Towards Enhancing Virtual Museums by Contextualizing Art through Interactive Visualizations

CHRISTOFER MEINECKE, Leipzig University, Germany
CHRIS HALL, Chris Hall Design, Denmark
STEFAN JÄNICKE, University of Southern Denmark, Denmark

In response to the COVID-19 pandemic, public spaces such as museums and art galleries are experiencing increased demands to offer virtual online access. While current solutions seek to replace or augment a real visit, online tours often suffer from being too passive and lack in-depth interactivity to keep virtual visitors meaningfully engaged with an exhibition. Museums and art galleries seeking to broaden and engage their audience more deeply should offer intriguing experiences that invite the visitor to explore, to be entertained, and to learn by interacting with the content. We propose a novel virtual museum experience that utilizes multiple visualizations to contextualize a gallery's digitized artworks with related artworks from large image archives. We make use of the WikiArt data set that includes more than 200,000 images and offers diverse metadata used for comparative visual exploration. In addition, we apply machine learning methods to extract multifaceted information about the objects detected in the images and to compute similarities across them. Visitors of our virtual museum can interactively explore the artworks using different search filters such as artist, style, or object classes detected within an image. The results are displayed through interactive visualizations offering different perspectives on artwork collections, leading to serendipitous discoveries and stimulating new insights. The utility of our concept was confirmed by an informal evaluation with virtual museum visitors.

CCS Concepts: • **Applied computing → Fine arts**; **Digital libraries and archives**; • **Computing methodologies → Computer graphics**; **Computer vision**.

Additional Key Words and Phrases: Visualization in the Humanities, Human-Computer interfaces for virtual and digital museums, Analytic tools to assist research on collections or artefacts, Digital Art History

ACM Reference Format:
Christofer Meinecke, Chris Hall, and Stefan Jänicke. 2022. Towards Enhancing Virtual Museums by Contextualizing Art through Interactive Visualizations. *ACM J. Comput. Cult. Herit.* 1, 1, Article 1 (January 2022), 27 pages. https://doi.org/10.1145/3527619

1 INTRODUCTION

As public institutions seek to attract as broad an audience as possible, there is an imperative to cater for the growing virtual online audience alongside regular visitors. This has been magnified by the closure of public institutions across the world during the COVID-19 pandemic, meaning that virtual online access has remained the only option for domestic and international visitors [5, 79]. Traditional art gallery exhibitions are typically limited to a finite number of works, arranged in a fixed order and accessed by visitors walking around within a physical space. The range of artworks shown can be further limited by considerations such as the fragility of the

Authors' addresses: Christofer Meinecke, cmeinecke@informatik.uni-leipzig.de, Leipzig University, Leipzig, Germany; Chris Hall, chris.impala@gmail.com, Chris Hall Design, Odense, Denmark; Stefan Jänicke, stjaenicke@imada.sdu.dk, University of Southern Denmark, Odense, Denmark.

2476-1249/2022/1-ART1 $15.00
https://doi.org/10.1145/3527619

items, visitor fatigue and the availability of items from existing collections or on loan from other institutions. These limits often prevent an institution from displaying the ideal content to support a given theme.

A virtual museum can be seen as an extension or a complement to a physical museum which removes some of these constraints by providing a solution to the space limitations of a physical exhibition and by giving a safe alternative to displaying and investigating fragile objects. Ideally, the design of a virtual museum should go beyond presenting only the content and information which the museum has digitally available from its own collections [1, 66]. "The term virtual museum has been defined as a logically related collection of digital objects composed in a variety of media which, because of its capacity to provide connectedness and various points of access, lends itself to transcending traditional methods of communicating and interacting with visitors; it has no real place or space, its objects and the related information can be disseminated all over the world" [66]. This core idea of the virtual museum goes back to André Malrauxs project the Imaginary Museum [53, 73] that he described as a museum without walls, which was a montage of photos from all around the world and from different time ranges.

While existing types of virtual experiences seek to replace or complement a real visit, online tours often suffer from being too passive and lack in-depth interactivity to keep virtual visitors meaningfully engaged with the content for more than a short time. We propose a virtual museum tour enhanced by direct access to various selected visualizations that contextualize the artworks within the gallery. Our approach allows the virtual visitor to explore and compare paintings using machine learning methods and visual interfaces which arrange related artworks as an array of icons. Following the concept of generous interfaces using multiple representations to reveal the complexity of the cultural collection [86].

This serves the visitor with diverse entry points to discover the underlying art collection and so to promote serendipitous discoveries [63]. In the context of our virtual exhibition, we offer different arrangements of art and visualization to give visitors a deeper and more fulfilling museum tour experience. Our intent is to encourage visitors to explore more of the virtual museum's exhibits and the latent space that the artworks are projected on. So that visitors are able to explore collections or exhibitions with individually targeted activities including searching, selecting, and comparing artworks.

We conducted an informal evaluation with 61 participants from different backgrounds to evaluate the concept of a virtual museum in a 3D environment combined with information visualization principles that contextualize the artworks. Logging all activities during the virtual museum visits gave us interesting insights concerning the visitors' adaptability to the visual interfaces, how long they observed the diverse visual depictions, their movement patterns inside the virtual space, and, in general, their interest in art. In summary, the contributions of this paper are:

- **A Virtual Museum Model & Tour** that connects the virtual version of a 'real' museum gallery with an adjoining exploration room, where the visitor is enabled to regard a painting in the context of a large image archive. Visitors are engaged in exploration bearing on machine learning methods to generate diverse perspectives on art collections through interactive visualizations which also allow serendipitous discoveries.
- **An Evaluation** that informs on behavioral aspects of museum visitors and their characteristics in a 3D space, delivering valuable insights regarding the adaptability of the virtual museum concept and for future improvements of virtual museums.

Our evaluation revealed that not only did it appeal to the general public, our approach also seems to serve the increasing desires of humanities scholars to quantitatively analyze art collections. Although the use of visualization methods in art history can be traced back to Aby Warburg's Mnemosyne Atlas in the 19th century, digital methods are rarely available [82]. Or as Drucker stated it "To date no research breakthrough has made the field of art history feel its fundamental approaches, tenets of belief, or methods are altered by digital work" [20].

Nevertheless, computational methods, e.g., machine learning and visualization can support working practices and can give new insights into the objects of interest and therefore help answer research questions. Furthermore, virtual museums could serve as an interdisciplinary object of investigation and extend conventional museum space by providing enhanced visitor experiences in terms of engagement and attraction [5].

2 RELATED WORKS

Our work focuses on virtual museums and more generally, visualizations of cultural heritage collections [87]. In particular, visualizing images and the objects contained within them, facilitated by neural networks and other machine learning methods.

2.1 Visualization of Cultural Heritage Collections

Images are the most used data type for visualizing Cultural Heritage Collections [87] as cultural objects are commonly digitized as images. Most of the time the collection is presented in an exploratory way by giving an overview of an image or document collection. For this, treemap and cluster algorithm are a popular method [3, 9, 81]. Furthermore, timelines [29, 30] are a common method to show the historical contexts of the artifacts, while maps are used to show the geographical context [21, 52]. Similar to us, other approaches combine multiple visualizations showing different facets of cultural heritage collections to give manifold perspectives on the objects of interest [6, 19, 32]. Cole et al. [11] showed how "to highlight the relationships between objects and their features within digital art collections and provide a means for visitors to explore these collections via interactive, narrated pathways", but they did not include interactive visualizations. Crissaff et al. [12] created a system inspired by traditionally used lightboxes to explore, compare, organize and annotate art image collections. Junginger et al. [43] displayed objects contained in photographic plates in a Close-Up cloud similar to our approach . The object categories can be further investigated and are in contrast to our approach manually annotated. Similar to us, the Bohemian Bookshelf by Thudt et al. [75] follows the serendipity principle to explore a collection of digital books. The visualizations give multiple access points to the data e.g. the book cover, author, or time.

Despite the wider application of machine learning methods on cultural heritage collections [23], they are rarely applied together with visualizations to allow further perspectives into the collections. When machine learning methods are applied to a collection they are mostly used to project the images onto a 2D space. Pflüger and Ertl [61] proposed a visualization system for large image sets based on clustering and projection methods. Crockett [13] plotted high dimensional clusters of images through dimensional reduction and arranged slices of images in histograms based on visual and non-visual features. Hochmann and Manovich [36] plotted photos from social media and based on features like hue, brightness and upload time. Similarly, Hristova [37] compared art from Aby Warburg's Mnemosyne and Yamaoka et al. [89] Time Magazine covers, Manga publications and paintings. In contrast to them, Strezoski et al. [69–71] presented multiple approaches to visualize art using the Omniart data set [72]. TindArt [69] an art recommender system where users are presented with art that they can then like or dislike. The image embeddings are also mapped to a 2D space in order to inspect the areas where likes and dislikes occurred and thus communicate the reasoning of the recommendation. ArtSight [70] is a query by color exploratory interface for images. The Art, Color and Emotion browser [71] visualizes the dominant sentiment and color in a specific time interval.

2.2 Virtual Museums

In recent years, different types of virtual museums have been proposed to display cultural heritage artifacts. Kabassi [44] evaluated the state of the art of museum websites including 3D environments and mobile apps, while Bekele et al. [4] surveyed AR, VR and Mixed Reality technologies for Cultural Heritage collections. Walczak et al. [80] presented the ARCO system to explore a virtual museum with AR and VR. Huang et al. [38] presented an augmented panorama approach. Lugrin et al. [50] presented a location-based virtual museum for over 100 users. For 3D artifacts, Loscos et al. [49] present a virtual museum where users can touch statues with haptic feedback and the Atalaya3D project [58] created 3D scans of sculptures and historical sites that can be visited in a 3D environment. Carvajal et al. [10] created a virtual museum based on 3D image acquisition and 3D modeling.

Some works focus on the creation of a personalised virtual museum. For this, Hayashi et al. [34] present an approach to automatically generate a virtual museum based on user's bookmarks, while VIRTUE [28] allows users to create and navigate within their own exhibition of 2D and 3D objects. DynaMus [45] is a 3D virtual museum framework to create virtual exhibitions and Liarokapis et al. [48] present a visualization framework to visualize cultural heritage artifacts in a virtual museum. All these approaches focus on the recreation of a museum in a virtual environment to present objects of interest but without visualizing additional information about them. Therefore, we see an opportunity to enhance virtual museums using visualization methods that contextualize the objects of interest and give new perspectives on cultural heritage collections by allowing visual exploration.

2.3 Art Classification and Object Detection

Various works focus on detecting objects in artworks using Convolutional Neural Networks [31, 83]. Previous works by Crowley et al. [14–16] showed that classifiers trained on natural images can retrieve paintings containing the selected category. Garcia and Vogiatzis [27] presented SemArt, an art training set for semantic understanding, while Strezoski and Worring [72] presented baselines for multiple art classification tasks on the Omniart data set. Wevers and Smiths [84, 85] applied neural networks to historic images to find similar images and to observe trends in the data. Arnold and Tilton [2] proposed a theoretical framework for large collections of visual material and applied object detection and image similarity measurements. Garcia et al. [26] computed context-aware embeddings of images using image features and meta-data presented in a knowledge graph. For medieval images, computer vision algorithms were applied for pattern spotting [78], to detect and classify crowns [90] or gestures [65] and for general similarity search [47]. Shen et al. proposed a method of finding near-duplicate patterns in images [67] But none of these works applied interactive visualizations to explore the results. Mao et al. presented DeepArt Search an art retrieval and annotation system trained on the Art500K data set [54], where users can upload their own images. Similar to us, they allow users to find similar images but they do not provide further contextualization methods.

