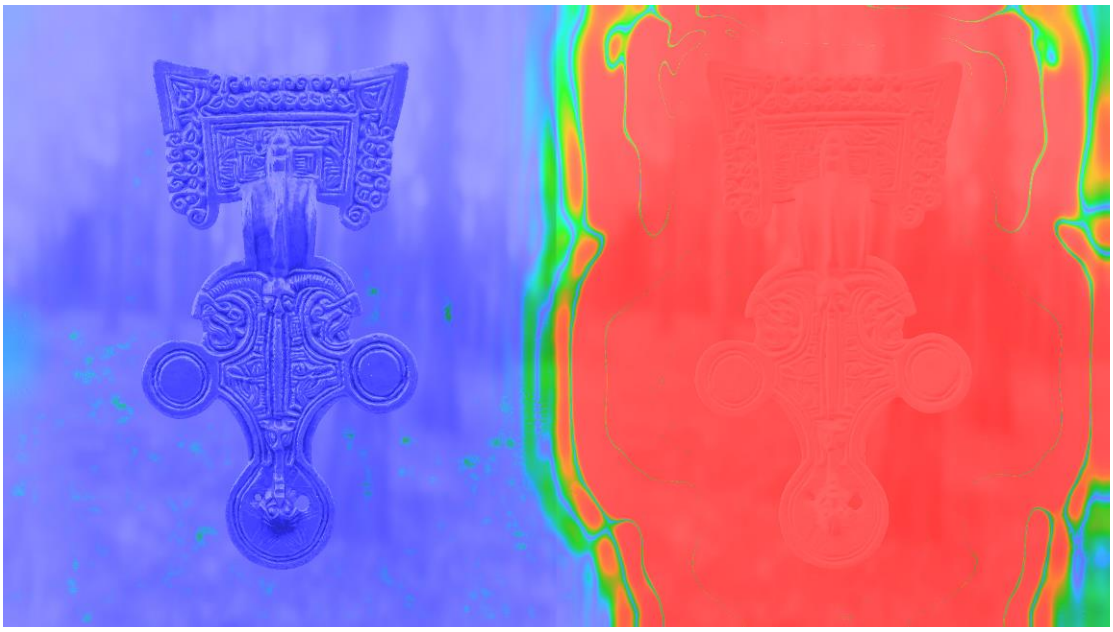

Figure 89: Reference model with stimuli at 70% polygonal resolution and no texture

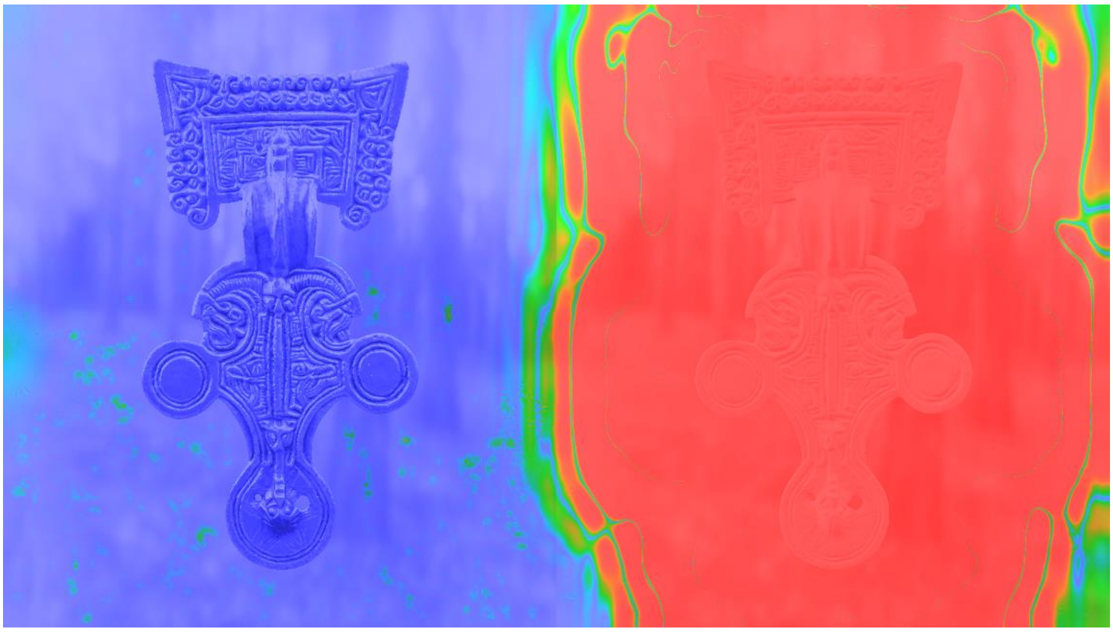

Figure 90: Reference model with stimuli at 100% polygonal resolution and no texture

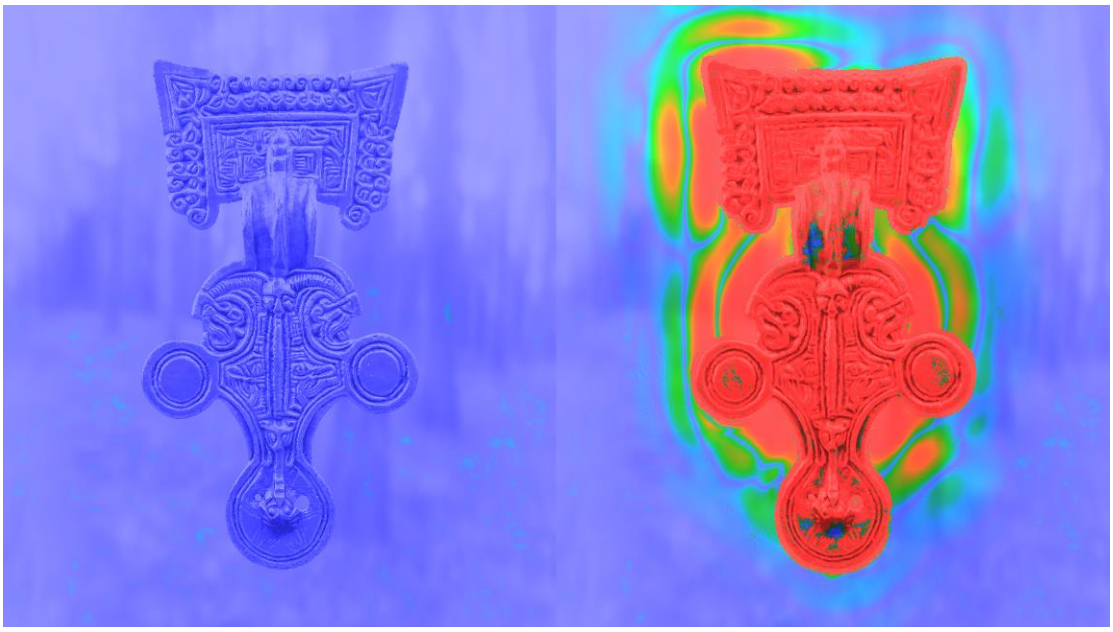

Figure 91: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 92: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 93: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 94: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

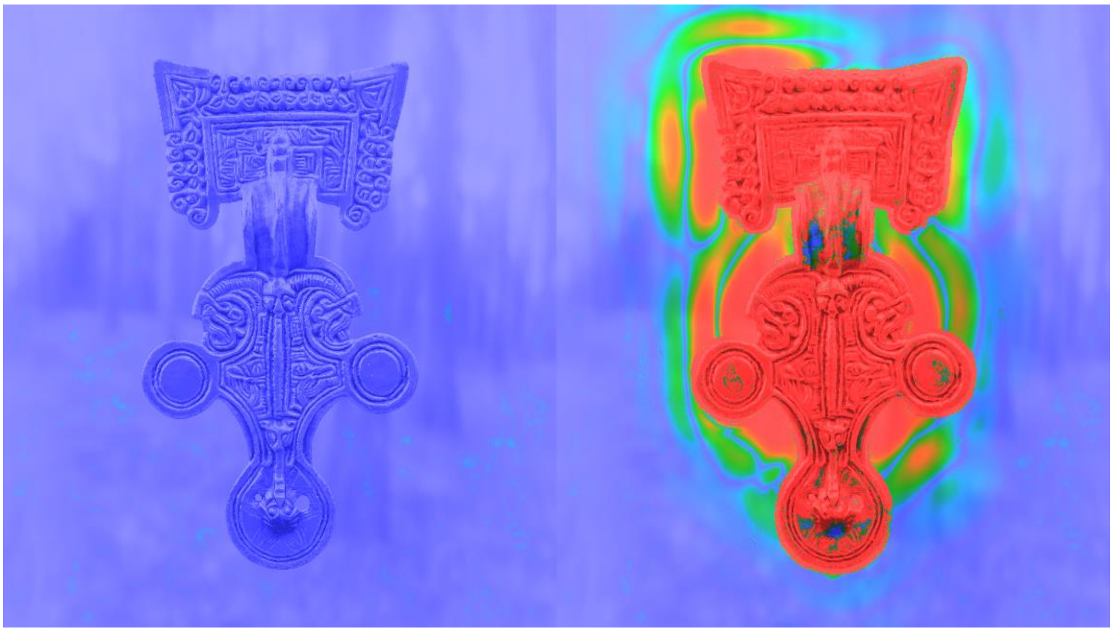

Figure 95: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 96: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 97: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 98: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

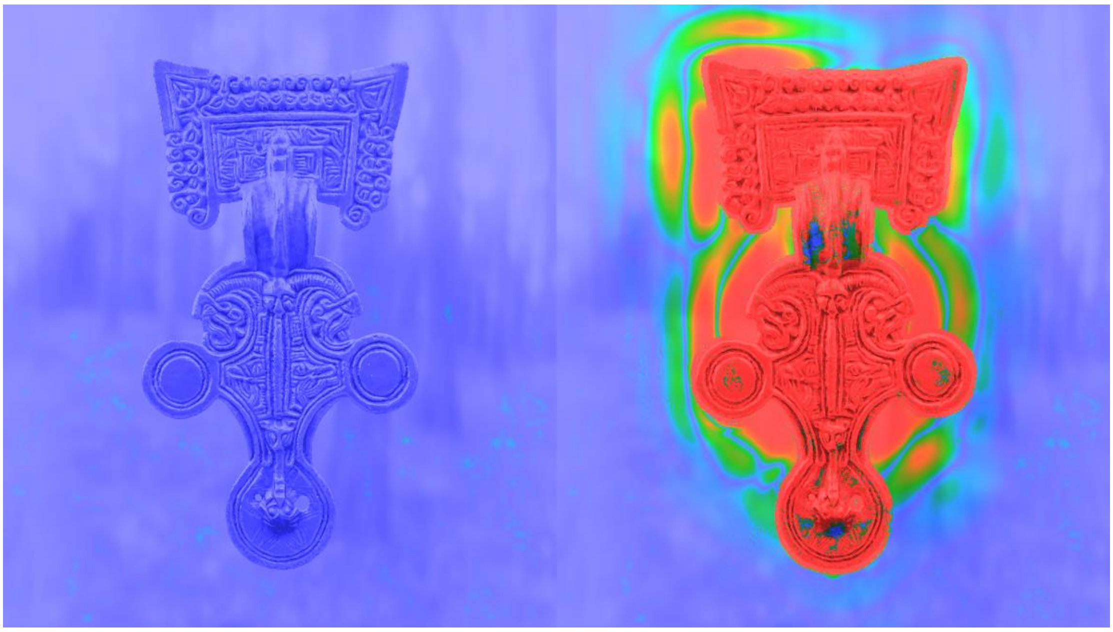

Figure 99: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 100: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 101: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the sRGB colour space

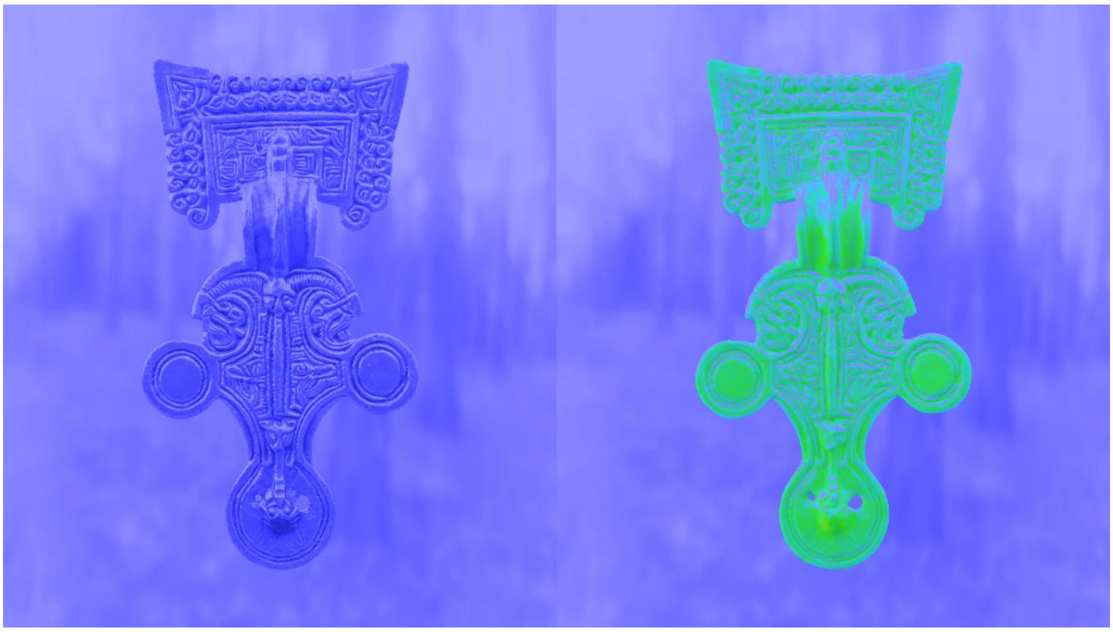

Figure 102: Reference model with stimuli at 10% polygonal resolution and no texture

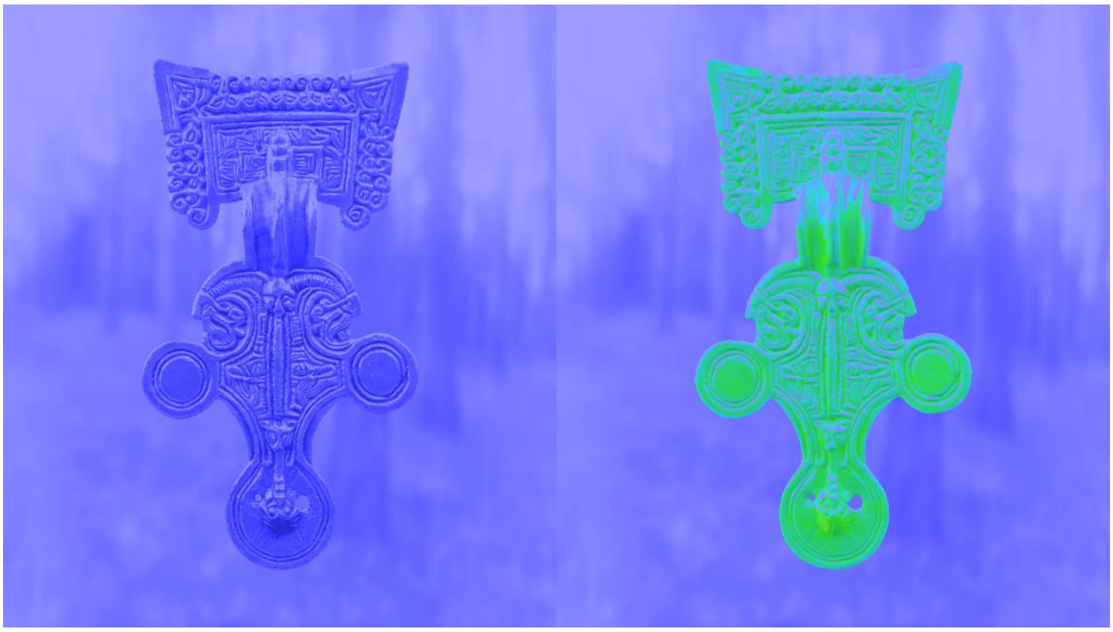

Figure 103: Reference model with stimuli at 40% polygonal resolution and no texture