3 VIRTUAL MUSEUM PROJECT

The project started with the idea of applying the concept of quantitative text analysis to image data, which had been introduced in 2019 by the term *Distant Viewing* [2]. The two visualization scholars co-authoring this paper started prototyping a series of 2D visualizations, each of which generated a new, quantitative perspective on art collections. Our team was later complemented by an expert in museum exhibition design, who, in collaboration, designed a virtual museum concept that made it possible to make the multiple exploration abilities of our visualizations accessible to the general public. We began by investigating current virtual museum solutions.

Practical Examples. Existing virtual online museum experiences can be categorized by two broad platform types: 360 degree Photo-stitched based methods like 'Google Arts and Culture's' Street View and those based on 3D scanned-models like Matterport. Google Arts and Culture tours are an extension of Google Maps Street View and offer the same functionality familiar to a global user group. A practical example is SMK-Statens Museum for Kunst in Copenhagen [68]. Virtual versions of the museum's interior are created using a series of photo-stitched 360 degree panoramas which gives an impression of being in a 3D environment but also caused distracting warping of the images as the user advances through the space. The 360-degree photography provides a series of fixed viewpoints that are geo-located within a building. A user can explore the spaces by scrolling, panning and clicking on navigation targets which are arrows or crosses on the ground. Some institutions offer additional clickable targets placed within the space which trigger a sidebar showing additional images and text or links. A shortcoming is that visitors do not have the freedom of navigating a 3D model as they are limited by the original Google street-view 360 degree viewpoints. Also, the tour is free-form, unstructured and the resolution of the images is too low to allow visitors to read label texts within the galleries. A practical example of a Matterport tour is The National Museum of History and Art in Luxembourg [60]. Matterport tours are created by 3D scanning of an institution's spaces using a range of different cameras combined with Matterport's proprietary service which provides 3D explorable photographic models and plans. This approach to virtual online tours offers high-quality images coupled with intuitive navigation via clickable circle targets within the spaces. There are three viewing modes, Doll's house view, plan and walkthrough. Doll's house view makes good use of double-sided surfaces with visible textures on the internal walls of the model and transparent textures externally. This allows internal detail to be seen and the model to be explored externally and entered from any point. These models offer a very responsive tour experience and allow the virtual visitor to zoom into small details and read nearly all levels of interpretive text. The Plan view is useful for switching locations but would benefit from further information when hovering the mouse. There is also a list of floors that allows the user to switch between levels of the museum. There are some clickable targets next to museum art and objects which reveal a small text label as well as a link to further information about that object on the museum's website.

Our Vision. Current virtual museum solutions mostly aim to recreate digital representations of already existing exhibitions. The major shortcoming is that they seldom take advantage of interactive visualization principles to give multi-perspective views on the objects of interest. In order to make virtual museums more interactive, we regarded them as an ideal environment in which to integrate our prototypical 2D visualizations for quantitative image analysis. Visualizations can further help in contextualizing the paintings and provide new information that is currently not supported by other virtual museums. Consequently, we wanted to find out if quantitative views are of interest for museum visitors and if they are seen as a valuable complement to real museums. Therefore, we formulated multiple abstract tasks and designed visualizations to fulfill the underlying information need.

Tasks. The exploration of the virtual museum can spark different questions in a visitor's mind. We formulated the following abstract tasks based on Munzner's task abstraction model [59]. First and foremost, a virtual museum's function is to *present* artifacts to an audience in an *enjoyable* way that empowers the visitor to *explore* the collection and therefore to *discover* new insights into art. In particular, the visualizations should allow the visitor to *discover* new images, styles, genres or artists. Furthermore, a virtual museum should allow the visitor to *query* for particular paintings or artists. Referencing the large data set allows the visitor to easily *discover* and *identify* similar images. When focusing on an artist of interest, a *summary* of the artist's development over time, their oeuvre, can be observed. The ability to *compare* images and artists with each other is also easier on a large scale through digital tools like a virtual museum than it would be in a real museum. We also allow interactions to perform different *search* tasks. Sometimes visitors want to *lookup* a specific image or an artist of interest. Other times they might want to interact with the collection freely and *explore* it, resulting also in serendipitous findings. The concrete tasks for each wall are explained together with the design in Section 4.2.

Data. We combined the paintings from Wikimedia and the WikiArt corpus. The WikiArt corpus consists of over 180,000 paintings from over 3,600 different artists ranging over 199 different styles. It is one of the largest publicly available digital art data sets and includes information about the creation year of the painting, the genre, the media used, the location and the series it is a part of. There are also some manually annotated tags for the content of the images and sometimes descriptions. The Wikimedia paintings set consists of around 20,000 images with over 1,400 different artists. Our total data set without duplicates contains around 200,000 images from over 4,800 artists. We applied machine learning methods in order to further contextualize this data set.

Object Detection and Pre-Processing. We applied the Faster R-CNN [62] trained on the Open ImagesV4 corpus [46] to the data. The Open Images data set consists of around 9 million images with different image-level annotations. For around 1.74 million images, annotations about object bounding boxes exist that were manually annotated. An image contains on average 8.4 bounding boxes, therefore, resulting in 14.6 million bounding boxes for the whole data set. For object classes, a hierarchy consisting of 600 different classes is used including parent and child class relations. The image feature extractor (backbone) of the network is an Inception Resnet V2 [74] trained on the ImageNet data set [17]. The ImageNet data set consists of over 14 million images and is therefore an appropriate training data set to create a generalized feature extractor. For each image in our data set, the Faster R-CNN predicts 100 bounding boxes with a confidence score between 0 and 1. For this, the region proposal part of the network proposes rectangular regions of interest and the detector part classifies these regions based on the object classes. This results in around 20 million bounding boxes, for which we only included bounding boxes with a confidence score of 0.5 or higher to reduce the errors in the image-level annotations of the data set. Our final set includes around 530,000 object bounding boxes. Furthermore, we use the top layer of the image feature extractor to compute image embeddings for each image and each bounding box in our data set. These image embeddings are vectors with a dimension of 1536 and are used to compute the nearest neighbors for the images and the bounding boxes. All image embeddings are added to a faiss [42] index structure, where similarity is based on the Euclidean distance between them. For each image, the k most similar images and for the bounding boxes the k most similar bounding boxes are queried by searching the index based on the vector of the image or bounding box. We also apply the nearest neighbor search to find duplicates in the two data sets. For this, we first find the nearest neighbors for each image and then we disregard all images by different artists. All neighbors with a Euclidean distance of 0 are treated as duplicates. For the remaining images, the titles are compared with string similarity to prevent removing a similar image by the same artist. The titles are cleaned from special characters and lower-case. When the titles are identical, we remove one of the images. For removed images, we aggregated the metadata of the two sources.

4 VIRTUAL MUSEUM DESIGN

We started with the creation of 2D visualizations of the paintings and their metadata. When considering the complexity of in-depth visual exploration of 2D paintings there is an advantage in presenting and organizing the results of the different interactions in an engaging way within a 3D exploration room. The objective is to provide an intuitive virtual experience that takes advantage of visitors' existing knowledge of navigating and exploring 3D spaces and also to make it easy to learn by exploring the space. It is proposed that the virtual visitor's spatial awareness makes it easier for them to visualize, comprehend and navigate a complex array of visual results in "the round" compared to displaying them on an infinite and more abstract 2D web page. We also see an advantage in the visitors' ability to conduct their exploration seamlessly back and forth between the gallery and a virtual exploration space resulting in longer visitation times with less chance of leaving the site.

As virtual museum tours are typically situated in authentic representations of real visitable places like museums or art galleries, we seek to extend the online experience by adding a virtual exploration space. Once a real building has been digitized in 3D the resulting model can be extended easily. In this way, a virtual portal to another space can be added as a new piece of the building either static or hidden and revealed by clicking on a target.

4.1 Exhibition Design

In order to evaluate this enhanced virtual museum tour, we created a working prototype consisting of a realistic simulation of a gallery space with an added exploration room accessed via a portal (Figure 1). We created a virtual exhibition for evaluation using 12 paintings selected from the WikiArt data set, from the period 1887-1939 and included the styles: Art Nouveau (Modern), Fauvism, Impressionism, Realism, Pointillism, Social Realism, Ukiyo-e