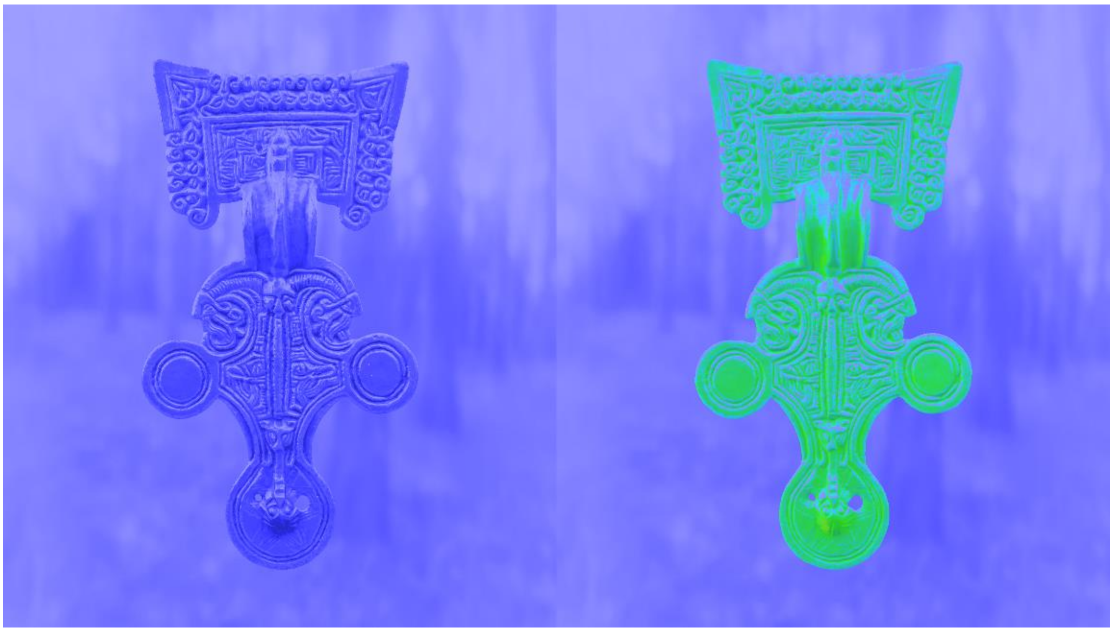

Figure 104: Reference model with stimuli at 70% polygonal resolution and no texture

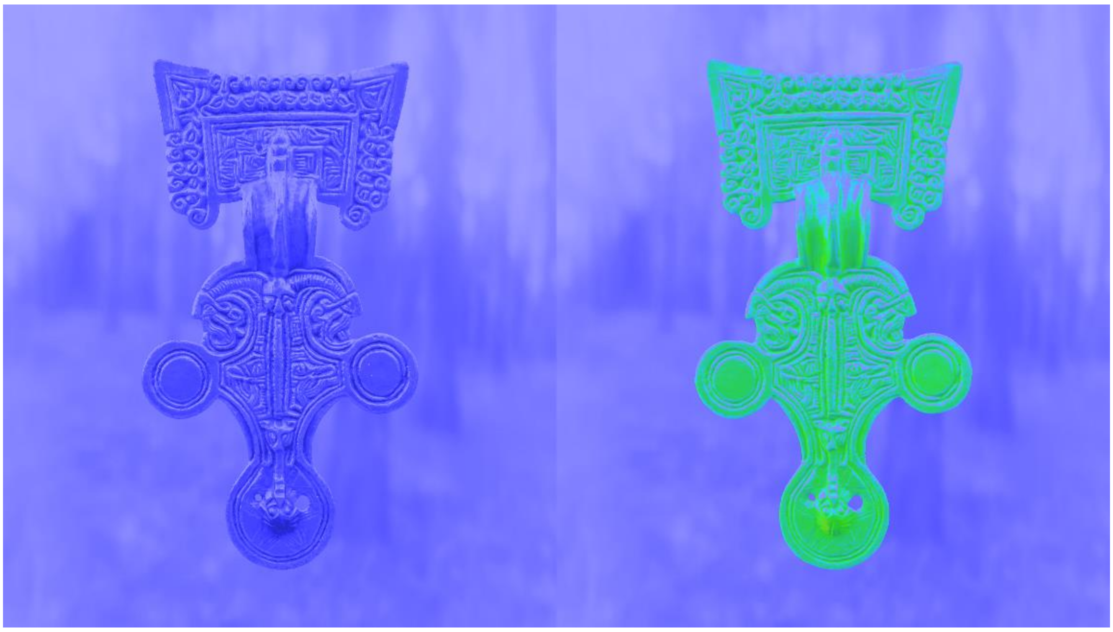

Figure 105: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 106: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

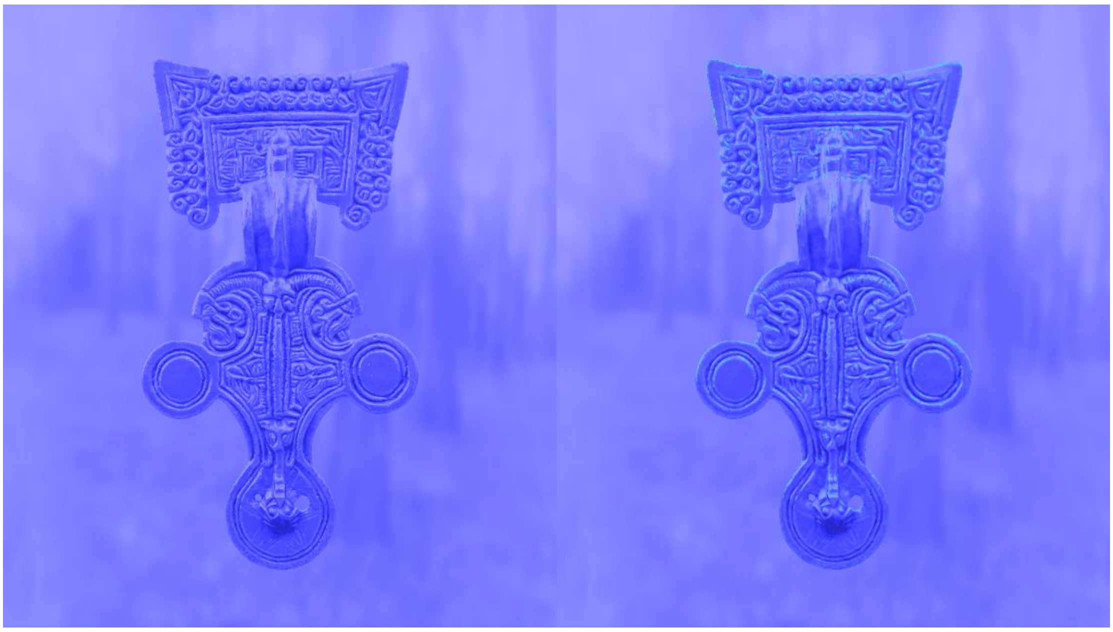

Figure 107: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

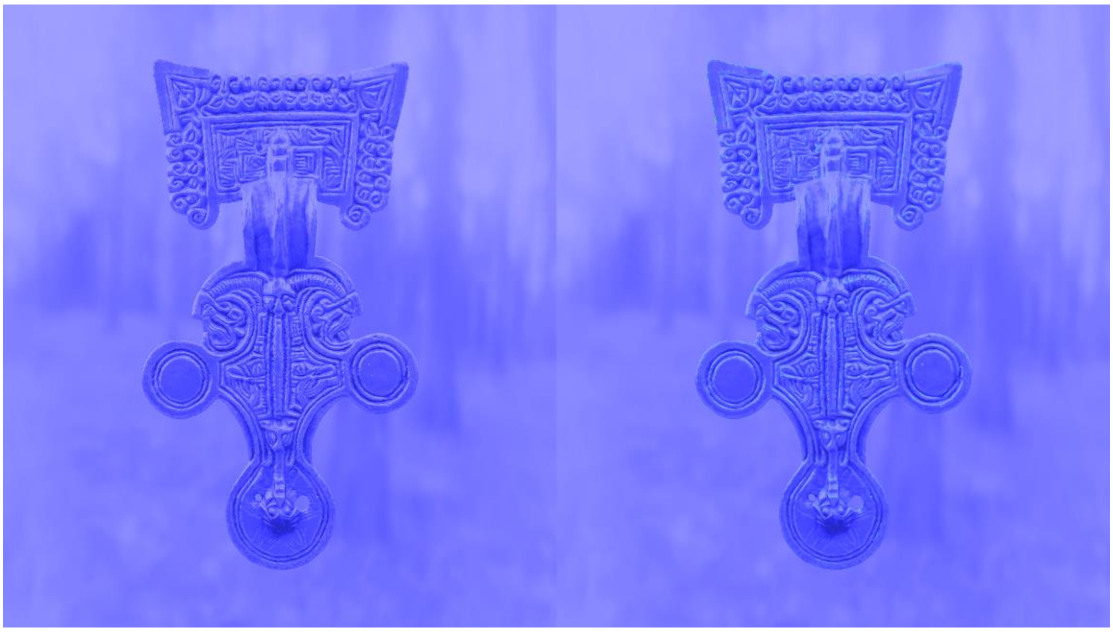

Figure 108: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

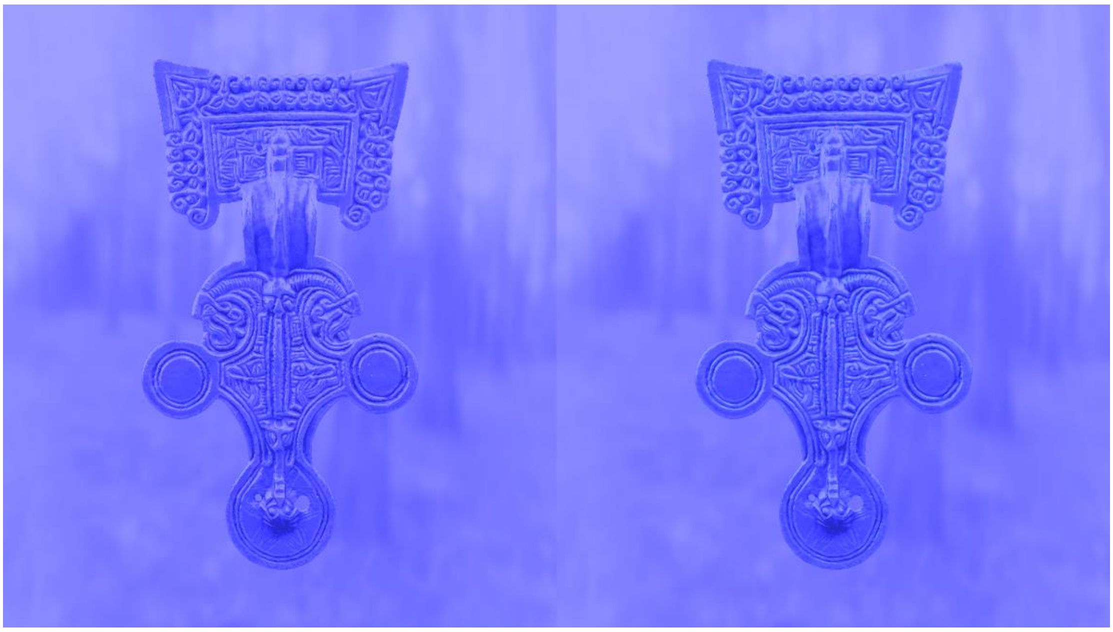

Figure 109: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 110: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