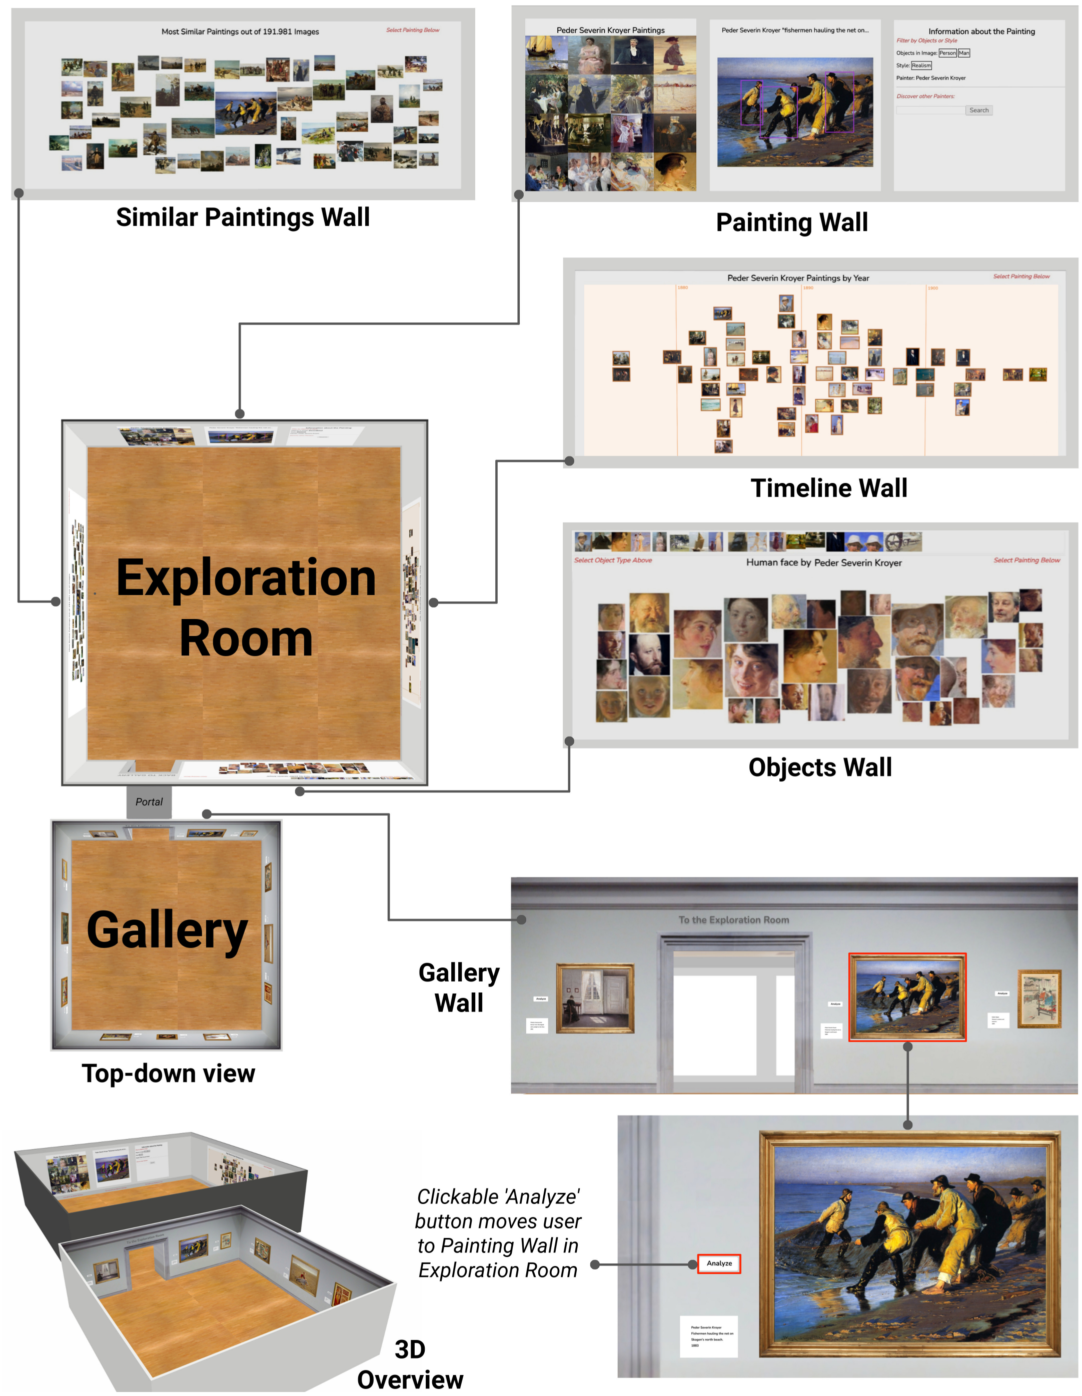

Fig. 1. The design of the virtual museum. Showing the visualizations on the different walls of the top-down view.

and Symbolism. In sizing the two spaces we took into account the number of paintings to be displayed, a natural angle of view, and the navigation by the visitor. The Gallery measures 12mx12m with a ceiling height of 4m, with space for 3 paintings on each wall and space on one wall for the portal to the exploration room. Each painting is set in a photo-realistic 3D frame with the centers at the visitor's eye-level and on the left-hand side of each painting is a label showing artist, the title and the year. The Exploration Room measures 18mx18m with a ceiling height of 5m. We found that a larger space is required so that the walls can accommodate the visualizations within the visitors' field of view.

The two rooms were rendered using photographic textures sampled from real buildings to give an impression of a gallery with realistic lighting in the same way other virtual tours represent real-world environments. The floor and ceiling in both spaces are the same to provide a feeling of continuity. The walls in the Gallery are sampled from a real well-lit gallery interior, however, those of the exploration room are stylized in plain grey-scale colors to give prominence to the visualizations. To facilitate navigation between the two spaces we added a way-finding signage above the portal. The final model is imported as a Collada file with the three.js library [18] to render it with WebGL in the browser. Furthermore, we add a CSS3D renderer to add interactive visualizations with d3 [7] to the 3D scene. This renderer applies 3D transformations based on the CSS transform property on the DOM elements. Each visualization is mapped to a wall of the exhibition. Visitors can move around in the virtual museum using the arrow keys or the WASD keys. Using a mouse or a touchpad they can rotate the camera by holding the left mouse click. Using the mouse-wheel or the touchpad scroll functionality they can zoom in and out. To prevent the visitors from walking through the floor or the ceiling, we restrict movements to the y-direction. Furthermore, we apply ray casting to prevent movements through walls. Visitors can interact with paintings and input components by clicking on them.

4.2 Virtual Museum Tour

The virtual museum displays artworks in different contexts based on specific attributes like objects contained in the image, artist, style, or year. This approach intends to extend the virtual experience, by exploring the objects of interest in a more faceted way, instead of purely present them to the audience. A visitor starts in the Gallery surrounded by artworks from different artists and various styles. The paintings we displayed in the Gallery can be seen in the Appendix. When clicking on an image of interest the visitor can choose to further analyze the image. For this, visitors are moved to the Exploration Room. The room consists of four walls, each one displaying interactive components containing visualizations to contextualize the artwork of the artist, either by similarity, creation year, depicted objects or other metadata. In the following, we describe a visit to the virtual museum based on the interactions of one participant in our evaluation. A video to follow the narrative can be seen in the Supplemental Material. The interactions of the visitor are followed by the concrete tasks that can be performed and the design rationales of the visualizations.

Painting Wall. Silke visits a virtual museum for the first time and has no prior experience with virtual 3D Environments. After observing the painting "Austria" by Alphonse Mucha, she decides to analyze it further. In the Exploration Room, the painting is shown together with some metadata and other paintings by Alphonse Mucha on the Painting Wall. After observing different paintings by the artist for a while, she searches for "Gustav Klimt". Now she is presented with "Adam and Eva (unfinished)". Next, she clicks on a bounding box in the image with the label "Woman".

Tasks. The main purpose of the Painting Wall is to display the painting of interest for enjoyment, but it also serves as a control panel to change the context i.e. artist, style, or depicted objects. The center of the wall enriches the image with information about the objects in it with their bounding boxes. The right-hand part of the wall can be used to lookup or search for a specific image, artist, or style. Furthermore, filters can be applied to find images that contain a specific class or multiple specific classes or have a specific style. The left-hand part of the wall gives an overview of some of the other images in the selected context and therefore a suggestion of other potential paintings of interest.

Design. The Painting Wall of the exhibition displays the currently selected image in the correct aspect ratio together with all of the WikiArt metadata about the image e.g. style, artist, year (Figure 1). When hovering over the image, the bounding boxes are shown with their class labels. Similarly, by hovering over the object filter, the bounding boxes for the specific objects are also displayed. In the metadata field, a visitor can also search for another artist to change the context of the room, or select objects contained in the image, or the style as a filter. The other paintings are also updated when the visitor interacts with the room e.g. search, filter or click on an image. By clicking on another image, that image is displayed on the Painting Wall and the context is changed if the painting is by a different artist. If a visitor is more interested in a particular style than an artist, a style can be used as the context. Further contexts are possible by using the objects contained in the image, the genre, the series the painting is part of, or manually annotated tags. For the evaluation prototype, we excluded the search for this additional metadata, to prevent the experience to become too complex [51].

Objects Wall. After selecting the bounding box, Silke sees the Picture Cloud of women painted by Gustav Klimt on the Objects Wall. She can now click on one Bounding Box of interest in the cloud or change the selected label to show another object class painted by Gustav Klimt.

Tasks. The Objects Wall gives insight into what classes of object an artist painted and which ones are the most frequent. Furthermore, it gives an overview of the variety of different depictions in each of these object classes.

Design. On the Objects Wall, the objects contained in the paintings are displayed in a Picture Cloud (Figure 1). At the top, the bounding box with the highest confidence is shown for each class. These are sorted by frequency in descending order and can be selected to change the content of the cloud to a specific class. The cloud shows all objects of the class for the given artist. This can give different perspectives on the paintings of the selected artist. The sizes of the images are based on the confidence of the bounding box and they are placed in descending order based on this value on an Archimedean spiral. When a style or an object class is selected on the Painting Wall, they are also applied to this Picture Cloud. Through this, objects of a specific style can be analyzed to see the broad range in a specific style. Furthermore, objects that are co-occurring with other objects can be found.

Similar Paintings Wall. Silke goes back to the Gallery to further look at the painting "Fishermen Hauling the net on Skagens north beach" by Peder Severin Krøyer. She decides to analyze the painting. In the Exploration Room, she chooses to look for similar paintings after looking at the image for some minutes. On the Similar Paintings Wall, the most similar paintings are displayed in a Picture Cloud. There she finds "Clam Diggers" by T. C. Steele and decides to further analyze this image. With the interactive option to take a closer look at similar images, she finds artists with whom she is not familiar.

Tasks. The Similar Paintings Wall supports serendipitous discoveries while exploring the data set of the virtual museum. Paintings similar to a painting of interest can be found and through this other similar artists that are not known by the visitor can also be found.

Design. The Similar Paintings Wall shows the most similar paintings of the currently selected painting (Figure 1), e.g., the k-nearest neighbors based on the Euclidean distance of the image embeddings of the paintings. The target painting is displayed at the center and the other paintings are scaled in size by the similarity (while preserving the aspect ratio) and placed on an Archimedean spiral around the painting. On mouse click, one of the nearest neighbors can be further analyzed on the Painting Wall. In contrast to the other visualizations, the nearest neighbor cloud does not just focus on one specific artist or style. This supports serendipitous discoveries of other paintings and new artists.

Timeline Wall. Later, Silke searches again for Gustav Klimt and decides to focus on his paintings on the timeline. In the timeline, she now focuses on some of his other works like "Girl with Long Hair, with a sketch for 'Nude Veritas'" and "The Virgin". It gives her an impression of the development and changes in Gustav Klimt's style and themes.

Tasks. The Timeline Wall gives visitors an overview of the development of an artist over time (Figure 1). It also helps in placing the painting of interest into the career of the artist. Furthermore, the development of different styles or genres can be observed and even the depiction of specific objects over time. In addition to selecting a specific attribute, it is also possible to analyze a specific time range.

Design. The Timeline Wall displays a set of images on a timeline using Timages [39]. For this, the context currently selected is used. This can be the artist, the style, or the objects that appear in the images. Through this, paintings of a specific style or a specific artist can be observed over time. The paintings are placed as thumbnails on a horizontal timeline by filling a polygonal region. While preserving the aspect ratios of the image sizes, the images are scaled and ordered decreasingly according to a custom relevance metric. Then, the thumbnails are horizontally placed as close as possible to the vertical center and the x-position that corresponds to the painting's year of origin. Varying sizes help to observe images with higher relevance to the topic of interest, e.g., by default they are scaled to a similar size, but the original image size, or other measurements are possible.

5 EVALUATION

In order to assess if our virtual museum concept would be accepted by the general public, we conducted an informal evaluation [77] suitable for the intended purpose. The evaluation took the form of an online experiment, which has been proven to deliver valuable results for studies in visualization [35]. As we designed our virtual museum in a participatory visual design approach assuring the adaptability of museum visitors, we did not conduct a pilot evaluation earlier. In order to mitigate the disadvantages compared to controlled laboratory studies [88], we requested participants to visit the virtual museum only once. In order to assure this condition, each participant had to register for a preferred time slot of two hours, in which they could freely decide when to enter and leave the museum, providing a suitable timeframe for casual exploration. We asked potential participants only if they would like to visit a virtual art museum. In order to observe how an enhanced 3D virtual art museum is perceived by general public visitors, we did not include any introduction to the project or any preview of our virtual museum concept or its visual interfaces. After finishing the virtual museum tour, the visitors were asked to fill in a questionnaire to reflect on the experience. Along with the logging data, we regard the provided feedback from different angles. The questionnaire can be found in the Appendix.