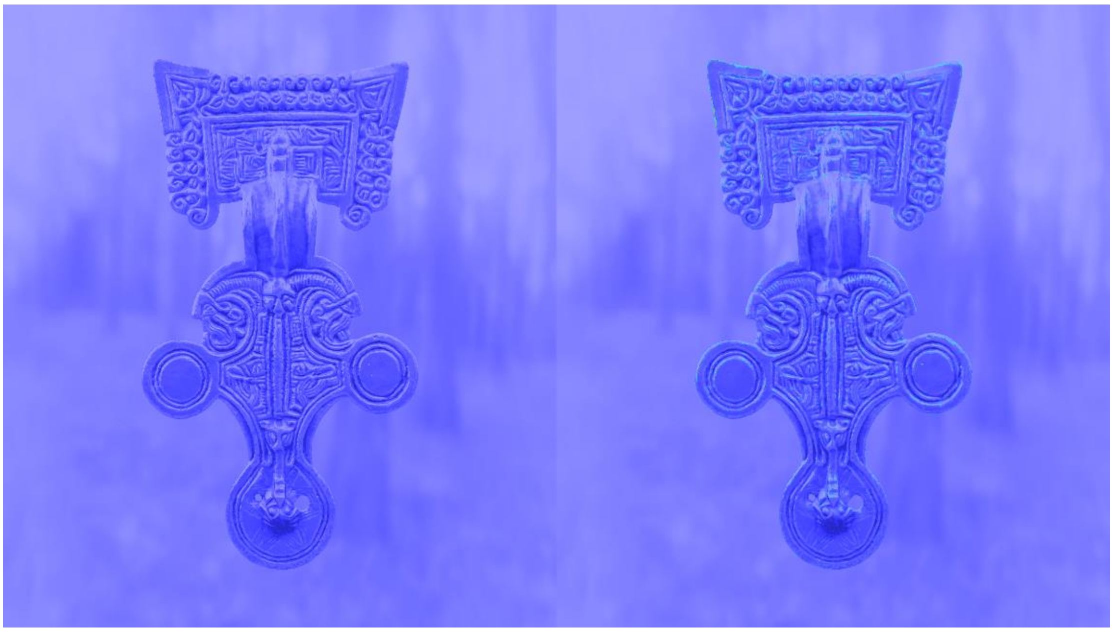

Figure 111: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

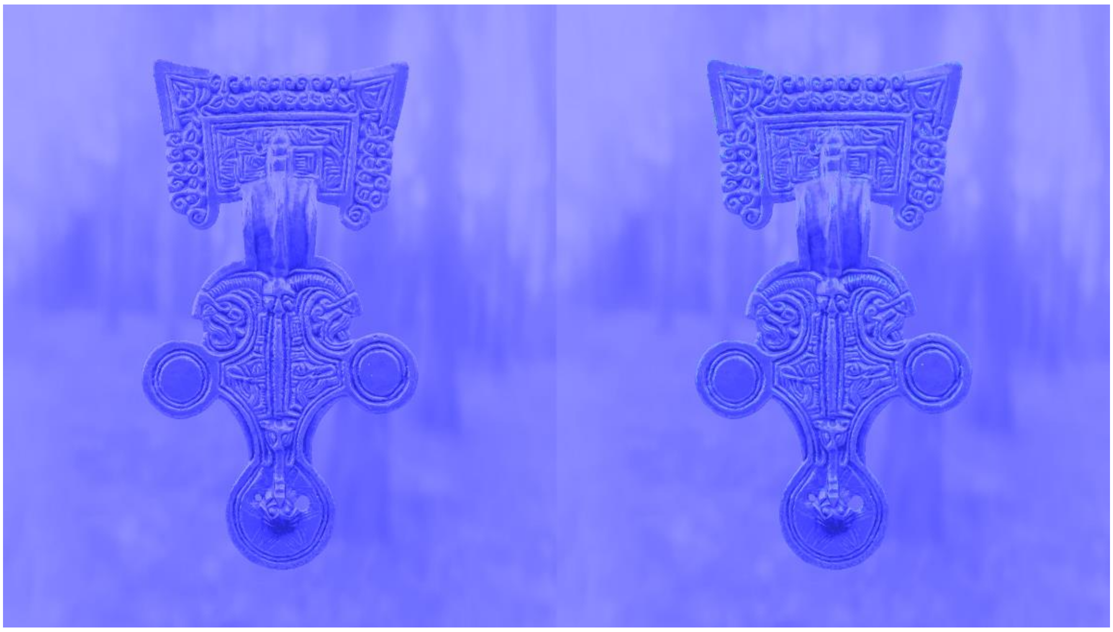

Figure 112: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

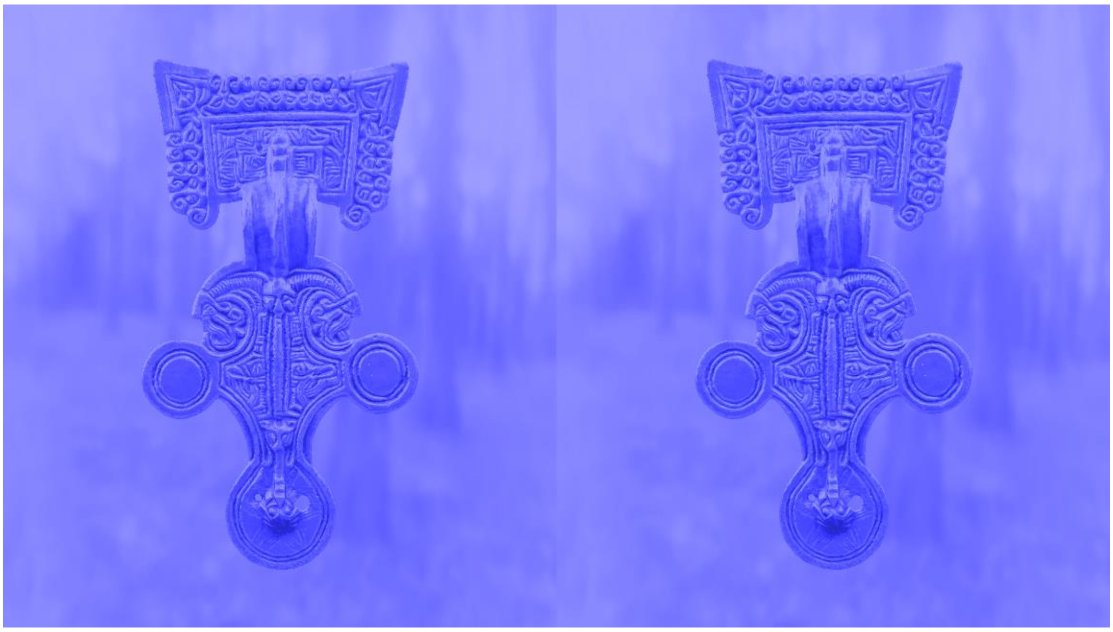

Figure 113: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 114: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

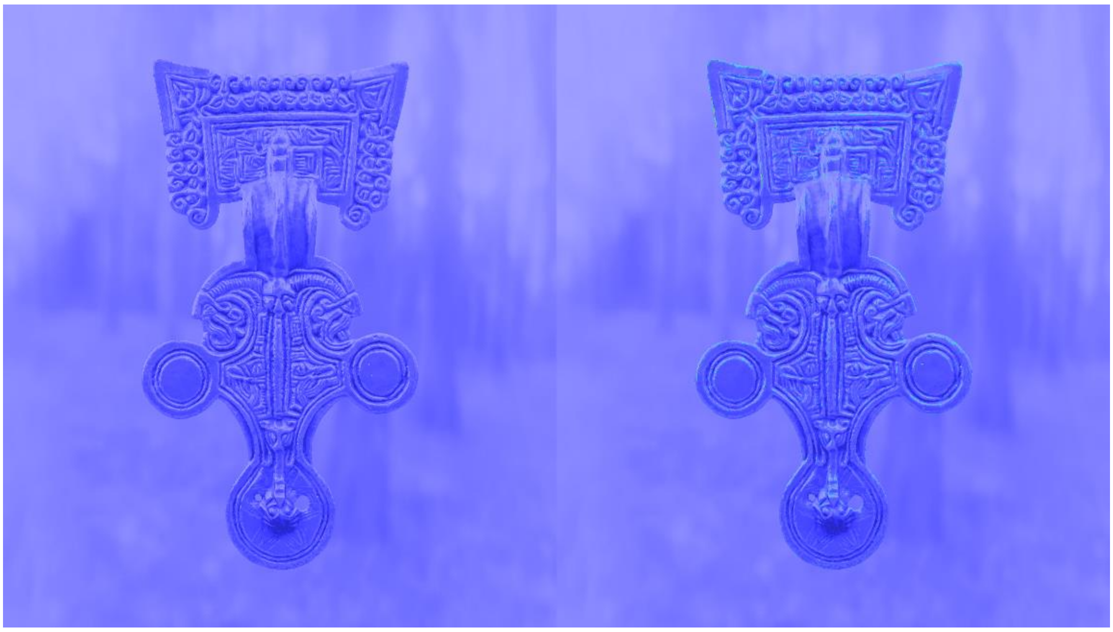

Figure 115: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

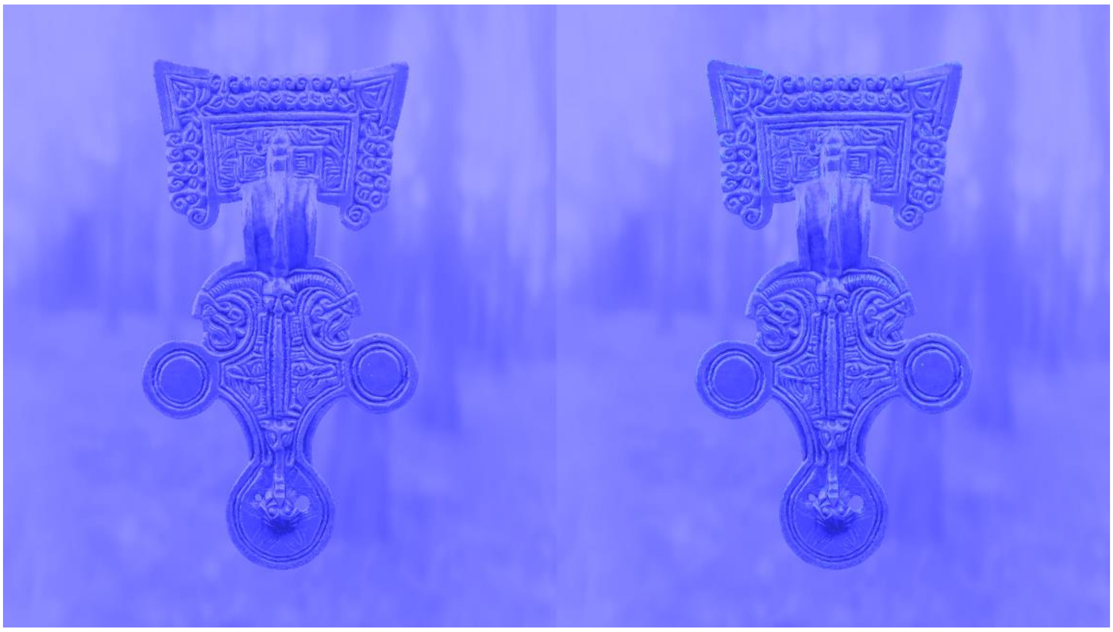

Figure 116: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Egyptian Relief

Rendered images of the different stimuli at different texture and polygonal resolutions

Figure 117: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 118: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 119: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 120: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 121: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 122: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 123: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 124: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 125: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 126: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 127: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 128: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 129: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 130: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 131: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Figure 132: Reference model with stimuli at 100% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

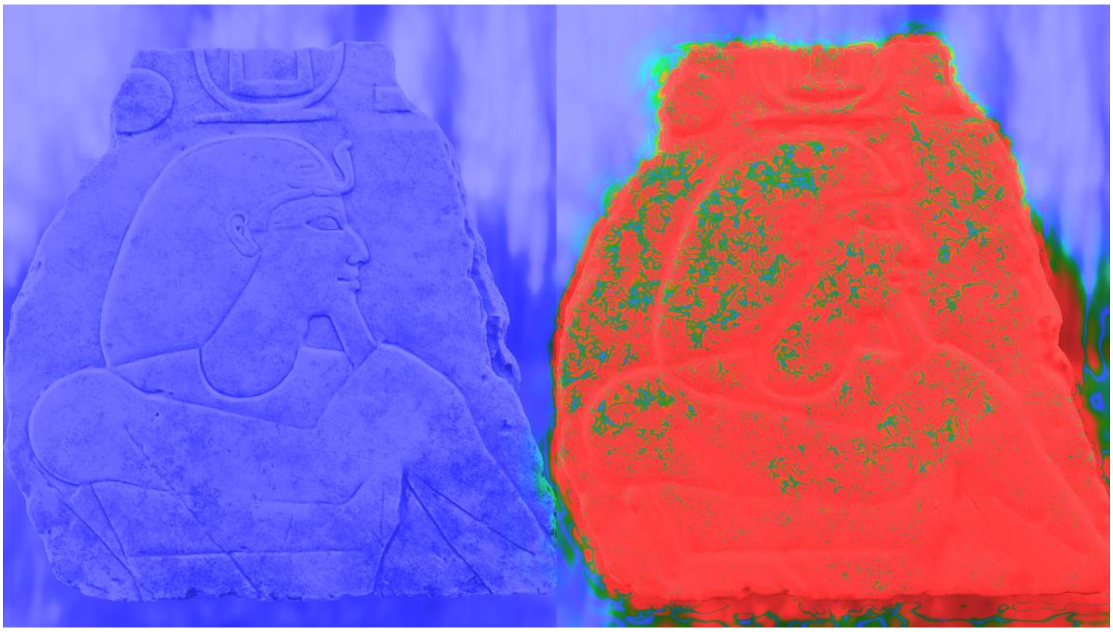

Figure 133: Reference model with stimuli at 10% polygonal resolution and no texture

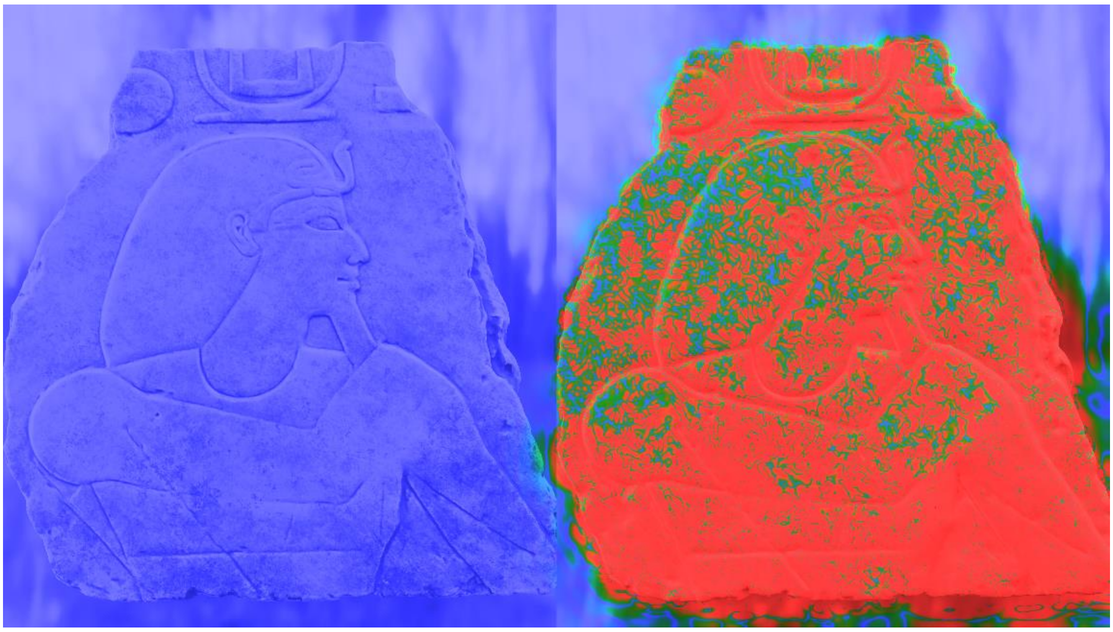

Figure 134: Reference model with stimuli at 40% polygonal resolution and no texture

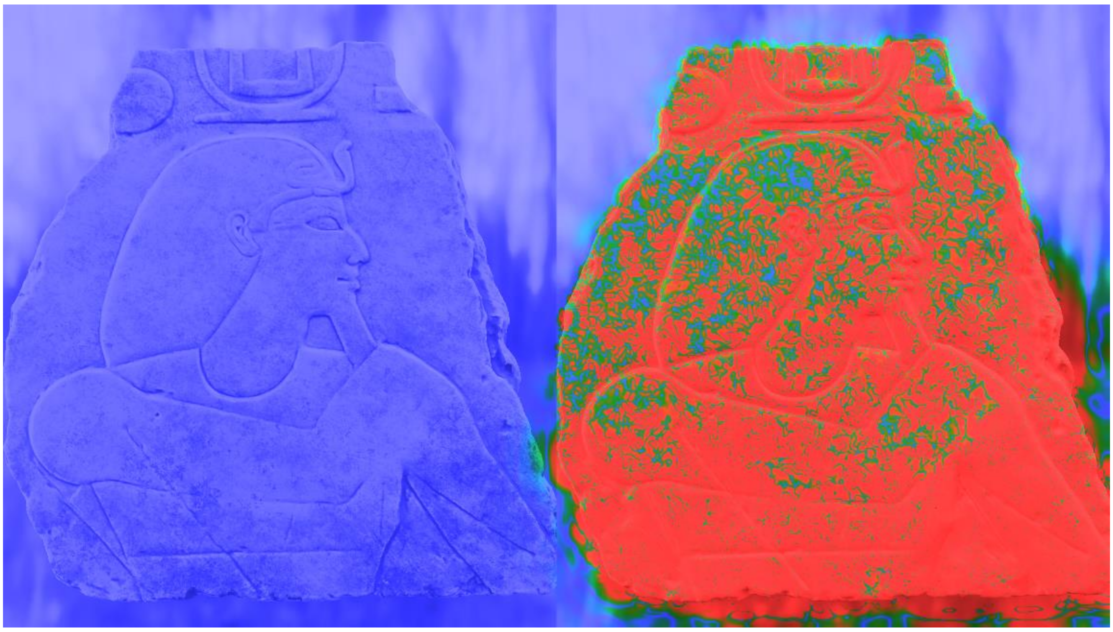

Figure 135: Reference model with stimuli at 70% polygonal resolution and no texture