Participants. The setup as an online evaluation helped us reach a large number of interested participants. To get a heterogeneous group of participants for our virtual museum evaluation to be conducted during one week, we invited scientific staff of different universities and research institutes, students and acquaintances via mailing lists. 61 of the invited confirmed their willingness to participate. The virtual museum visitors, of which 24 were female and 37 were male, have varying backgrounds and belong to different age groups. An overview is given in Figure 2. The age group from 18 to 29 years included most participants, followed by 19 aged between 30 and 39, 9 between 40 and 50, and 7 virtual museum visitors older than 50 years. 28 visitors had a scientific background (7 with a background in humanities), 13 were students (4 with a background in humanities) and 20 were non-scientific, general museum visitors.

Logging. For each participant, we monitored all movements in the virtual museum, all their interactions and the actual duration of the virtual museum visit. For the movements, we saved the position after moving and the new rotation of the camera. Furthermore, we applied ray casting to compute the duration the participants looked at a specific image in the gallery or a wall with visualizations. For each interaction we saved the type of interaction, e.g., searching for an artist, applying a filter, clicking on a bounding box or an image.

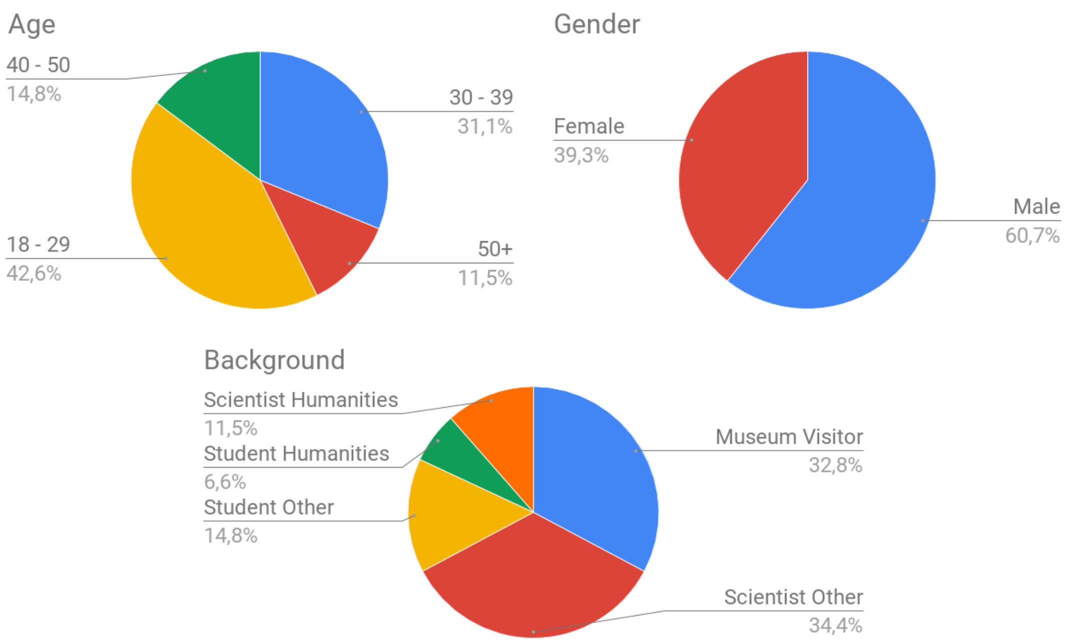

Fig. 2. The age, gender and background distribution of the participants.

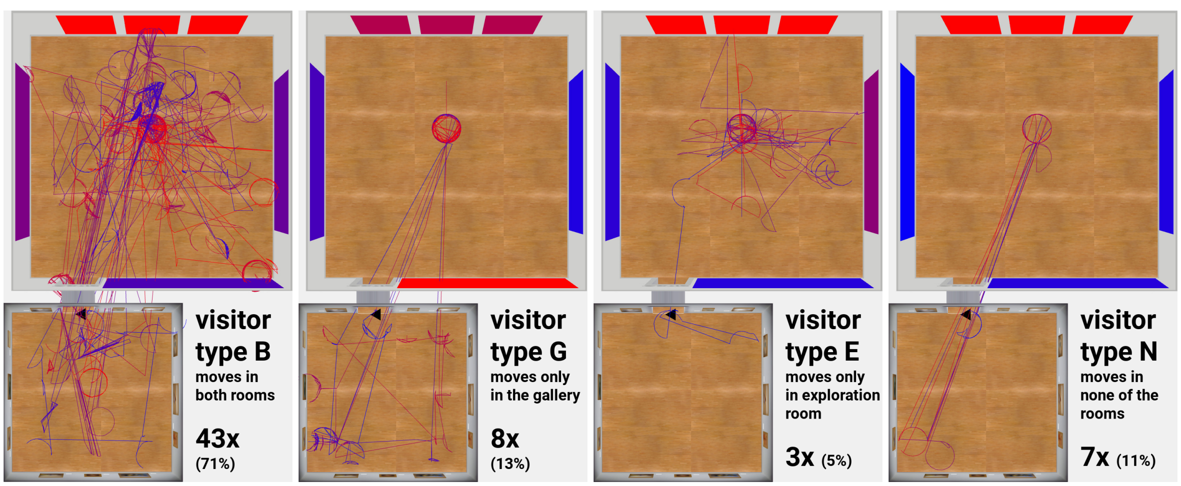

Fig. 3. Different movement patterns of museum visitors. The entry point is depicted with a black triangle. The color gradient of the trajectory flows from blue to red over time. The amount of time a visitor focused on particular walls in the exploration room is also shown on a color gradient from blue (less time) to red (more time).

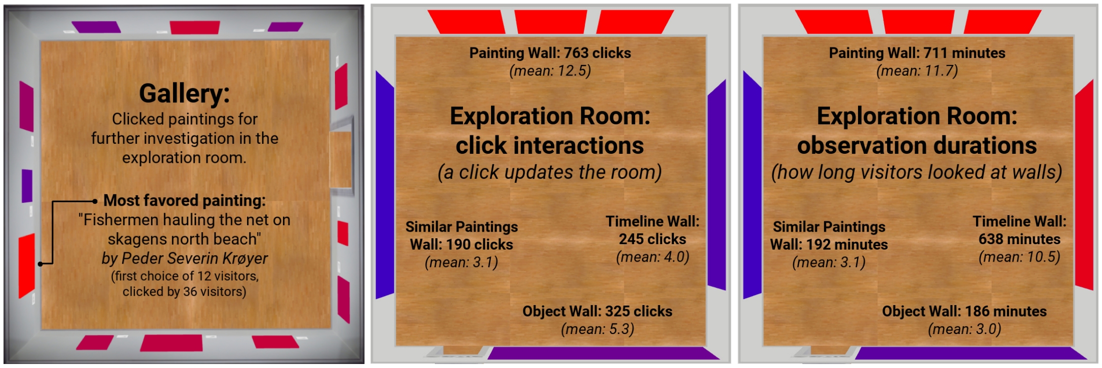

Fig. 4. Engagement, color-coded on a linear scale from blue (less engagement) to red (more engagement), with paintings and walls in the virtual museum and the mean per participant.

5.1 Movement & Focus

The gathered logging information of a participant gave us valuable insight on how our virtual museum as a whole and the individual interfaces in particular are used.

Movement Patterns. In order to get an overview of how museum visitors move in the virtual space, we computed a trajectory for each visitor in the form of a line string that depicts graphically the movements during the entire stay in the museum. We color-coded individual lines according to the time of the movement using a color gradient from blue (entering the museum) to red (leaving the museum). We observed four different movement patterns, which are illustrated in Figure 3. Circular artifacts occurring in all patterns denote rotations of a museum visitor. The straight lines through the walls are results of a direct transportation from the gallery to the exploration room.

- **Visitor type B:** 43 visitors used the entire space in **B**oth rooms to move, turn around and explore the different visual interfaces. They spent an average of 39 minutes in the museum, which is equal to the mean duration of all visitors. A majority of 71% of visitors showing this movement pattern, in other words, they made use of the degrees of freedom we provided them with, confirming our decision to offer a 3D space for both rooms.
- **Visitor type G:** Eight visitors moved almost entirely in the **G**allery, which shows typical movements one would expect from a real museum visit. Inside the exploration room, the visitors typically lingered at the initial position and only rotated to observe all walls from a distant perspective. The color gradient of the movement trajectory (rather blue in the gallery and red in the exploration room) and the relatively high mean of 47 minutes museum visit however indicate that these visitors required some time to adapt to the virtual space.
- **Visitor type E:** Three visitors disregarded the gallery and went directly into the **E**xploration room. They spent less than 30 minutes each (mean: 18 minutes), and in their qualitative reflections focused more on the technological concept of the museum.
- **Visitor type N:** Seven visitors made limited use of the 3D environment, indicating unfamiliarity with 3D technology. They only rotated to observe and interact with the walls resulting in **N**o movement. The mean time spent in the museum was 25 minutes. However, all visitors of that type confirmed their interest to return to the virtual museum.

It has to be said that the movement patterns do not show any significant influence on the ratings of the museum visit, the ratings of information content or intuitiveness of the visualizations provided. Thus, the patterns tend to reflect the heterogeneous visitor group and different behaviors in virtual spaces, which has been observed in previous studies [64].

Engagement with Visual Interfaces. We recorded the number of visitor interactions with each wall and the time they spent observing each wall, in other words, their overall engagement with the visual interfaces offered. Therefore, we color-coded the walls from blue (less engagement) to red (more engagement) on a linear scale; the results are depicted in Figure 4. According to click interactions, most activity has been registered for the Painting Wall (763 clicks), followed by the Objects Wall (325), the Timeline Wall (245), and the Similar Painting Wall (190). Each of these interactions changed the context of the Exploration Room, by either focusing on a new painting by the same artist on the Painting Wall, affecting also the Similar Paintings Wall, or by selecting a painting of a new artist, thereby updating the whole Exploration Room. The visitors looked at the Painting Wall for 711 minutes in total, followed by the Timeline Wall for 638 minutes, the Similar Paintings Wall for 192 minutes, and 135 minutes for the Objects Wall. The Painting Wall reached the highest values for engagement, which is not surprising as visitors initially look towards this wall when entering the room or clicking a painting elsewhere. Surprisingly, the Objects Wall registered more click interactions than the other two walls, which can be seen as an indicator for the visitors' interest in focusing on the individual elements of paintings. Furthermore, the Timeline Wall was looked at for the second longest time, explainable by its capacity to tell stories the oeuvre of artists. Finally, we compared the time visitors spent in each of the two rooms. The Gallery was visited for 682 minutes (mean: 11.2), whereas the Exploration Room was visited for 1,727 minutes (mean: 28.3). Considering the circumstance that our invitation neutrally asked participants to "visit a virtual museum", and that the visitors started in the Gallery, those numbers underpin the value of our solution to be an important complement to real museums.

5.2 The Value of our Virtual Museum

We asked the museum visitors to evaluate the utility and importance of our virtual museum. The results are depicted in Figure 5. Firstly, we asked them to rate their visit on a 5-point Likert scale from *boring* (1) to *exciting* (5). While only four visitors rated their visit as rather boring (7%), a majority of 39 visitors found it to be a rather exciting experience (64%). The remaining 18 visitors did not express a tendency. With a median of 4 (rather exciting) and a mean of 3.72, most visitors enjoyed the stay in our virtual museum. Secondly, we asked if our virtual museum is seen as a valuable complement to real museums. 46 visitors (75%) supported this capacity of our solution, reaching a median of 4 (rather agree) and a mean of 4.05. According to the qualitative feedback, this relates first and foremost to the exploratory functionality serving the visitors with new means to engage with art. Thirdly, accounting for the closure of real museums in times of the pandemic, we wanted to know the visitors' opinions on whether our virtual museum would be an important replacement. While 9 visitors rather disagreed, the majority of 39 visitors (64%) expressed their gratitude for having a museum-like space in which they were able to appreciate and discover art. A representative comment of an excited visitor aligns with the aims of our solution: *"The advantage of digital interaction possibilities in virtual museums compared to real classical exhibitions is immense from my perspective and increases the attraction for me to participate in exhibitions."* To get a more detailed overview of positive and negative aspects of our virtual museum prototype, we asked the visitors if they would return. Both positive and negative feedback is discussed below.

Positive Feedback. A majority of participants (47 of 62, 76%) declared interest in visiting our virtual museum again, pointing out the benefits of the exploratory environment. A representative comment was *"It is exciting that you can 'find your way'deeper and deeper into different artists and styles."* One participant valued this richness of accessible information and remarked that it is *"difficult in a real museum"* to link different artists and their works. Another participant pointed out more clearly that the means to explore art are limited in real museums (*"From my perspective, this makes it more appealing than many of my previous classic visits to museums."*), confirming the

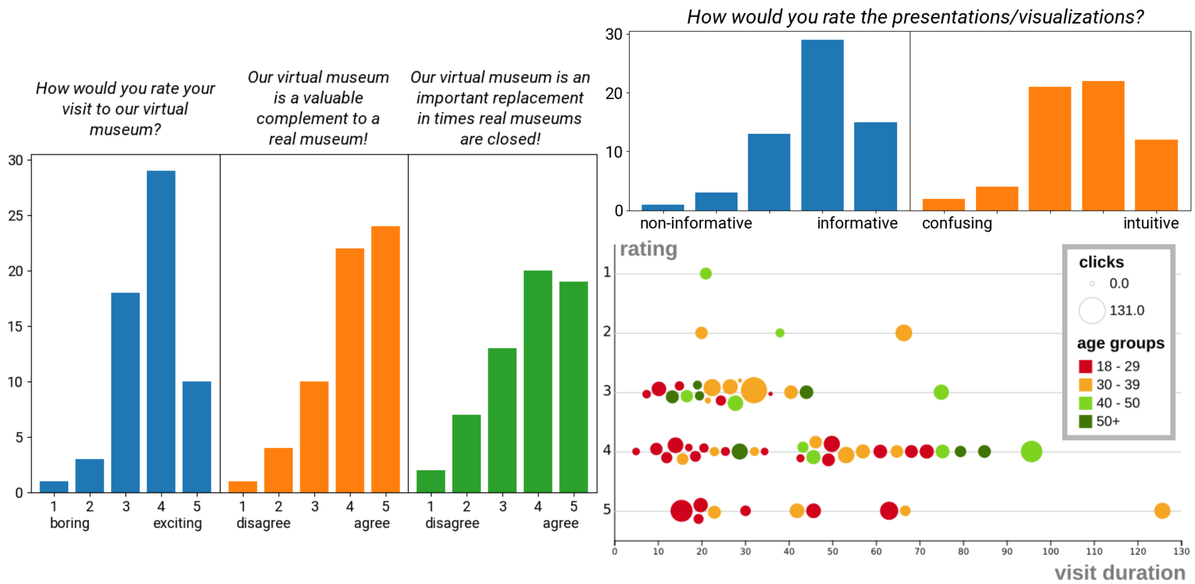

Fig. 5. Participants' ratings on the value of the virtual museum and on information content and intuitiveness of the visualizations, also excitement relates to age group and engagement.

capacity of the virtual museum as a valuable complement. Furthermore, the vastness of the data set was seen as an attractive reason for visitors to come back: *"There was far too little time to even come close to looking at the whole treasure trove of pictures."*

Negative Feedback. 15 out of 62 participants (24 %) indicated that they would not re-visit the virtual museum. The reasons given relate to aspects that real museums provide and those which virtual ones do not. One comment targets two limitations of our current solution: *"Compared to a real museum I miss the atmosphere and the contemplative - and on the other hand in-depth information."* First, our virtual museum does not include a curated collection. It only offers one room mimicking a real museum, in which we arranged paintings that generate diverse entry points to the analysis room. On the other hand, the vast image collection suffers from incomplete metadata, a remainder of duplicates and an error-prone object detector. Second, it does not include social aspects that emulate the atmosphere of real museums at all. Lastly, some participants expressed their general disinterest in art (*"I'm not really into paintings."*)

5.3 Acceptance of our Virtual Museum Concept

We analyzed the information provided by the participants to discover if any characteristics correlate with the acceptance of a virtual museum solution. Next to the general assessment of the virtual museum's value, we asked the visitors to rate on a 5-point Likert scale if the provided visualizations were informative and intuitive (see Figure 5). Both aspects were evaluated positively, reaching mean values of 3.89 for information content and 3.62 for intuitiveness.

Concept Addresses Needs in Humanities Research. Museum visitors with a background in the humanities (seven scientists and four students) rated the experience more favorably compared to other groups of the study. With a mean value of 4.45 (eight times rating 5), the assessment of the museum as a valuable complement to real museums is significantly higher compared to the entire group. This group also gave the highest scores for information content (mean: 4.18) and intuitiveness (mean: 4.45) of the presented visualizations. This might relate to the concept of distant viewing [2] that transfers the idea of quantitative text analysis (distant reading) to image collections. By employing temporal and similarity analysis, and by extracting objects from paintings on a large scale and making the results explorable, our virtual museum provides novel "distant viewing" avenues for humanities research.

Acceptance Correlates with Age and Engagement. Figure 5 provides a multifaceted picture on factors influencing the given museum visit ratings from *boring* (1) to *exciting* (5). It is evident that all visitors who gave the highest rating were younger than 40. None of the visitors of the group with the youngest participants (aged 18 to 29 years) gave a rating worse than 3. This age group also reaches the highest mean values for rating the visit (4.03), the virtual museum being a complement (4.38) or a replacement during times real museums are closed (4.11). These results may relate to the digital native generation feeling more at home in a virtual environment, having more interaction skills and being more technologically adept [41]. However, the chart also shows that the longer the duration time of older visitors, the better their rating was. The circle size reflects the number of clicks, in other words, the number of images selected in the Exploration Room. As there is no clear tendency visible, one can conclude that visitors may be of different types: the ones who aim to gather more information and those who spend more time observing the paintings. However, what the distribution clearly indicates is that ratings correlate with actual visit durations. While the actual mean visit duration of visitors who rated the virtual museum as (rather) exciting (with 4 or 5) is 43.5 minutes, the other visitors spent only 28.5 minutes on average.

Perceived vs. Actual Duration of Museum Visit. Due to consistent logging of all visitor activities inside the virtual museum, we were able to determine the actual duration of museum visits for all participants. While the mean actual duration of the visits was 39 minutes with an average entry delay of 25 minutes, 13 visitors spent more than one hour in the museum. As well as logging the time, we also asked the visitors in the final questionnaire how much time they spent in the museum. 13 visitors underestimated while only four visitors overestimated their duration of stay. Visitors tended to underestimate the duration of their stay, especially when they spent more time in the museum. According to studies in psychology [24, 25], active participation and higher levels of motivation lead to perceiving time to be shorter than it appears to last. Thus, we consider this as an indicator for casual entertainment.

Does it Have to Be 3D?. Two visitors remarked that the 3D environment in its current form is superfluous (*"The 3D-stuff has no value for me."*), or suggested to take advantage of immersive technologies (*"The use in VR glasses would make the 3D element more important."*). In our design phase, we evaluated opportunities and drawbacks of diverse possible implementations of our prototype visualizations. We decided on a 3D representation of the gallery to emulate the real space as best as possible. On the one hand, this consideration serves the desire of partaking visitors of our evaluation to reproduce the atmosphere of real museums. On the other hand, it is in line with existing virtual museum implementations. As those are often only available as desktop applications to reach a large audience, we did not target a virtual reality solution for our prototype implementation. Nevertheless, we added an option to deactivate the 3D control for visitors who prefer exploring in a 2D environment.

6 LIMITATIONS & FUTURE WORK

The evaluation results underpin the utility of our virtual museum as a complement to real museums offering new avenues for the general public to engage with art collections. We registered some limitations of our current solution, some of which we were aware of prior to the evaluation. However, our main focus was connecting quantitative analysis of paintings to virtual museums and evaluating how visitors adopt and value such *distant* *viewing* concepts. Below, we discuss how our concept might be improved to make virtual museums even more engaging for the general public.

A Curated Virtual Museum. Some of the evaluation participants remarked that they were missing a curated exhibition where a story is told based on a selection of objects and artworks to target a particular set of goals For example, a goal could be to reveal new information about an artist's life and work. In this context, the exploration room's search menu might feature custom filters and preset search parameters written by the curator resulting in visualizations that support a particular curatorial goal. As the story of an artist is inherent in the Timeline Wall this could explain why it was observed longer than the other visualizations. Some visitors also addressed the need for "a tour guide's insight on the artwork". These results are not unsurprising, as public audiences prefer curated collections that have a story to tell [22, 76], while the inclusion of information retrieval principles e.g. search options to fulfill different information needs can lead to designing for an expert audience. Nevertheless, our system allows serendipitous discoveries without expert knowledge. The general intention of our concept is to connect the Exploration Room to real museum exhibitions that are already curated. In addition, we could also improve the storytelling in the Exploration Room by improving the data quality and by including external sources like Wikipedia. It should be noted that the digital versions of the paintings provide the correct aspect ratio of the real painting, the resolution is lower and the real-world size is missing for many of the paintings from the data set.

The Social Virtual Museum. Some participants pointed out that much of the atmosphere of a real museum (the architecture of the building, background noises or even the smell) is missing, leading to discomfort being "alone in the virtual space". Currently, our solution does not support shared visits, in other words, it does not generate or strengthen social and emotional connections between visitors. Partaking visitors expressed their wish to interact with other people, "to discuss the art and share impressions". This lack of social aspects is common for most of current online and virtual cultural experiences [79]. A solution would be to allow multiple people to visit the virtual museum together for example by introducing avatars to the virtual space and therefore allowing synchronous interactions. There is also potential for increasing degrees of "gamification" as well as the integration of social media interactions. To further include more of the actual atmosphere, background noises from real museums including sounds of footsteps could be included.