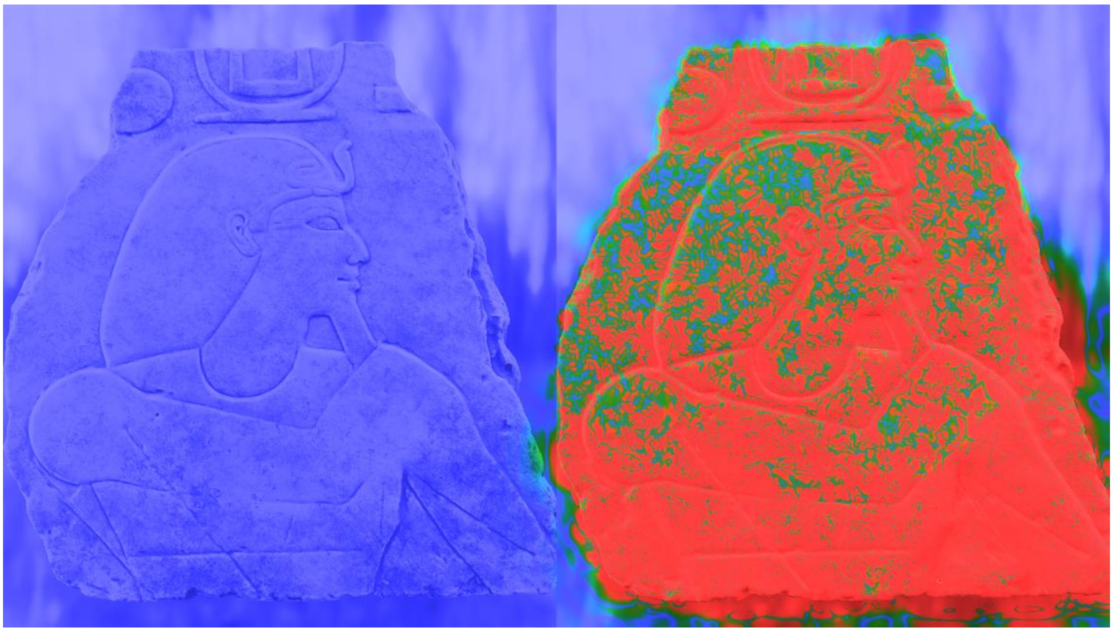

Figure 136: Reference model with stimuli at 100% polygonal resolution and no texture