Personalization of Virtual Museum Visit. Currently, our interface offers a search option for artists the visitors are interested in, then filter results by painting styles. Although we developed a series of other means to explore the art collection (e.g., searching for particular objects and compare those on the timeline), we did not include those in the evaluation prototype in order to keep the interface as intuitive as possible. The problem of additional representations and metadata was also reported by Ma et al. [51] "when designing visualizations for museums, additional representations should be carefully considered and secondary data may need to be left out". However, a few visitors wished for more personalized search functionality. A further possibility to motivate visitors to interact with the provided interfaces would be allowing them to feed in their own data, for example, by uploading images that they are interested to compare with our art collection [54].

Moving Beyond Paintings. Our museum only focuses on paintings. Visitors expressed the desire to extend our collection with 3D artifacts such as statues, musical instruments, or tools. However, reconstructions of 3D objects are more expensive, more complex and more error-prone [73]. In addition, our methods are tailored for processing 2D image data, which we would have to adapt to 3D sources.

Cross Depiction Problem & Incompleteness. One limitation of our virtual museum relates to the machine learning approach and the data that was used to train the methods. The neural networks that we applied were trained on ImageNet and OpenImages. Both data sets only contain real photographs instead of fine-art paintings. This leads to the cross-depiction problem [33]. For example, neither an abstract cat nor a cubist cat resembles a real cat. Because of this, and a general error-proneness of automatic approaches, the bounding boxes of the objects in the images can be wrongly labeled, and not all objects have been detected. Also, hierarchies like OpenImages suffer from incompleteness, because they do not contain all kinds of objects. To improve automatic object detection, context-aware approaches [26] combining image features and meta-data can increase the accuracy. Furthermore, a crowdsourcing approach that combines manual annotation of museum visitors with machine learning approaches like few-shot or zero-shot learning can extend hierarchies like OpenImages and can correct wrong objects. Based on the feedback of the evaluation, we started to add means to annotate images with bounding boxes and so to create new labels. Incompleteness is not only given in the class hierarchies but also by the data. Although we included around 200,000 paintings, there are still many missing paintings for different artists and some artists are not yet part of the data set.

Profiling of Artists. Another future direction is to put more emphasis on artists instead of paintings. One visitor suggested to include more "information about the artist and how they relate to artists from a similar group/time". For this purpose, we could include biographical information from external sources and apply profiling techniques, proven to deliver valuable results for prosopographical data sets [40], aiming to discover artists similar to an artist of interest. The similarity can be computed based on metadata like activity period, style, genre, or social relations to other artists. This can be even expanded by including similarity metrics based on depicted objects, themes, used colors and image embedding similarity.

7 CONCLUSION

Our work contributes a virtual museum model that connects a virtual version of a Gallery with an Exploration Room that contextualizes artworks with a large image archive based on WikiArt. For this purpose, we take advantage of machine learning techniques to extract object information from artworks and determine similarities among them. The results are depicted as interactive visualizations that provide a novel virtual museum experience to visitors, allowing them access to paintings beyond those that are exhibited in a real museum. From our evaluation, we can conclude that a virtual museum is not a replacement for a real museum, but during times when real museums can not be visited, they are a viable alternative. Our solution was further regarded as a valuable complement to a real museum, exemplified by how the interactive visualizations, composed for objects of interest, can intrigue new visitors, give them new thoughts and address the information needs of general visitors and humanities scholars. Lastly, it is important to note that designing a virtual museum is not about replicating a real museum in a virtual space, but more about extending the notion of a museum taking advantage of digital methods.

REFERENCES

[1] James Andrews and Werner Schweibenz. 1998. New Media for Old Masters: The Kress Study Collection Virtual Museum Project. *Art Documentation: Journal of the Art Libraries Society of North America* 17, 1 (1998), 19–27.
[2] Taylor Arnold and Lauren Tilton. 2019. Distant viewing: analyzing large visual corpora. *Digital Scholarship in the Humanities* 34, Supplement_1 (2019), i3–i16.
[3] Benjamin B Bederson. 2001. PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. In *Proceedings of the 14th annual ACM symposium on User interface software and technology.* 71–80.
[4] Mafkereseb Kassahun Bekele, Roberto Pierdicca, Emanuele Frontoni, Eva Savina Malinverni, and James Gain. 2018. A survey of augmented, virtual, and mixed reality for cultural heritage. *Journal on Computing and Cultural Heritage (JOCCH)* 11, 2 (2018), 1–36.
[5] Bernadette Biedermann. 2021. Virtual museums as an extended museum experience: Challenges and impacts for museology, digital humanities, museums and visitors–in times of (Coronavirus) crisis. *DHQ: Digital Humanities Quarterly* 15, 3 (2021).
[6] Mark-Jan Bludau, M Dörk, and F Heidmann. 2019. Relational perspectives as situated visualizations of art collections. In *Proceedings of the ADHO Conference on Digital Humanities (DH), Utrecht, The Netherlands.* 9–12.

[7] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D^3 data-driven documents. *IEEE transactions on visualization and computer graphics* 17, 12 (2011), 2301–2309.
[8] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. 2000. LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*. 93–104.
[9] Paolo Brivio, Marco Tarini, and Paolo Cignoni. 2010. Browsing large image datasets through voronoi diagrams. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 1261–1270.
[10] Daniel Alejandro Loaiza Carvajal, María Mercedes Morita, and Gabriel Mario Bilmes. 2020. Virtual museums. Captured reality and 3D modeling. *Journal of Cultural Heritage* 45 (2020), 234–239.
[11] Richard J Cole, Frithjof Dau, Jon Ducrou, Peter W Eklund, and Tim Wray. 2019. Navigating Context, Pathways and Relationships in Museum Collections using Formal Concept Analysis. *International Journal for Digital Art History* 4 (2019), 5–13.
[12] Lhaylla Crissaff, Louisa Wood Ruby, Samantha Deutch, R Luke DuBois, Jean-Daniel Fekete, Juliana Freire, and Claudio Silva. 2017. ARIES: enabling visual exploration and organization of art image collections. *IEEE computer graphics and applications* 38, 1 (2017), 91–108.
[13] Damon Crockett. 2016. Direct visualization techniques for the analysis of image data: the slice histogram and the growing entourage plot. *International Journal for Digital Art History* 2 (2016).
[14] Elliot J Crowley, Omkar M Parkhi, and Andrew Zisserman. 2015. Face painting: querying art with photos. In *Proceedings of the British Machine Vision Conference (2015)*. BMVA Press.
[15] Elliot J Crowley and Andrew Zisserman. 2014. In search of art. In *European Conference on Computer Vision*. Springer, 54–70.
[16] Elliot J Crowley and Andrew Zisserman. 2014. The state of the art: Object retrieval in paintings using discriminative regions. In *Proceedings of the British Machine Vision Conference (2014)*. BMVA Press.
[17] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
[18] Jos Dirksen. 2013. *Learning Three. js: the JavaScript 3D library for WebGL*. Packt Publishing Ltd.
[19] Marian Dörk, Christopher Pietsch, and Gabriel Credico. 2017. One view is not enough: High-level visualizations of a large cultural collection. *Information Design Journal* 23, 1 (2017), 39–47.
[20] Johanna Drucker. 2013. Is there a "digital" art history? *Visual Resources* 29, 1-2 (2013), 5–13.
[21] Bruno Dumas, Bram Moerman, Sandra Trullemans, and Beat Signer. 2014. ArtVis: combining advanced visualisation and tangible interaction for the exploration, analysis and browsing of digital artwork collections. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*. 65–72.
[22] Timothy AM Ewin and Joanne V Ewin. 2016. In defence of the curator: maximising museum impact. *Museum Management and Curatorship* 31, 4 (2016), 322–330.
[23] Marco Fiorucci, Marina Khoroshiltseva, Massimiliano Pontil, Arianna Traviglia, Alessio Del Bue, and Stuart James. 2020. Machine Learning for Cultural Heritage: A Survey. *Pattern Recognition Letters* 133 (2020), 102 – 108. https://doi.org/10.1016/j.patrec.2020.02.017
[24] Paul Fraisse. 1963. *The psychology of time.* Harper & Row.
[25] Paul Fraisse. 1984. Perception and estimation of time. *Annual review of psychology* 35, 1 (1984), 1–37.
[26] Noa Garcia, Benjamin Renoust, and Yuta Nakashima. 2019. Context-aware embeddings for automatic art analysis. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval*. 25–33.
[27] Noa Garcia and George Vogiatzis. 2018. How to read paintings: semantic art understanding with multi-modal retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. 0–0.
[28] Ivan Giangreco, Loris Sauter, Mahnaz Amiri Parian, Ralph Gasser, Silvan Heller, Luca Rossetto, and Heiko Schuldt. 2019. VIRTUE: a virtual reality museum Experience. In *Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion*. 119–120.
[29] Katrin Glinka, Christopher Pietsch, and Marian Doerk. 2017. Past Visions and Reconciling Views: Visualizing Time, Texture and Themes in Cultural Collections. *DHQ: Digital Humanities Quarterly* 11, 2 (2017).
[30] Doron Goldfarb, Max Arends, Josef Froschauer, and Dieter Merkl. 2011. Revisiting 3D information landscapes for the display of art historical web content. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology*. 1–8.
[31] Nicolas Gonthier, Yann Gousseau, Said Ladjal, and Olivier Bonfait. 2018. Weakly supervised object detection in artworks. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 0–0.
[32] Flavio Gortana, Franziska von Tenspolde, Daniela Guhlmann, and Marian Dörk. 2018. Off the Grid: Visualizing a Numismatic Collection as Dynamic Piles and Streams. *Open Library of Humanities* 4, 2 (2018).
[33] Peter Hall, Hongping Cai, Qi Wu, and Tadeo Corradi. 2015. Cross-depiction problem: Recognition and synthesis of photographs and artwork. *Computational Visual Media* 1, 2 (2015), 91–103.
[34] Masaki Hayashi, Steven Bachelder, and Masayuki Nakajima. 2016. Automatic generation of personal virtual museum. In *2016 International Conference on Cyberworlds (CW)*. IEEE, 219–222.
[35] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 203–212.

[36] Nadav Hochman and Lev Manovich. 2013. Zooming into an Instagram City: Reading the local through social media. *First Monday* 18, 7 (2013).

[37] Stefka Hristova. 2016. Images as data: cultural analytics and Aby Warburg's Mnemosyne. *International Journal for Digital Art History* 2 (2016).

[38] Chun-Rong Huang, Chu-Song Chen, and Pau-Choo Chung. 2005. Tangible photorealistic virtual museum. *IEEE computer graphics and applications* 25, 1 (2005), 15–17.

[39] Stefan Jänicke. 2018. Timages: Enhancing Time Graphs with Iconographic Information. In *LEVIA'18: Leipzig Symposium on Visualization in Applications*.

[40] Stefan Jänicke, Josef Focht, and Gerik Scheuermann. 2015. Interactive visual profiling of musicians. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 200–209.

[41] Ali Jarrahi, Mohammad Hossein und Eshraghi. 2019. Digital Natives vs Digital Immigrants. *Journal of Enterprise Information Management* 32, 6 (2019).

[42] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data* 7, 3 (2021).

[43] Pauline Junginger, Dennis Ostendorf, Barbara Avila Vissirini, Anastasia Voloshina, Timo Hausmann, Sarah Kreiseler, and Marian Dörk. 2020. The Close-up Cloud. *International Journal for Digital Art History* 5 (2020), 6–2.

[44] Katerina Kabassi. 2017. Evaluating websites of museums: State of the art. *Journal of Cultural Heritage* 24 (2017), 184–196.

[45] Chairi Kiourt, Anestis Koutsoudis, and George Pavlidis. 2016. DynaMus: A fully dynamic 3D virtual museum framework. *Journal of Cultural Heritage* 22 (2016), 984–991.

[46] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Tom Duerig, et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision* 128, 7 (2020), 1956––1981.

[47] Sabine Lang and Björn Ommer. 2018. Attesting similarity: Supporting the organization and study of art image collections with computer vision. *Digital Scholarship in the Humanities* 33, 4 (2018), 845–856.

[48] Fotis Liarokapis, Stella Sylaiou, Anirban Basu, Nicholaos Mourkoussis, Martin White, and Paul F Lister. 2004. An Interactive Visualisation Interface for Virtual Museums.. In *VAST*. 47–56.

[49] Céline Loscos, Franco Tecchia, Antonio Frisoli, Marcello Carrozzino, Hila Ritter Widenfeld, David Swapp, and Massimo Bergamasco. 2004. The Museum of Pure Form: touching real statues in an immersive virtual museum.. In *VAST*. 271–279.

[50] Jean-Luc Lugrin, Florian Kern, Ruben Schmidt, Constantin Kleinbeck, Daniel Roth, Christian Daxer, Tobias Feigl, Christopher Mutschler, and Marc Erich Latoschik. 2018. A location-based VR museum. In *2018 10th International Conference on Virtual Worlds and Games for Serious Applications (VS-Games)*. IEEE, 1–2.

[51] Joyce Ma, Kwan-Liu Ma, and Jennifer Frazier. 2019. Decoding a Complex Visualization in a Science Museum–An Empirical Study. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 472–481.

[52] Eetu Mäkelä, Eero Hyvönen, and Tuukka Ruotsalo. 2012. How to deal with massively heterogeneous cultural heritage data–lessons learned in CultureSampo. *Semantic Web* 3, 1 (2012), 85–109.

[53] André Malraux. 1954. *Le musée imaginaire de la sculpture mondiale*. Vol. 3. Gallimard.

[54] Hui Mao, Ming Cheung, and James She. 2017. Deepart: Learning joint representations of visual arts. In *Proceedings of the 25th ACM international conference on Multimedia*. 1183–1191.

[55] Leland McInnes, John Healy, and Steve Astels. 2017. hdbscan: Hierarchical density based clustering. *Journal of Open Source Software* 2, 11 (2017), 205.

[56] L. McInnes, J. Healy, and J. Melville. 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *ArXiv e-prints* (Feb. 2018). arXiv:stat.ML/1802.03426

[57] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. UMAP: Uniform Manifold Approximation and Projection. *The Journal of Open Source Software* 3, 29 (2018), 861.

[58] Francisco J Melero, Jorge Revelles, and Maria Luisa Bellido. 2018. Atalaya3D: Making Universities' Cultural Heritage Accessible Through 3D Technologies.. In *GCH*. 31–35.

[59] Tamara Munzner. 2014. *Visualization analysis and design*. CRC press.

[60] Matterport Inc. Musée National d'Histoire et d'Art. 2017. Musée National d'Histoire et d'Art - Matterport Model. https://my.matterport.com/show/?m=hPsqTjD3PNQ. (Retrieved 2021-03-31).

[61] Hermann Pflüger and Thomas Ertl. 2016. Sifting through visual arts collections. *Computers & Graphics* 57 (2016), 127–138.

[62] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.

[63] James Rice. 1988. Serendipity and holism: The beauty of OPACs. *Library Journal* 113, 3 (1988), 138–41.

[64] Corina Sas, Gregory O'Hare, and Ronan Reilly. 2004. A performance analysis of movement patterns. In *International Conference on Computational Science*. Springer, 954–961.

[65] Joseph Schlecht, Bernd Carqué, and Björn Ommer. 2011. Detecting gestures in medieval images. In *2011 18th IEEE International Conference on Image Processing*. IEEE, 1285–1288.
[66] Werner Schweibenz. 2004. Virtual museums. *The Development of Virtual Museums, ICOM News Magazine* 3 (2004), 3.
[67] Xi Shen, Alexei A Efros, and Mathieu Aubry. 2019. Discovering visual patterns in art collections with spatially-consistent feature learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9278–9287.
[68] Google LLC Statens Museum for Kunst. 2012. SMK-Statens Museum for Kunst in Copenhagen, Google Arts and Culture tours. https://www.google.com/maps/@55.6887473,12.5786094,2a,75y,91.15h,81.73t/data=!3m6!1e1!3m4!1sWSukjJxQiuc0fX6hOdAzsg!2e0!7i13312!8i6656. (Retrieved 2021-03-31).
[69] Gjorgji Strezoski, Lucas Fijen, Jonathan Mitnik, Dániel László, Pieter de Marez Oyens, Yoni Schirris, and Marcel Worring. 2020. TindART: A Personal Visual Arts Recommender. In *Proceedings of the 28th ACM International Conference on Multimedia*. 4524–4526.
[70] Gjorgji Strezoski, Inske Groenen, Jurriaan Besenbruch, and Marcel Worring. 2018. Artsight: an artistic data exploration engine. In *Proceedings of the 26th ACM international conference on Multimedia*. 1240–1241.
[71] Gjorgji Strezoski, Arumoy Shome, Riccardo Bianchi, Shruti Rao, and Marcel Worring. 2019. Ace: Art, color and emotion. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1053–1055.
[72] Gjorgji Strezoski and Marcel Worring. 2018. Omniart: a large-scale artistic benchmark. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 14, 4 (2018), 1–21.
[73] Sylaiou Styliani, Liarokapis Fotis, Kotsakis Kostas, and Patias Petros. 2009. Virtual museums, a survey and some issues for consideration. *Journal of cultural Heritage* 10, 4 (2009), 520–528.
[74] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*.
[75] Alice Thudt, Uta Hinrichs, and Sheelagh Carpendale. 2012. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1461–1470.
[76] Martin Tröndle, Steven Greenwood, Konrad Bitterli, and Karen van den Berg. 2014. The effects of curatorial arrangements. *Museum Management and Curatorship* 29, 2 (2014), 140–173.
[77] Michael Twidale. 1993. Redressing the balance: the advantages of informal evaluation techniques for Intelligent Learning Environments. *Journal of Artificial Intelligence in Education* 4 (1993), 155–155.
[78] Ignacio Úbeda, Jose M Saavedra, Stéphane Nicolas, Caroline Petitjean, and Laurent Heutte. 2020. Improving pattern spotting in historical documents using feature pyramid networks. *Pattern Recognition Letters* 131 (2020), 398–404.
[79] Maria Vayanou, Akrivi Katifori, Angeliki Chrysanthi, and Angeliki Antoniou. 2020. Cultural Heritage and Social Experiences in the Times of COVID 19.. In *AVI²CH@ AVI*.
[80] Krzysztof Walczak, Wojciech Cellary, and Martin White. 2006. Virtual museum exbibitions. *Computer* 39, 3 (2006), 93–95.
[81] Chaoli Wang, John P Reese, Huan Zhang, Jun Tao, Yi Gu, Jun Ma, and Robert J Nemiroff. 2015. Similarity-based visualization of large image collections. *Information Visualization* 14, 3 (2015), 183–203.
[82] Martin Warnke and Lisa Dieckmann. 2016. Prometheus meets Meta-Image: implementations of Aby Warburg's methodical approach in the digital era. *Visual Studies* 31, 2 (2016), 109–120.
[83] Nicholas Westlake, Hongping Cai, and Peter Hall. 2016. Detecting people in artwork with cnns. In *European Conference on Computer Vision*. Springer, 825–841.
[84] Melvin Wevers and Thomas Smits. 2018. Seeing History: The Visual Side of the Digital Turn. (2018).
[85] Melvin Wevers and Thomas Smits. 2020. The visual digital turn: Using neural networks to study historical images. *Digital Scholarship in the Humanities* 35, 1 (2020), 194–207.
[86] Mitchell Whitelaw et al. 2015. Generous interfaces for digital cultural collections. *DHQ: Digital Humanities Quarterly* 9, 1 (2015).
[87] Florian Windhager, Paolo Federico, Günther Schreder, Katrin Glinka, Marian Dörk, Silvia Miksch, and Eva Mayr. 2018. Visualization of cultural heritage collection data: State of the art and future challenges. *IEEE transactions on visualization and computer graphics* 25, 6 (2018), 2311–2330.
[88] Andy T Woods, Carlos Velasco, Carmel A Levitan, Xiaoang Wan, and Charles Spence. 2015. Conducting perception research over the internet: a tutorial review. *PeerJ* 3 (2015), e1058.
[89] So Yamaoka, Lev Manovich, Jeremy Douglass, and Falko Kuester. 2011. Cultural analytics in large-scale visualization environments. *Computer* 44, 12 (2011), 39–48.
[90] Pradeep Yarlagadda, Antonio Monroy, Bernd Carque, and Björn Ommer. 2010. Recognition and analysis of objects in medieval images. In *Asian Conference on Computer Vision*. Springer, 296–305.

A GALLERY PAINTINGS

The gallery includes the following 12 paintings placed at four walls depicted in Figure 6:

- Peder Severin Kroyer "Fishermen hauling the net on skagen s north beach" (1883)
- Ogata Gekkoprint "From series women s customs and manners" (1895)
- Camille Pissarrothe "Harvest of hay in eragny" (1887)
- Albert Blochthe "Garden of asses II" (1939)
- Anna Ancher "Harvesters" (1905)
- Hans Andersen Brendekilde "Worn out" (1889)
- Alphonse Mucha "Austria" (1899)
- Vilhelm Hammershoi "Interior from strandgade with sunlight on the floor" (1901)
- Andre Derain "The basin of london" (1906)
- Georges Seurat "Harbour at port en bessin at high tide" (1888)
- Homer Watson "The flood gate" (1901)
- Henri Fantin Latour "Peaches" (1903)

Selecting a painting in the Gallery takes the visitor to the Exploration Room. Due to space restrictions, we could not include an exhaustive example in the paper. Figure 7 shows the composition of visualizations in the Exploration Room when clicking Henri Fantin Latour's painting "Peaches"; the large amount of fruits and flowers is notable.

B CLUSTERING & OUTLIER

We also computed a 2D embedding with UMAP [56, 57] for each image based on the image embeddings of the Inception Resnet V2 [74]. The UMAP embeddings are used to compute clusters for each object class with HDBSCAN [55] and outliers with Local Outlier Factor [8]. To allow for uncertainty in the cluster assignment, we apply a soft clustering strategy. For each image, a probability distribution for the cluster assignment is computed, which allows us to assign an image to multiple clusters and therefore to assign outliers that are not part of a specific cluster to multiple clusters. To prevent large and heterogeneous clusters, we are applying a leaf selection method for the cluster selection, more likely resulting in smaller and homogeneous clusters. In order to visualize the clusters later we compute the centroid for each cluster in the 2D space. Additional to the dimension reduction and clustering based on the different classes we compute the same for artists and styles, as they are the most frequent metadata.