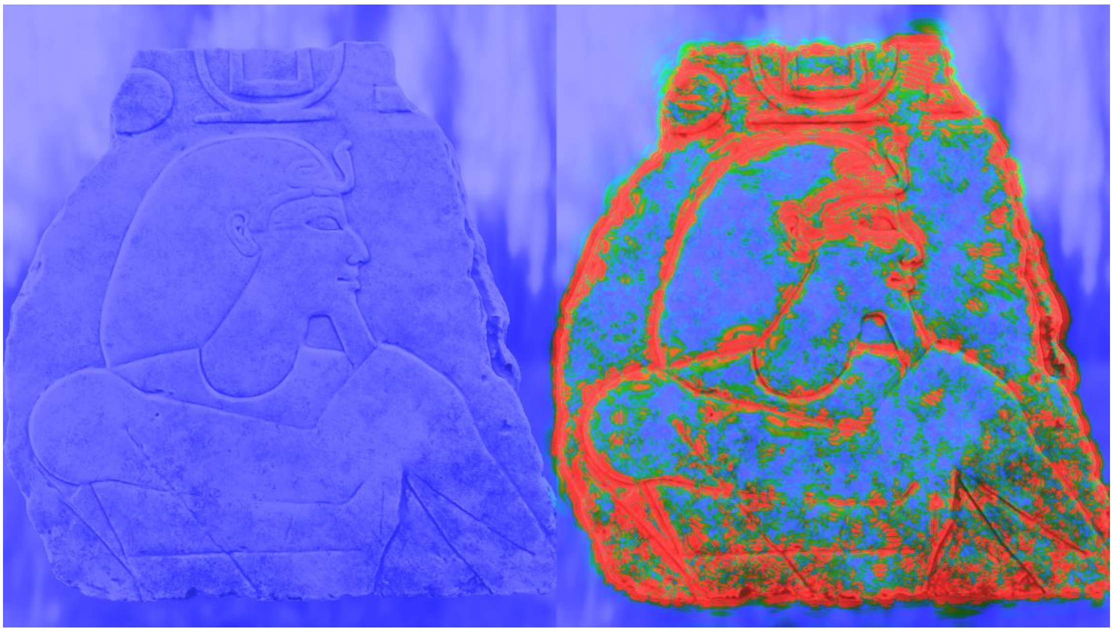

Figure 137: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

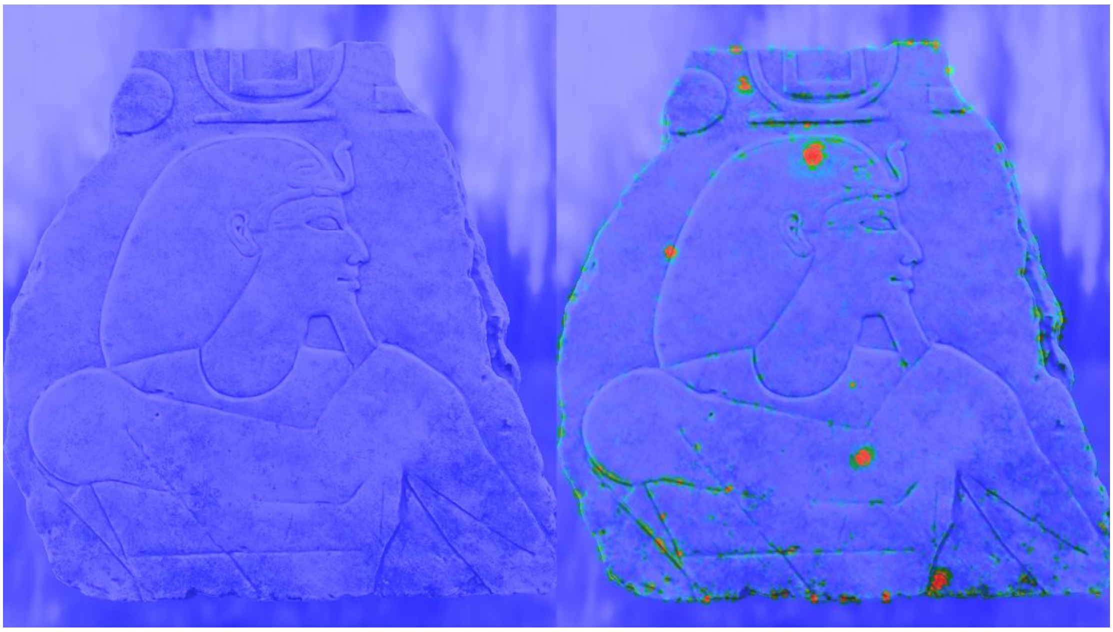

Figure 138: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

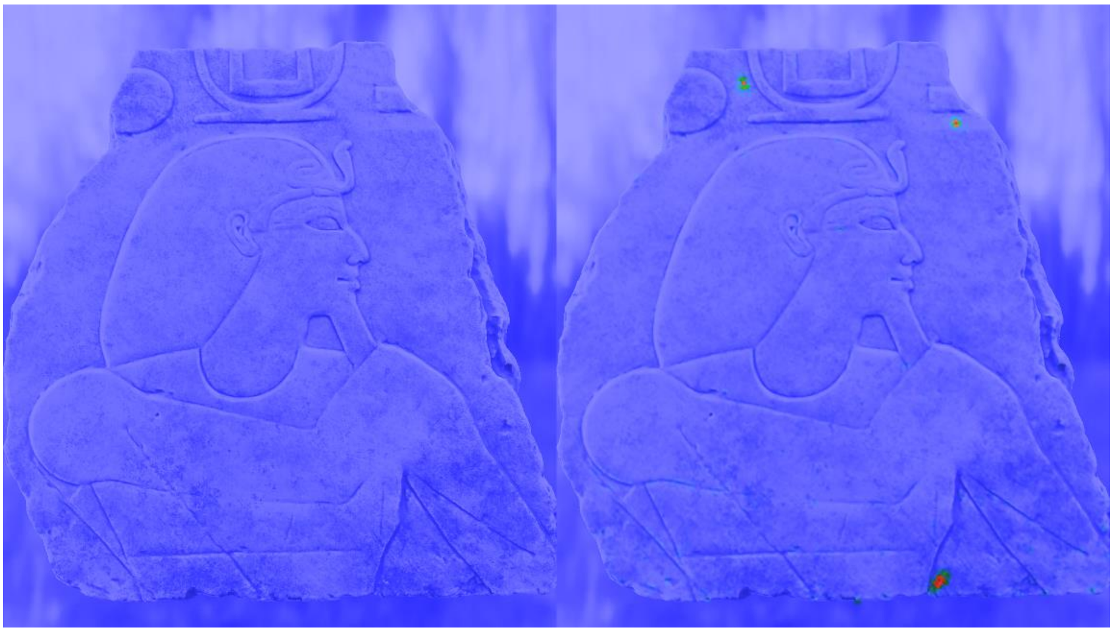

Figure 139: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

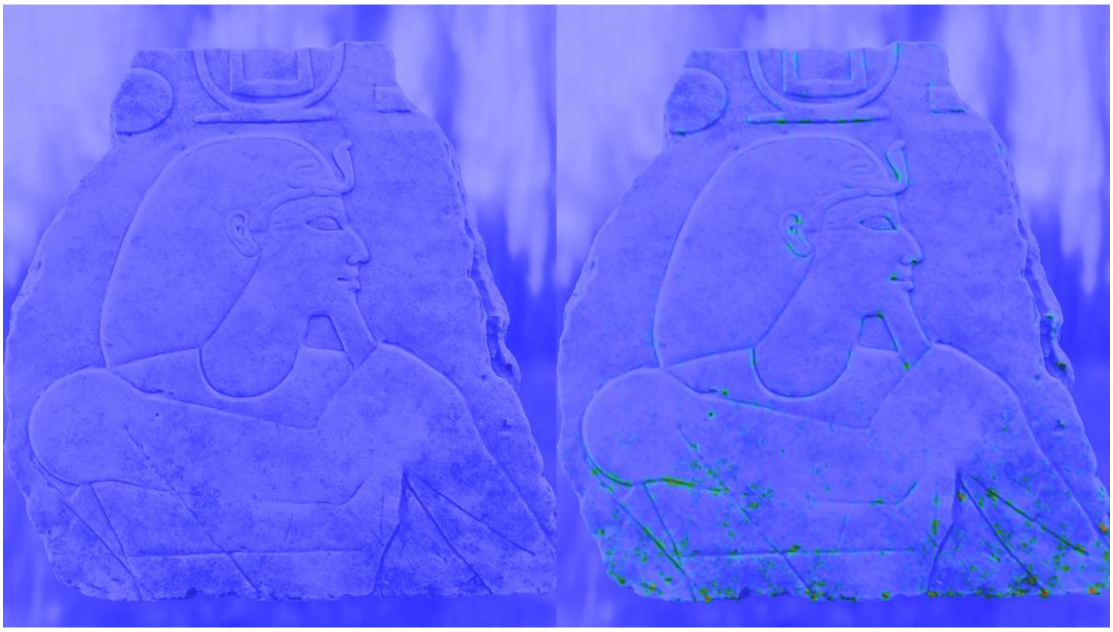

Figure 140: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

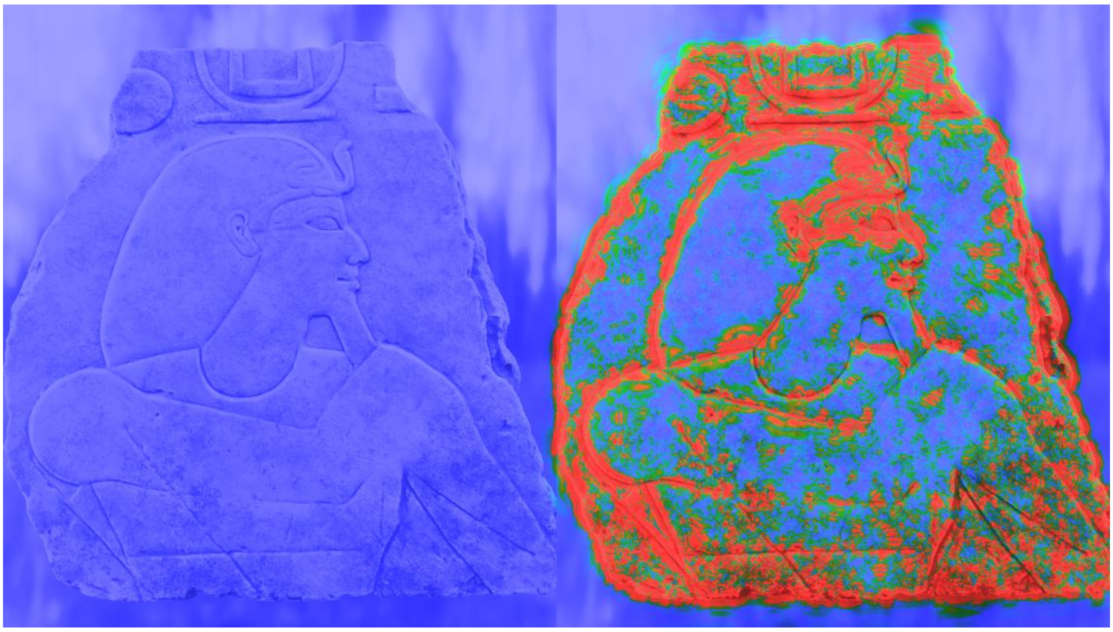

Figure 141: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

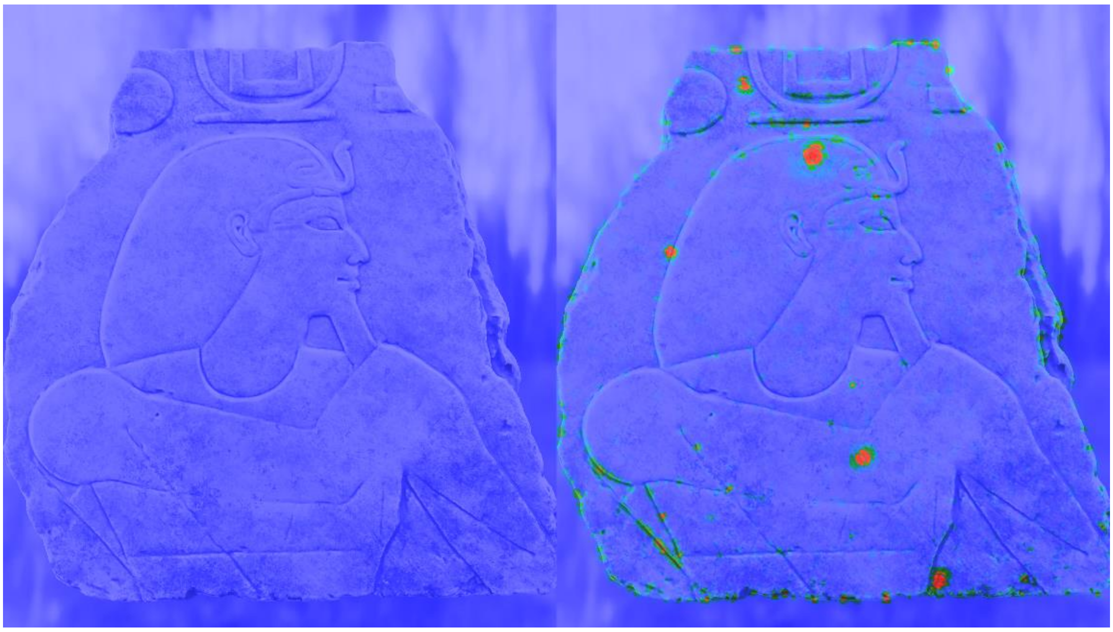

Figure 142: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

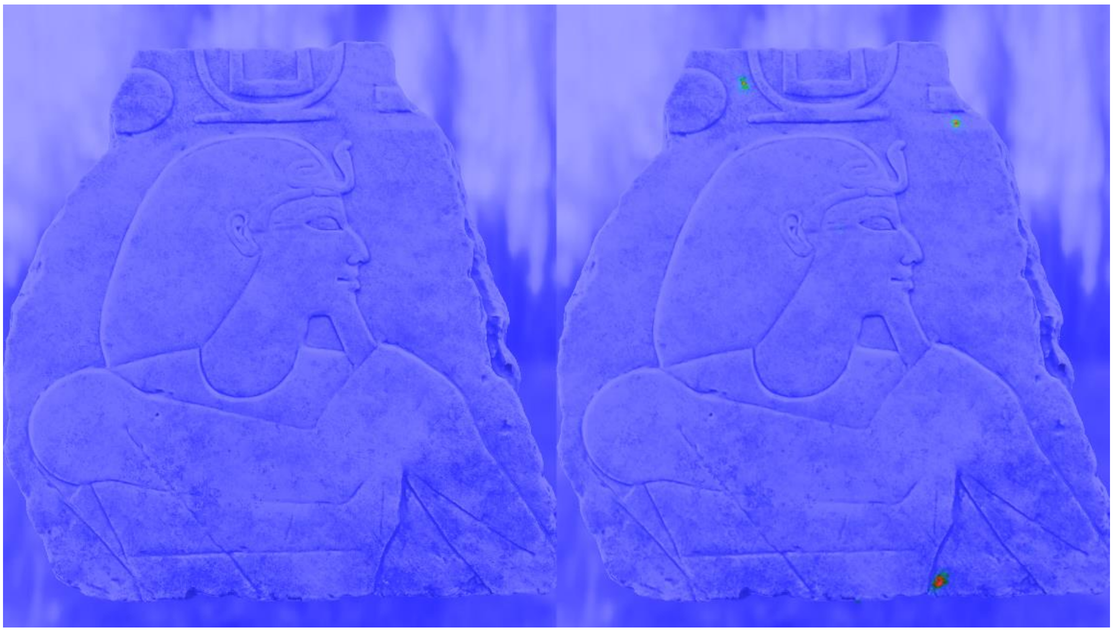

Figure 143: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 144: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

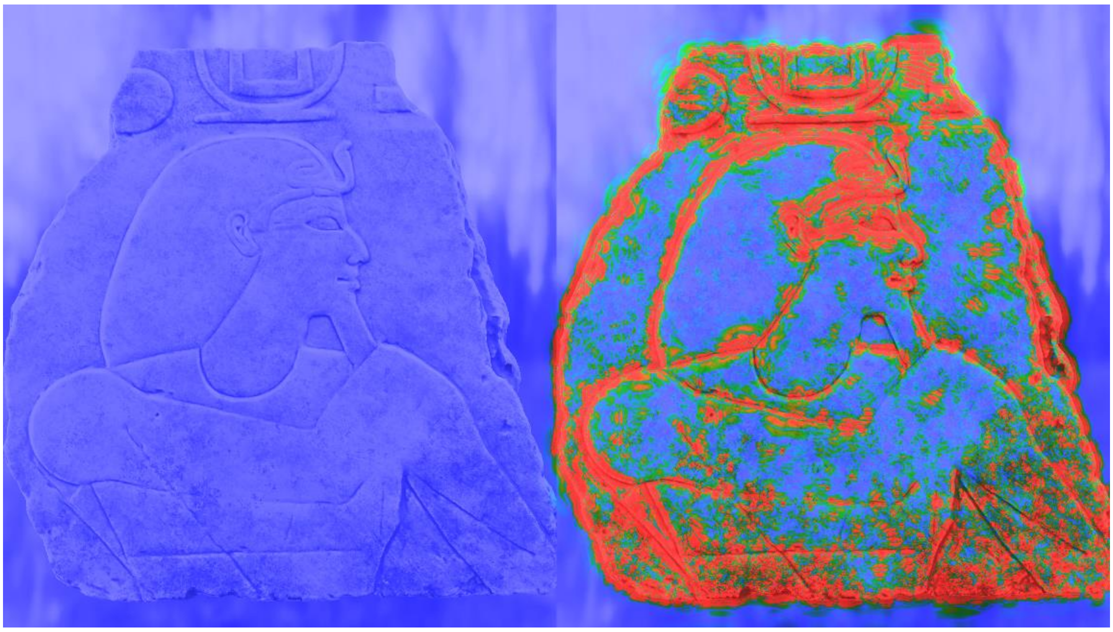

Figure 145: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

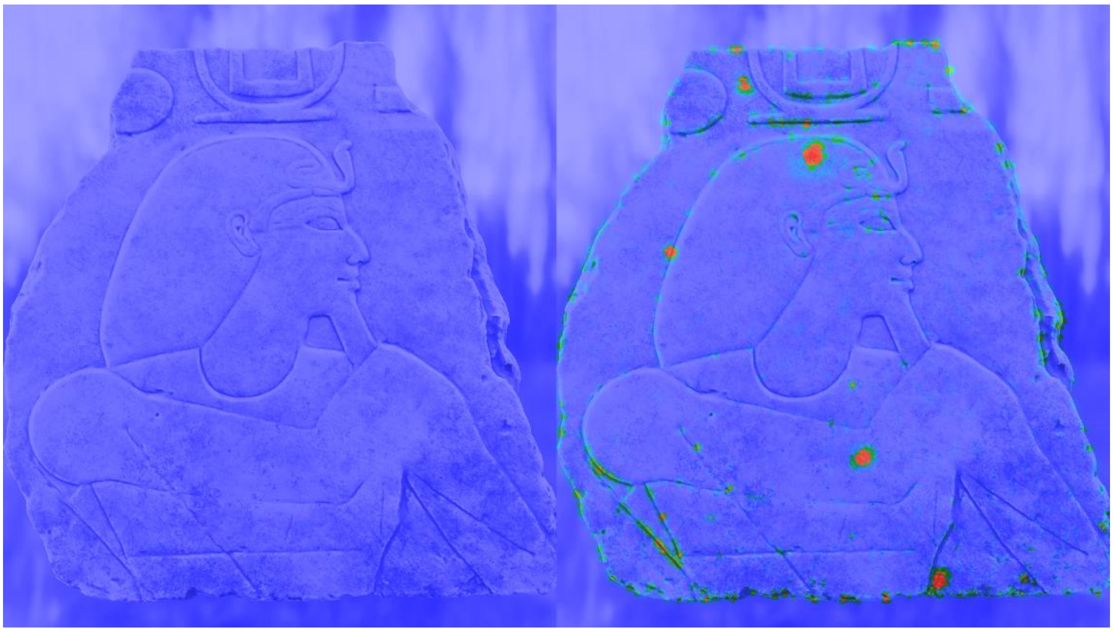

Figure 146: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

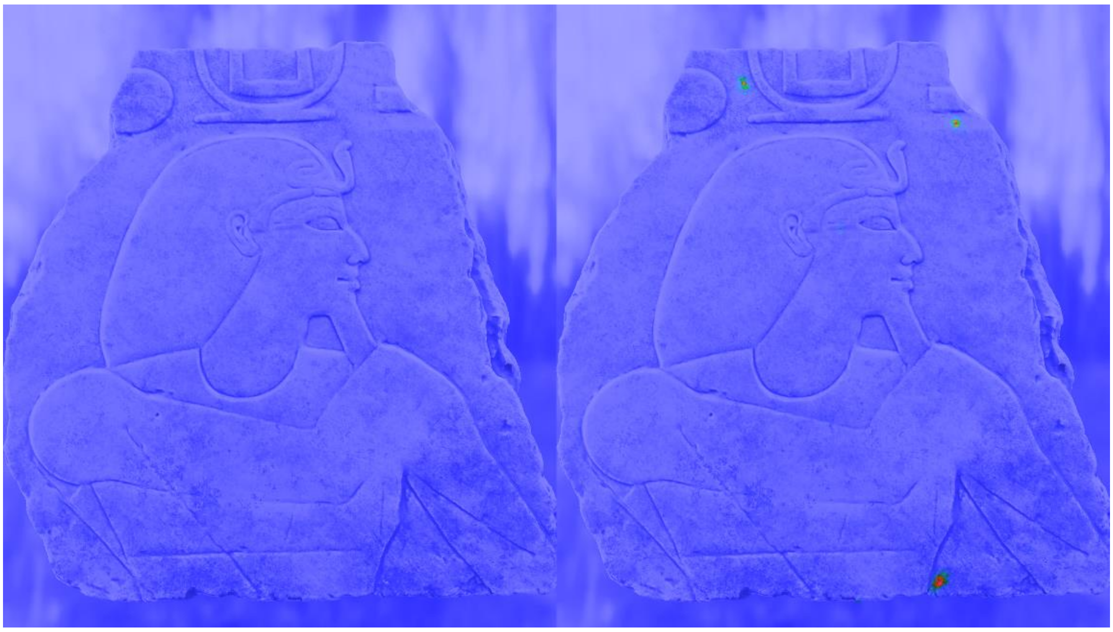

Figure 147: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

Figure 148: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 149: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 150: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 151: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 152: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 153: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 154: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 155: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 156: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 157: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 158: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 159: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 160: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 161: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 162: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the sRGB colour space

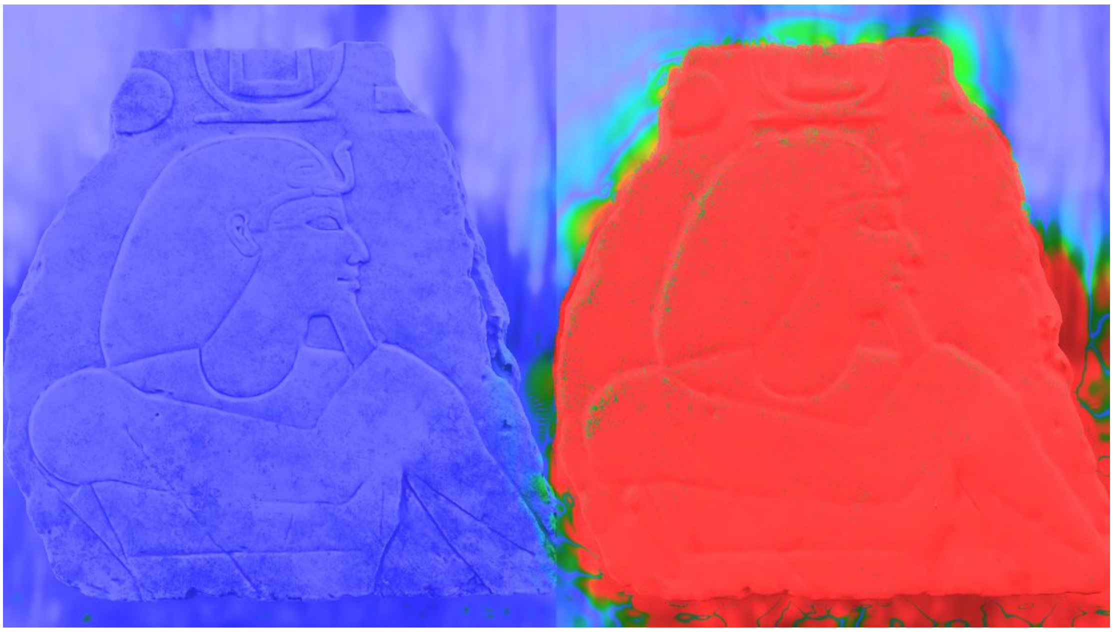

Figure 163: Reference model with stimuli at 10% polygonal resolution and no texture

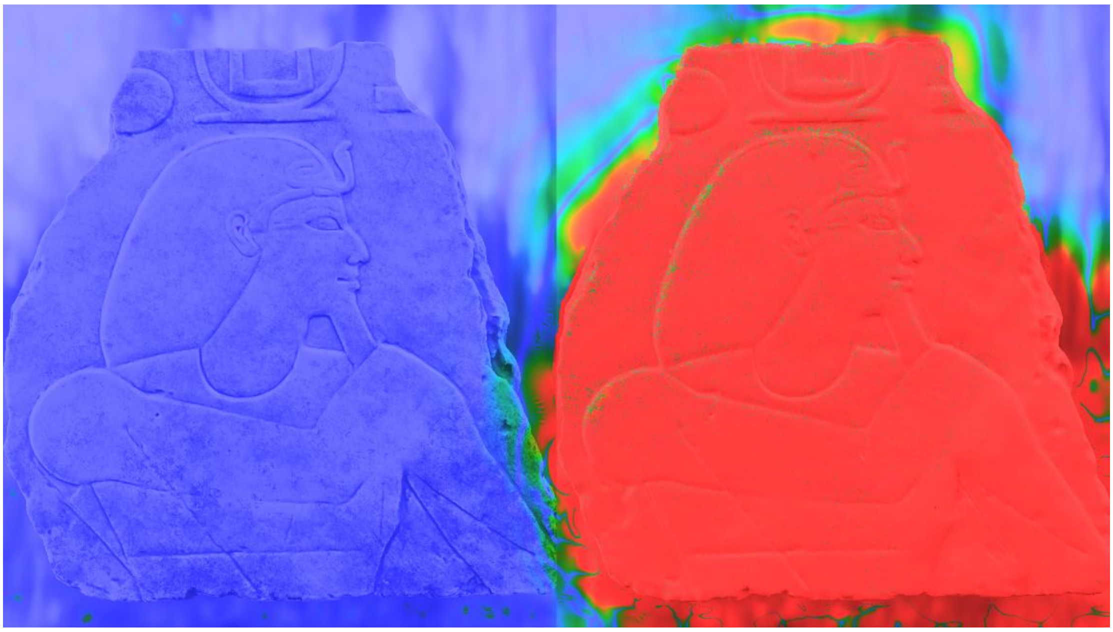

Figure 164: Reference model with stimuli at 40% polygonal resolution and no texture

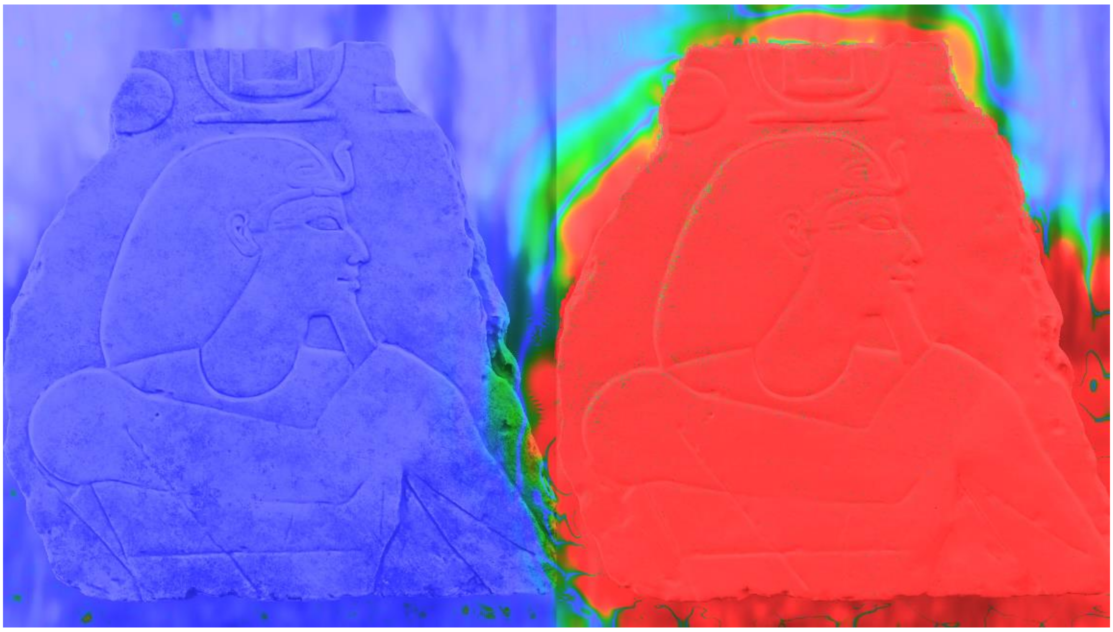

Figure 165: Reference model with stimuli at 70% polygonal resolution and no texture

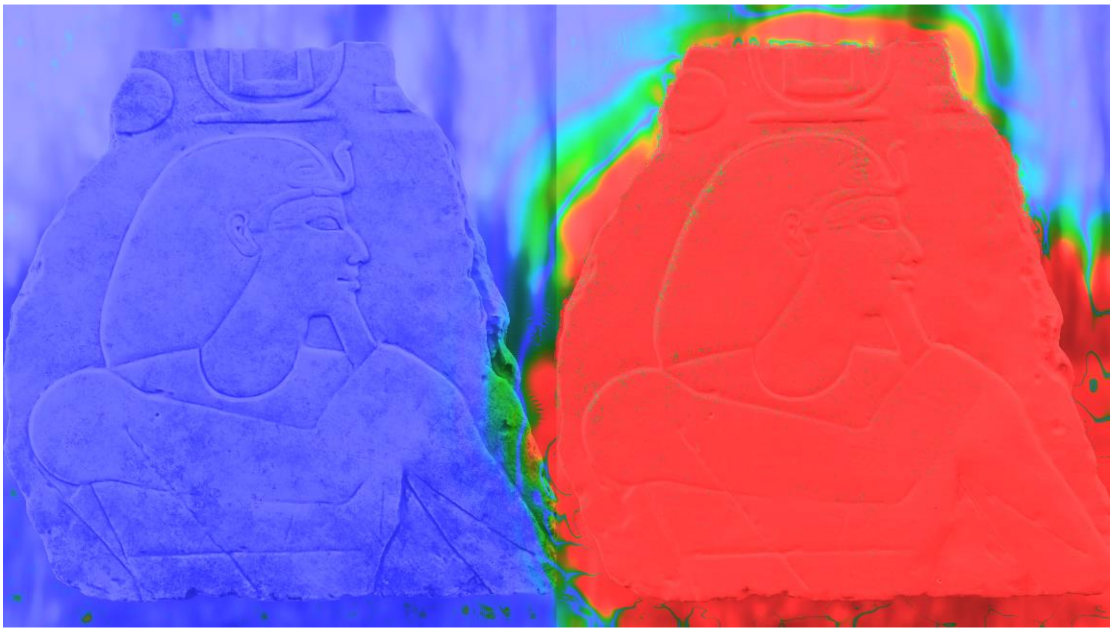

Figure 166: Reference model with stimuli at 100% polygonal resolution and no texture