To give a better overview of a specific object class, artist or style, we presented each cluster by the image closest to the centroid of the cluster. To prevent overlaps, we apply a collision detection that applies a small offset to the x and y-coordinate. When clicking on a cluster the images of the clusters move to the outer part of the space, while the cluster results appear in the center. If the cluster is too large to fit all results on the canvas an additional soft-clustering is applied on the cluster, presenting only the centroids of each new cluster. Images can be part of multiple clusters, by mouse-over the other clusters are highlighted. Additionally a tool-tip shows the class, style and artist distribution of the cluster. The timeline visualization can be linked to the cluster map to also allow the user to investigate the time distribution of a cluster of interest.

We also computed outliers in the collection through Local Outlier Factor [8] for the styles, artists and objects depicted in the images. Similar to the nearest neighbors of an image of interest, they can be displayed in a Picture Cloud with the distance as a scaling factor to see the most unusual paintings of an artist, a style or a specific object class. An example of the class Elephant can be seen in Figure 8.

We excluded visual depictions resulting from the above described technique due to their complex composition and the limited intuitiveness of results. Dimensional projections of a large number of high-dimensional vectors results for images in a lot of clutter, which can be seen in Figure 9. The Figure shows an example of the Ukiyo-e

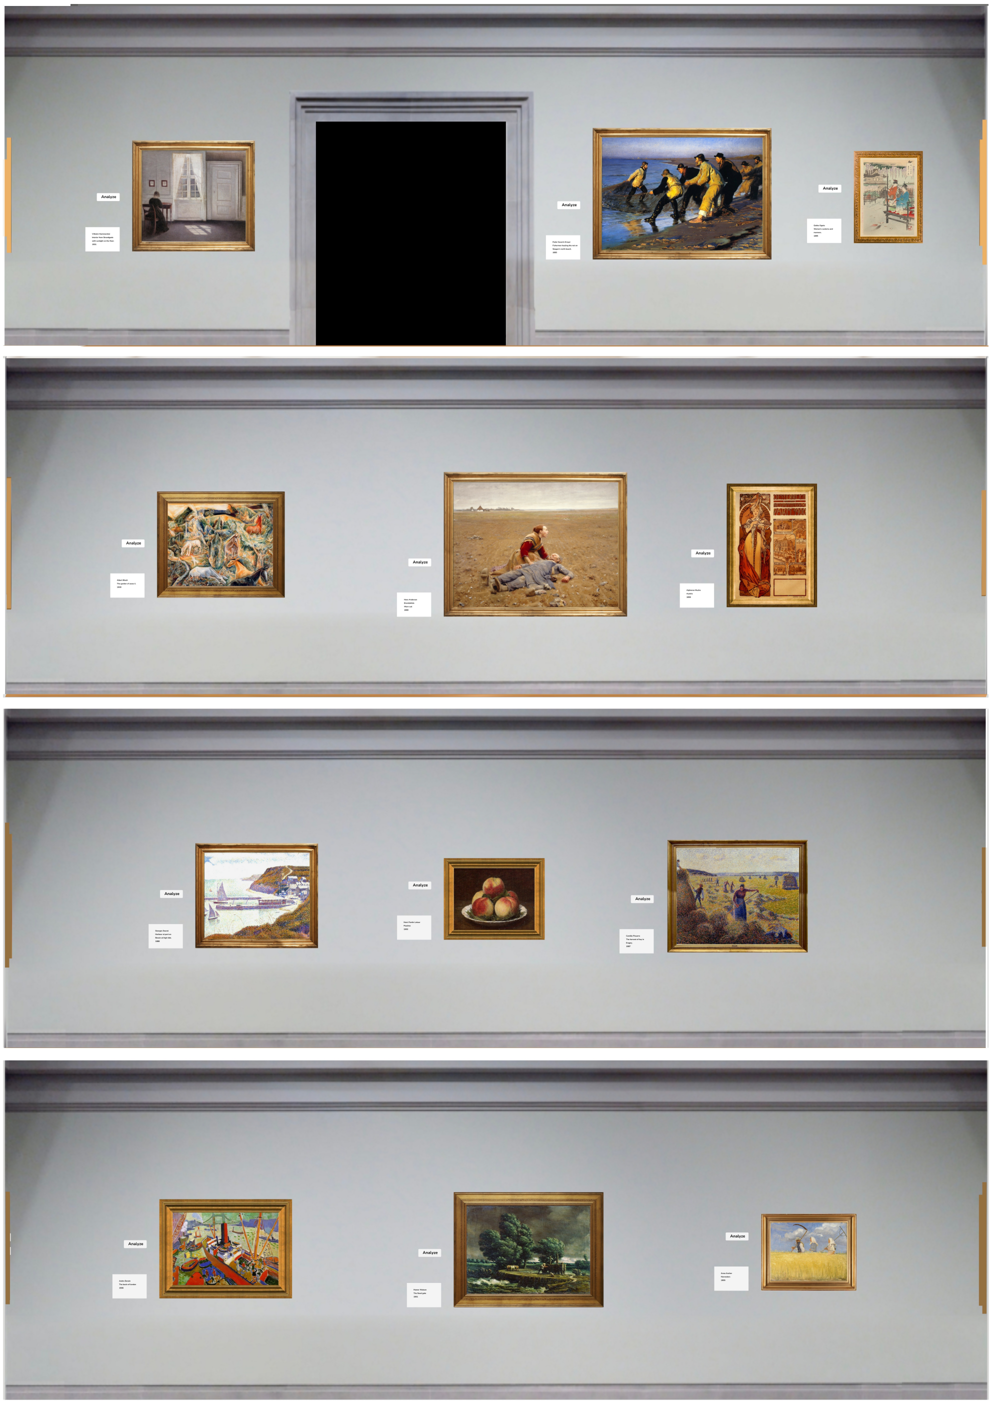

Fig. 6. The walls of the gallery room.

Fig. 7. The walls in the exploration room focusing on Henri Fantin Latour - Peaches.

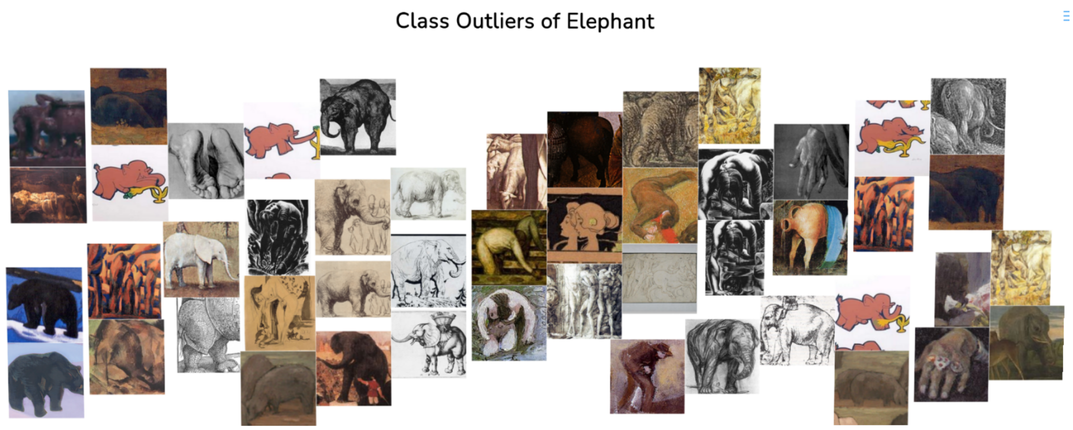

Fig. 8. Picture Cloud showing outliers of the class Elephant. Revealing multiple wrong classified instances.

style. To the upper right, images depict nature and water, while the left side shows humans in different situations. We also applied cluster mechanism to reduce the clutter, but the results were not satisfying.

C FOCUSED ART.

From our Gallery, 12 visitors started their tour by selecting the Danish painting "Fishermen hauling the net on Skagens north beach" by Peder Severin Krøyer (see Figure 4), 9 chose the Japanese painting "Print from series womens customs and manners" by Ogata Gekko. These paintings were also analyzed by the majority of people, but none of the paintings was disregarded. Except for the gallery paintings, the most analyzed paintings were "Interiør med syende pige ved vinduet" by Carl Holsøe (21 times), "Camille Monet in the Garden at Argenteuil" by Claude Monet (16 times), "Girl on the Beach" by Peder Severin Krøyer (13 times), "Besog hos bedstemor" by Hans Andersen Brendekilde (13 times) and "Sunlight in the blue room" by Anna Ancher (10 times). Except for French artist Claude Monet, most of the mentioned artists have Danish nationality, and are, like Carl Holsøe who was related to Krøyer and Hammershøi, closely connected in themes and style. Thus, these results reflect the virtual museum's capacity to "get deeper and deeper into different artists and styles", as one visitor remarked. Over 700 different images, painted by 202 different artists, have been observed on the Painting Wall. The most frequent artists that were not part of the Gallery were Carl Holsøe (48), Claude Monet (43), Salvador Dali (34), and Pablo Picasso (30). In total, the visitors searched for 105 different artists, the most frequent artists were Claude Monet (22), Vincent Van Gogh (17), Leonardo Da Vinci (9), Pablo Picasso (8) and Edvard Munch (6). The most frequent applied object filters were *Human face* (24), *Person* (23), *Clothing* (19), *Woman* (15), *Tree* (14), *Man* (11) and *Boat* (10). Similar statistics are recorded when visitors clicked on a bounding box on the Painting Wall.

Fig. 9. Two dimensional representation of the images with the style Ukiyo-e.

D QUESTIONNAIRE

Virtual museum visitor survey

1. What is your age?

- ◯ 18 - 29
- ◯ 30 - 39
- ◯ 40 - 50
- ◯ 50+

2. What is your gender?
○ Male
○ Female
○ Non-binary

3. You are a ... ?
○ Scientist (Humanities)
○ Scientist (Other)
○ Student (Humanities)
○ Student (Other)
○ Museum Visitor

4. How often did you visit a museum in a year before the COVID-19 restrictions?
○ 0
○ 1 - 2
○ 3 - 4
○ 5+

5. Have you ever visited a virtual museum before?
○ Yes
○ No

6. Have you used any of the following technologies before?
☐ Augmented Reality
☐ Virtual Reality
☐ 3D Games
☐ 3D Movies
☐ Other 3D Environments
☐ None of them

7. How long was your visit to our virtual museum?
○ less than 10 minutes
○ 10 - 30 minutes
○ 30 - 60 minutes
○ over 60 minutes

8. How would you rate your visit to our virtual museum?

Boring ○—○—○—○—○ Exciting

9. Our virtual museum is a valuable complement to a real museum!

Disagree ○—○—○—○—○ Agree

10. Our virtual museum is an important replacement in times real museums are closed.

Disagree ○—○—○—○—○ Agree

11. What did you miss compared to a real museum?

____
____
____

12a. Would you visit this virtual museum again?
○ Yes
○ No

12b. Why or why not?

____

____

____

13. How would you rate the presentations/visualizations?

Non-informative ○–○–○–○–○ Informative
Confusing ○–○–○–○–○ Intuitive

14. How would you rate the navigation/orientation?

Confusing ○–○–○–○–○ Intuitive

15. What did you learn from this experience?

____

____

____

16. What information about the paintings were missing?

____

____

____

17. Do you have any suggestions to extend and/or improve the virtual museum experience?

____

____

____