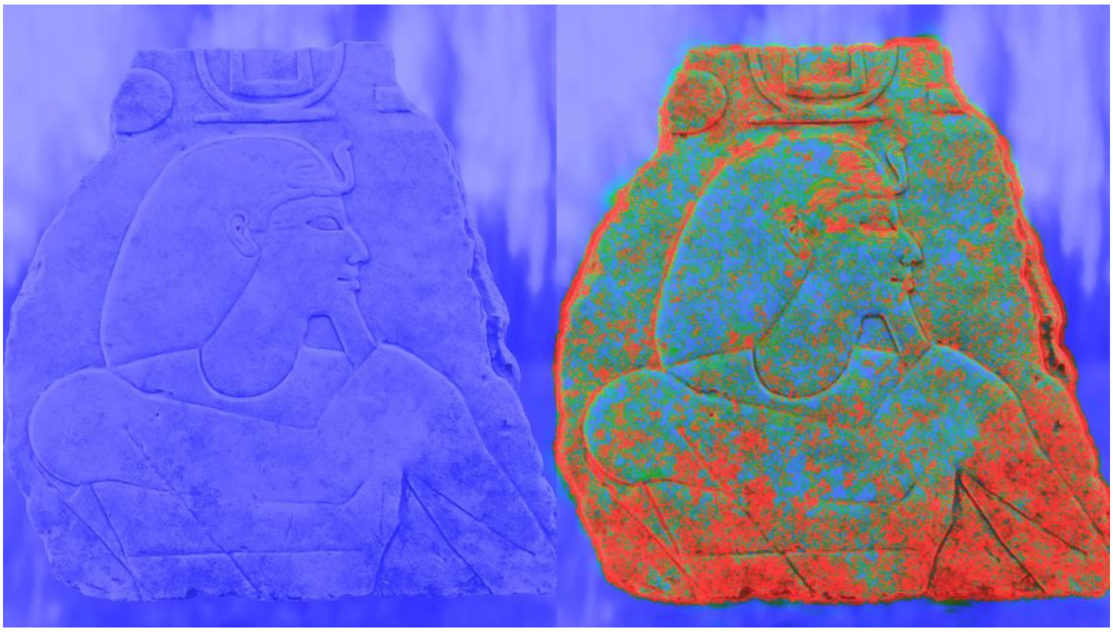

Figure 167: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

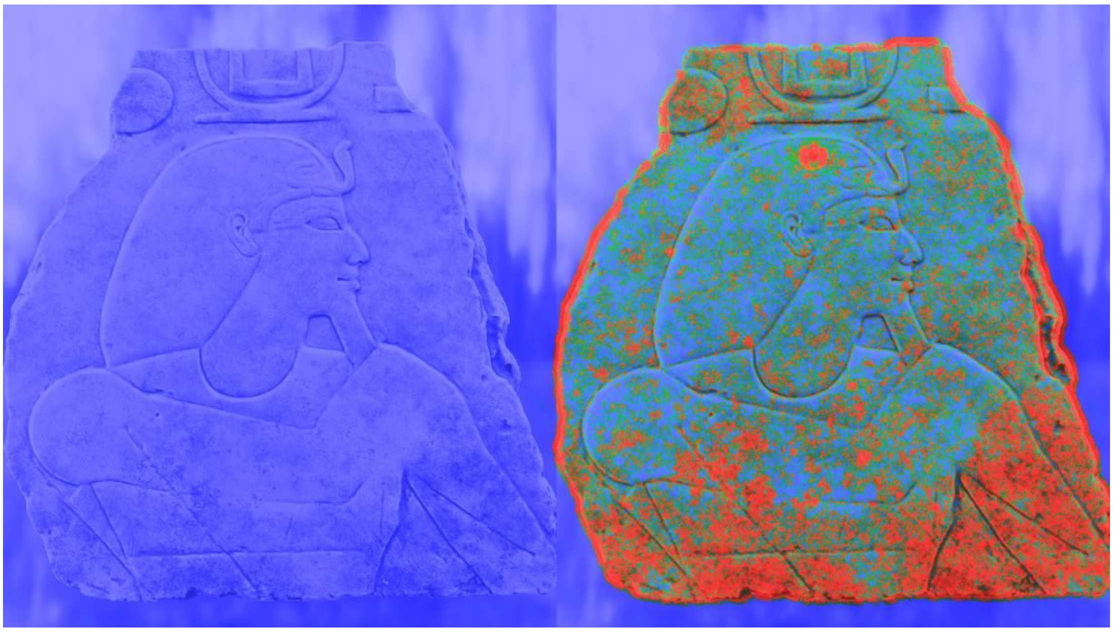

Figure 168: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

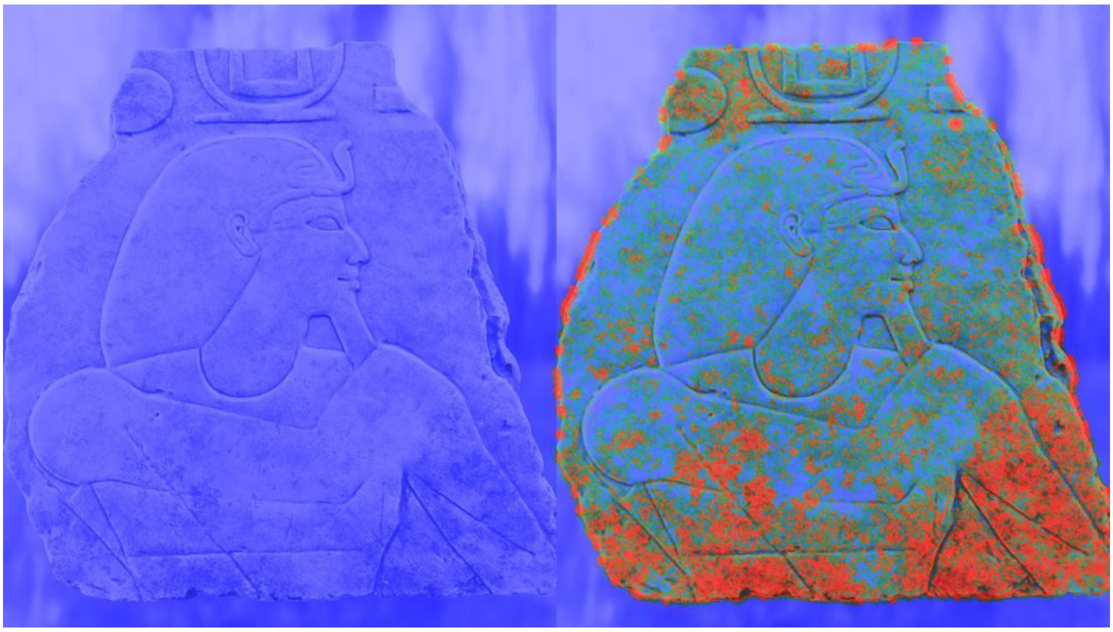

Figure 169: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

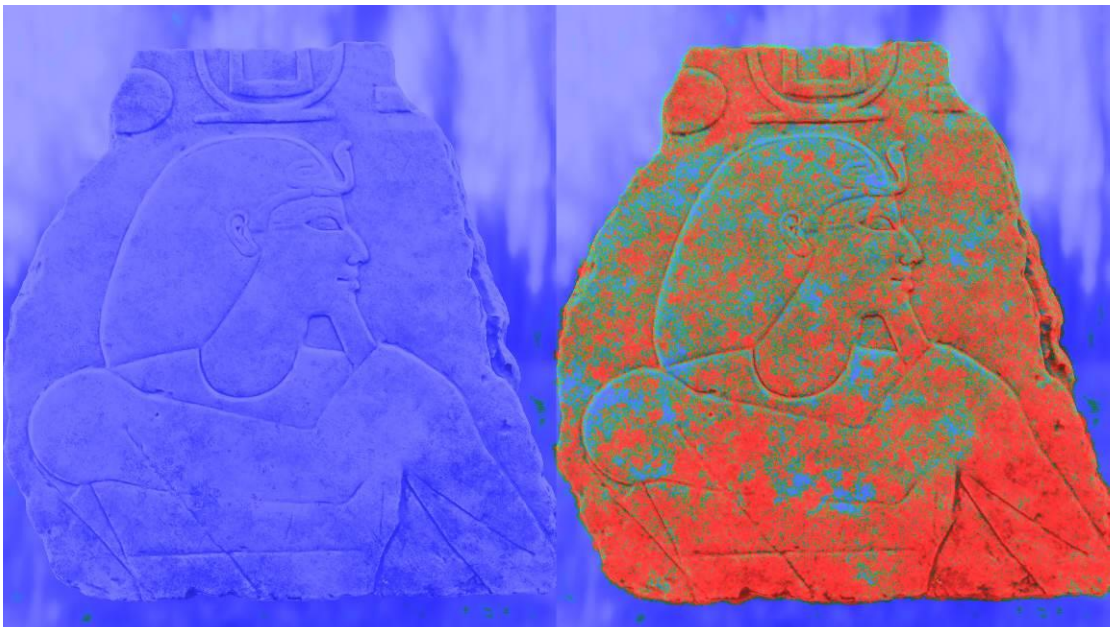

Figure 170: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

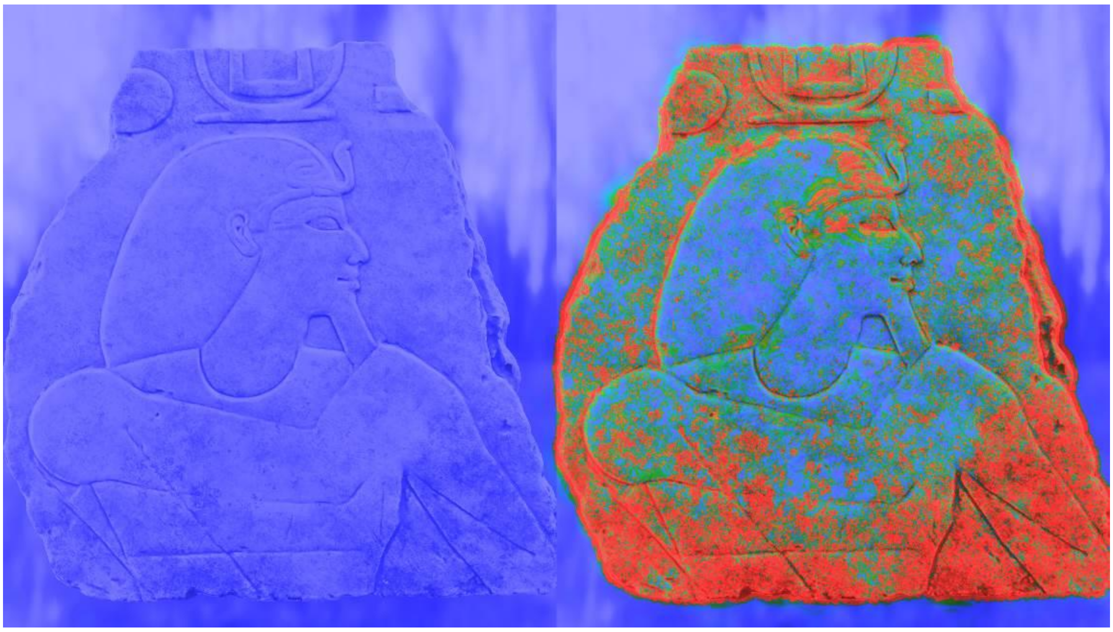

Figure 171: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

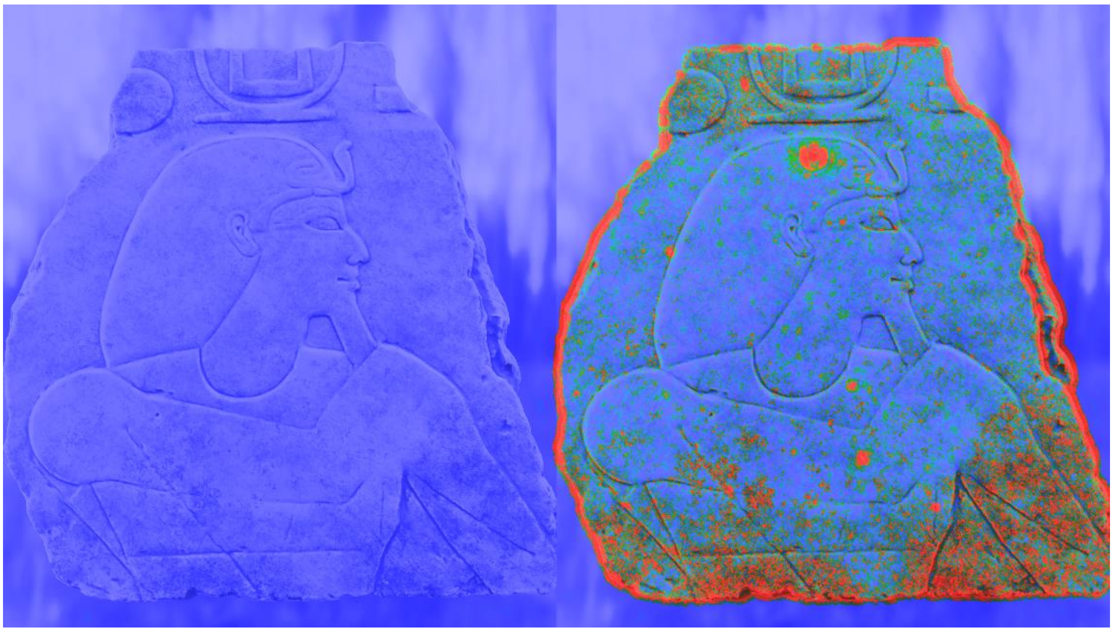

Figure 172: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

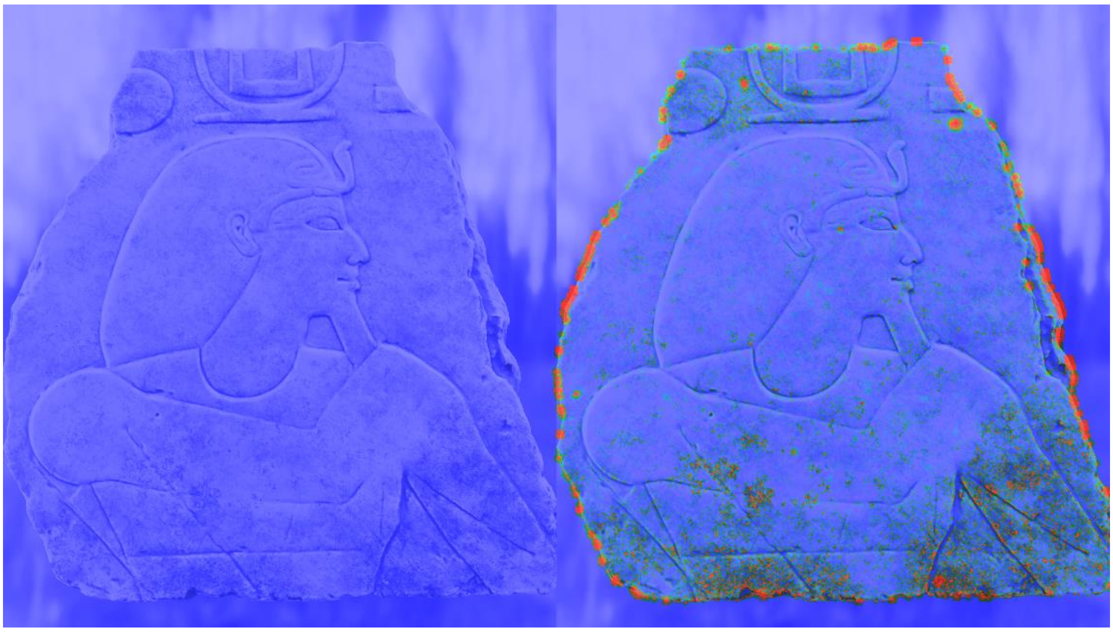

Figure 173: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

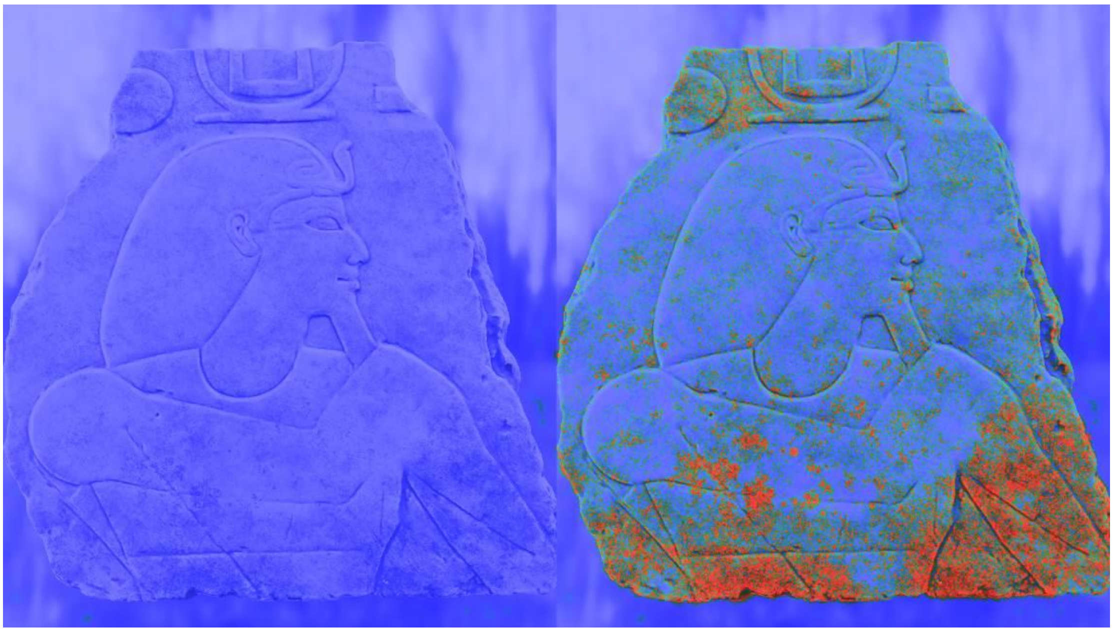

Figure 174: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

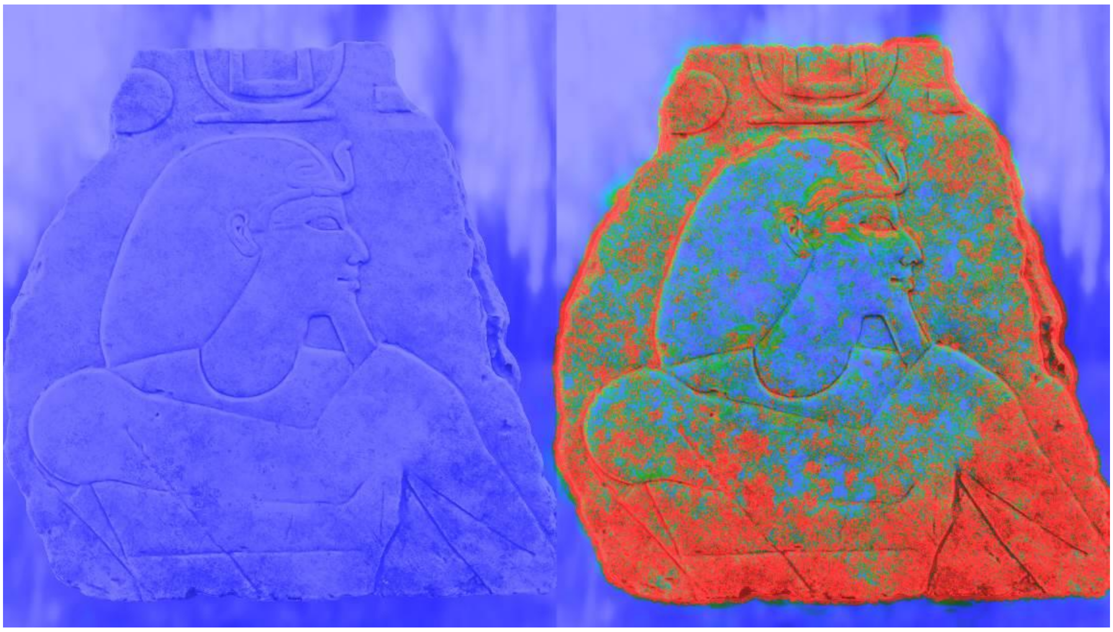

Figure 175: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

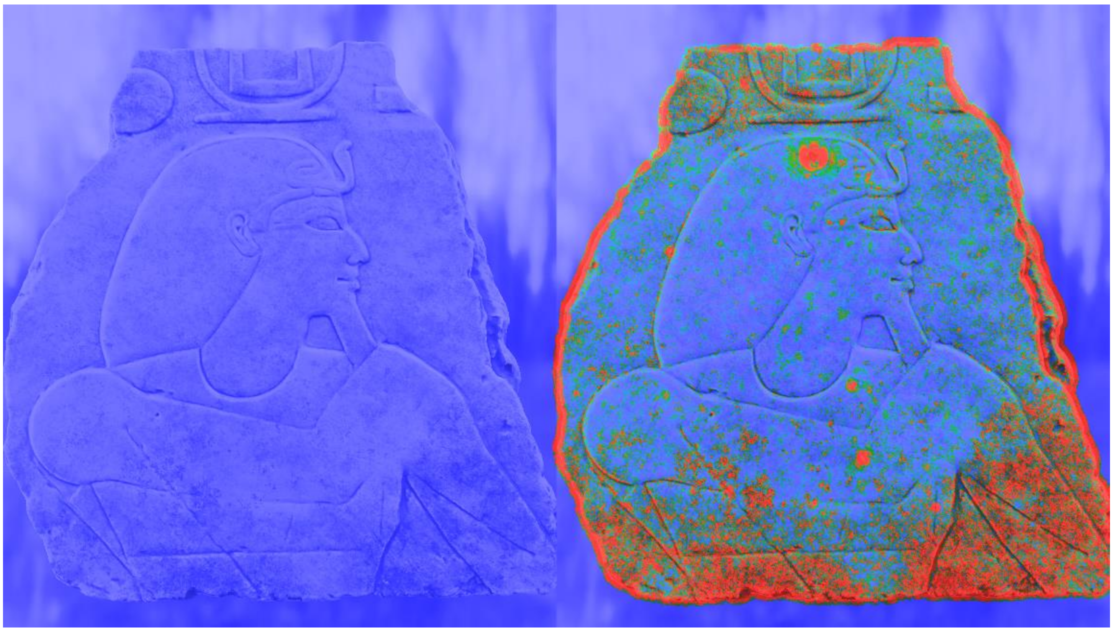

Figure 176: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

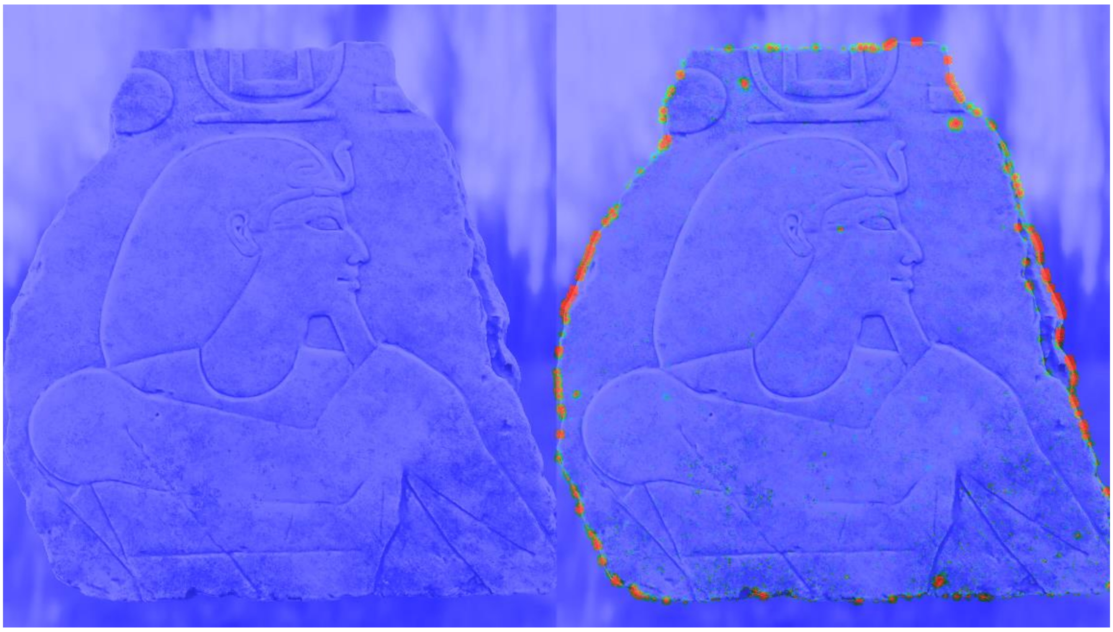

Figure 177: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 178: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 179: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 180: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 181: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 182: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 183: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 184: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 185: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 186: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 187: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 188: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 189: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 190: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 191: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 192: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Zeus Ammon

Rendered images of the different stimuli at different texture and polygonal resolutions

Figure 193: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 194: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 195: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 196: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 197: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 198: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 199: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 200: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 201: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 202: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 203: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 204: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 205: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 206: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 207: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Figure 208: Reference model with stimuli at 100% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

Figure 209: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 210: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 211: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 212: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 213: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 214: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 215: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 216: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 217: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 218: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 219: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 220: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 221: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 222: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 223: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

Figure 224: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 225: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 226: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 227: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 228: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 229: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 230: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 231: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 232: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 233: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 234: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 235: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 236: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 237: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 238: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 239: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 240: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 241: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 242: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 243: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 244: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 245: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 246: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 247: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 248: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 249: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 250: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 251: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 252: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 253: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 254: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 255: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 256: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 257: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 258: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 259: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 260: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 261: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 262: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 263: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 264: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 265: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 266: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 267: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 268: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Shakespeare Bust

Rendered images of the different stimuli at different texture and polygonal resolutions

Figure 269: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 270: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 271: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 272: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 273: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 274: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 275: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 276: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 277: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 278: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 279: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 280: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 281: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 282: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 283: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Figure 284: Reference model with stimuli at 100% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

Figure 285: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 286: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 287: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 288: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 289: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 290: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 291: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 292: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 293: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 294: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 295: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 296: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 297: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 298: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 299: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the RBG.BT 709 colour space

Figure 300: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 301: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 302: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 303: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 304: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 305: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 306: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 307: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 308: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 309: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 310: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 311: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 312: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 313: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 314: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric code with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 315: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 316: Reference model with stimuli at 40% polygonal resolution and no texture

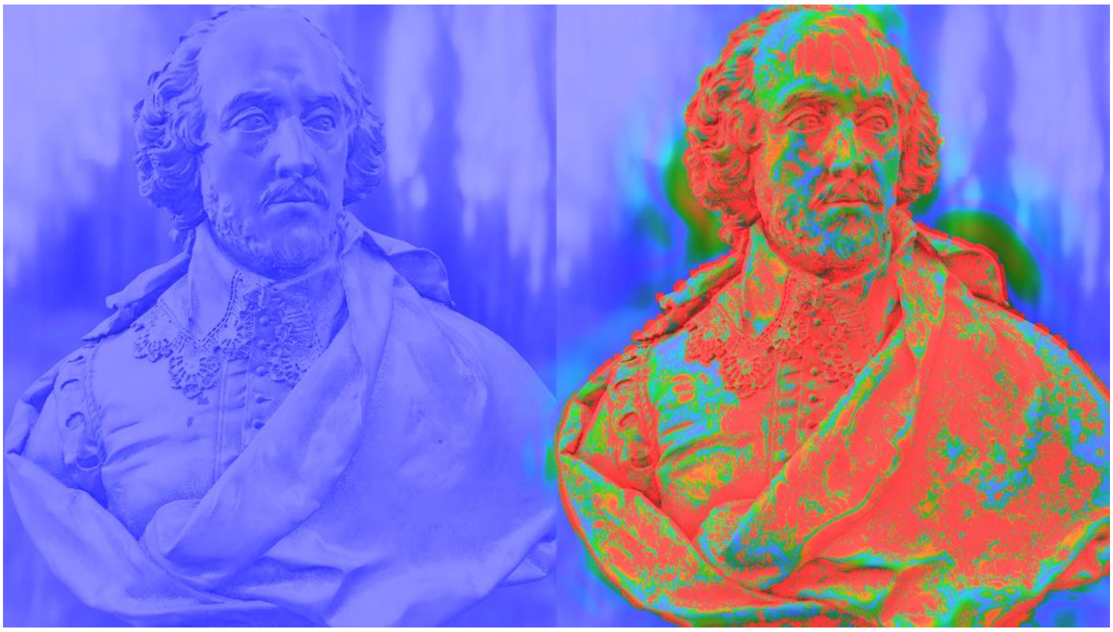

Figure 317: Reference model with stimuli at 70% polygonal resolution and no texture

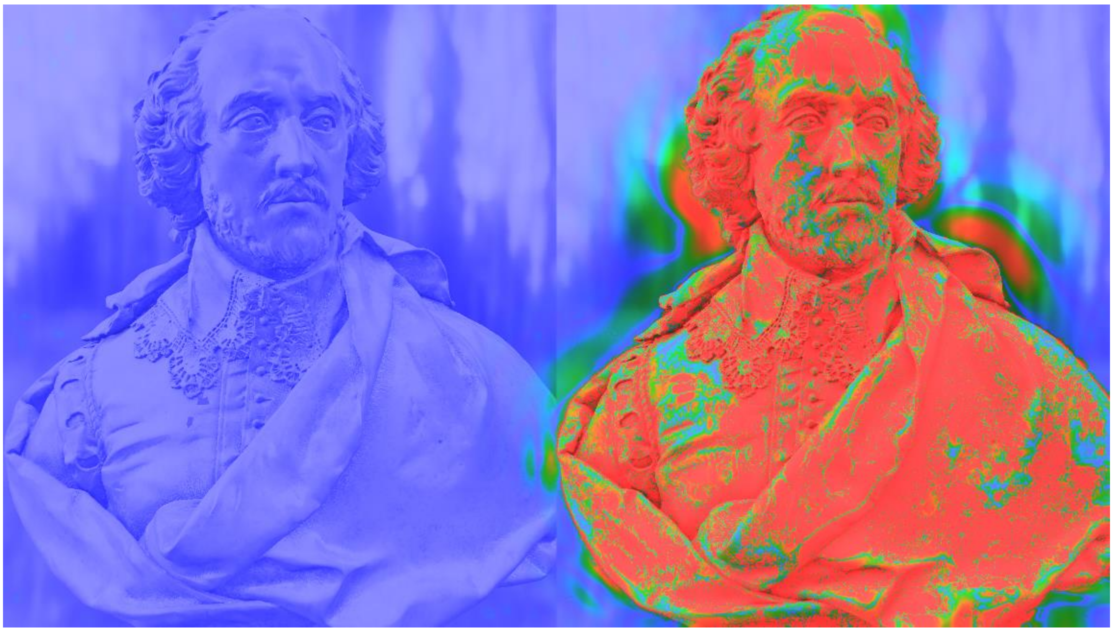

Figure 318: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 319: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 320: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 321: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 322: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 323: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 324: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 325: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 326: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 327: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 328: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 329: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

Results of the HDR-VDP2 image metric with calculated differences with the stimuli at different polygonal and texture resolution in the sRGB colour space

Figure 330: Reference model with stimuli at 10% polygonal resolution and no texture

Figure 331: Reference model with stimuli at 40% polygonal resolution and no texture

Figure 332: Reference model with stimuli at 70% polygonal resolution and no texture

Figure 333: Reference model with stimuli at 100% polygonal resolution and no texture

Figure 334: Reference model with stimuli at 10% polygonal resolution and 512x512px texture resolution

Figure 335: Reference model with stimuli at 40% polygonal resolution and 512x512px texture resolution

Figure 336: Reference model with stimuli at 70% polygonal resolution and 512x512px texture resolution

Figure 337: Reference model with stimuli at 100% polygonal resolution and 512x512px texture resolution

Figure 338: Reference model with stimuli at 10% polygonal resolution and 1024x1024px texture resolution

Figure 339: Reference model with stimuli at 40% polygonal resolution and 1024x1024px texture resolution

Figure 340: Reference model with stimuli at 70% polygonal resolution and 1024x1024px texture resolution

Figure 341: Reference model with stimuli at 100% polygonal resolution and 1024x1024px texture resolution

Figure 342: Reference model with stimuli at 10% polygonal resolution and 2048x2048px texture resolution

Figure 343: Reference model with stimuli at 40% polygonal resolution and 2048x2048px texture resolution

Figure 344: Reference model with stimuli at 70% polygonal resolution and 2048x2048px texture resolution

G Subjective Questionnaire Results

The following contains the results of the questionnaires conducted as part of the subjective test. They contain information for each question, age, sex, and experience with 3D graphic and other information.

G: Shakespeare Bust

Age	Gender	Experience with 3D	Score	Texture Material	Importance of Texture	Want textured & untextured	Would you prefer orginal, replica or Digital 3D	Want to learn more	Comments
41	Male	None	6	Clay	Sort of Important	yes	both	yes	
28	Female	None	7	Marble	Very	No	3D Model	yes	
35	Female	None	6	Clay	Quite Important	yes	Original	no	
26	Male	none	7	Stone	Very	yes	Original	yes	
32	Male	None	6	Stone	Kind of	yes	both	no	
20	Male	N/a	7	Clay	very	yes	3D model	yes	
28	Male	None	6	Clay	Sort of	yes	Both	no	
18	Male	Yes	8	clay	very	yes	3D Model	yes	
22	Male	This is new for me	7	Clay	Very - Excellent details visibl	Not important	Original/replica if possible, excellent if not available	yes	
25	Female	None	6	clay	quite	yes	original	no	
35	Male	None	7	clay	Sort of	no	original	no	
23	Female	Low, I prefer original artefacts	5	clay	sort of	no	real thing	no	
21	Male	medium	8	clay	very	yes	3D Model	yes	
41	Female	none	7	Stone/marble	quite	yes	3D (cos its new)	I guess	
18	Female	None	6	Clay	very	yes	3D Model	no	
22	Female	None	6	Clay	quite	yes	original/replica	no	
29	Female	N/a	8	Stone	quite	yes	3D model	yes	
27	Male	Some	6	Clay	Very	yes	3D model	no	
26	Male	None	6	Marble	very important	Always	Original	Yes	
21	Female	Little	8	Clay	very	yes	3D model	yes	
		Mean	6.65						
		Standard Deviation Sample	0.875094						
		Standard Deviation	0.875094						
		Sample Size	20						
		Confidence Coeff	2.093						
		Margin of Error	0.409552						
		Upper Bound	7.059552						
		Lower Bound	6.240448						
		Max	8						
		Min	5						
		Range	3						

G: Anglo Saxon Brooch

Age	Gender	Experience with 3D	Score	Texture Material	Importance of Texture	Want textured & untextured	Would you prefer orginal, replica or Digital 3D	Want to learn more	Comments
27	Female	None	6	Gold	Very important	Yes	The actual object	yes	
22	Female	None	8	Copper	Very improtant	yes	The actual object	yes	
36	Female	Very vague	10	Bronze painted Gold	Very important	yes	3D replica	yes	
21	Female	Little	5	Gold or Bronze	10, Very important	yes	Original	yes	
57	Male	Competiant with PC, none with 3D	7	Copper	It is very useful	Yes	Original	yes	
23	Male	Very Good	6	Gold	Very Important	yes	Original and Digital	No	
24	Female	None	8	Gold	yes very	yes	both	yes	
26	Male	None	10	Bronze	Very Important	With texture only	The original	No	
42	Female	None	7	Gilded Bronze	Quite Important	Yes	Both	yes	
42	Male	Yes	9	Gold	It is important	Yes	3D model	yes	Very Interesting
52	Male	Not much	10	Copper	Not important	No	Actual origina	yes	
60	Male	Some	6	Gold	Very Important	yes	both	yes	
54	Male	Some	4	Gold plated metal	Very Important	Very	original	yes	
29	Male	Some	8	Bronze	Very important, but like it wit	Yes	Replica and digital	yes	
24	Female	None	7	Copper/bronze	Very Important	yes	Both	yes	
		Mean	7.4						
		Standard Deviation Sample	1.843909						
		Standard Deviation	1.843909						
		Sample Size	15						
		Confidence Coeff	2.145						
		Margin of Error	1.021224						
		Upper Bound	8.421224						
		Lower Bound	6.378776						
		Max	10						
		Min	4						
		Range	6						

G: Egyptian Relief

Age	Gender	Experience with 3D	Score	Texture Material	Importance of Texture	Want textured & untextured	Would you prefer orginal, replica or Digital 3D	Want to learn more	Comments
19	Male	Ipads, Laptops, Computers	7	Stone/Sandstone/Lim	Very	yes	3D model	yes	
22	Female	None	6	Stone	Very	yes	3D model	yes	
35	Female	Some	7	Stone	Sort of Important	yes	3D model	no	
18	Female	None	8	Stone	Quite	yes	original	yes	
36	Male	None	7	Plaster	Quite	yes	3D model	no	
24	Female	None	6	Stone	Quite	yes	original	no	
26	Female	None	8	Stone	Sort of Important	yes	3D model	no	
32	Female	None	5	Stone	Quite	yes	Original	no	
25	Female	None	6	Stone	Very	Yes	3D model	no	
23	Female	None	6	Plaster	quite	yes	original	yes	
20	Male	Some	8	Stone	Very	yes	3D model	yes	
34	Male	None	7	Stone	Quite	yes	3D model	yes	
26	Female	Not a lot	5	Stone	Very	yes	original	yes	
30	Male	Very little	6	Plaster	quite	Yes	original	yes	
21	Female	None	7	Limestone	Quite Important	yes	original	yes	
		Mean	6.6						
		Standard Deviation Sample	0.985611						
		Standard Deviation	0.985611						
		Sample Size	15						
		Confidence Coeff	2.145						
		Margin of Error	0.545867						
		Upper Bound	7.145867						
		Lower Bound	6.054133						
		Max	8						
		Min	5						
		Range	3						

G: Zeus Ammon

Age	Gender	Experience with 3D	Score	Texture Material	Importance of Texture	Want textured & untextured	Would you prefer orginal, replica or Digital 3D	Want to learn more	Comments
26	Male	none	6	Bronze	Sort of Important	yes	both	No	
27	Male	Some experience	5	Metal	Quite	yes	original	no	
24	Female	None	8	copper	quite	yes	both	yes	
32	Female	None	8	Bronze	Quite	yes	Original	yes	
33	Female	Some	2	Copper	Very Important	Yes	The original	yes because I am a nerd	Why are there no nu
38	Male	No	8	Bronze	Very Important	yes	3D model	yes	
36	Female	Very good	6	Metal	Quite	yes	both	no	
47	Male	none	6	Metal	Very Important	yes	3D	yes	
40	Male	Very good	7	Iron	Very	yes	3D model	yes	
24	Female	None	5	bronze	sort of	yes	original	no	
22	Female	None	5	Metal	quite important	yes	both	no	
32	Male	None	7	Bronze	Sort of	yes	both	yes	
18	Female	none	6	Metal	quite	yes	both	no	
18	male	none	6	Stone	sort of	yes	both	no	
18	Female	Little	8	Bronze	Very	yes	3D model	yes	no
25	Male	None	4	Metal	Quite	yes	Original	no	
21	Female	Little	3	Bronze	very	yes	original	yes	
32	male	none	8	Stone? Maybe some	Very Important, it shows dep	No	Most likely the real thing but replications are impor	Yes	
20	Male	None	7	Metal	Sort of	yes	Original	no	no
54	Male	A little	6	Bronze	quite important	yes	both	yes	n/a
		Mean	6.05						
		Standard Deviation Sample	1.700619						
		Standard Deviation	1.700619						
		Sample Size	20						
		Confidence Coeff	2.093						
		Margin of Error	0.795905						
		Upper Bound	6.845905						
		Lower Bound	5.254095						
		Max	8						
		Min	2						
		Range	6						