Article

The Replica Project: Co-Designing a Discovery Engine for Digital Art History

Isabella di Lenardo

Collège des Humanités, École Polytechnique Fédérale de Lausanne, (EPFL-CDH), 1015 Lausanne, Switzerland; Isabella.dilenardo@epfl.ch

Abstract: This article explains how the Replica project is a particular case of different professionals coming together to achieve the digitization of a historical photographic archive, intersecting complementary knowledge specific to normally unconnected communities. In particular the community of Art History researchers, brought together here in relation to their common methodologies in the practice of visual pattern research, became protagonists in the construction of a specific tool, the Morphograph, to navigate through the archive's photos. A specific research problem, the recognition of visual patterns migrating from one work to another, became the key to developing a new technology initially intended for a specific community of users, but with such a generic character in its approach that it could easily be made available to other uninformed users as learning by doing tools. The Morphograph tool also made it possible to demonstrate how, within a community, the partial expertise of individuality needs to be related to each other and benefits enormously from the knowledge densification mechanism made possible by the sharing. The digital context easily makes it possible to create tools that are specific in terms of content but generic in form that can be communicated and shared with even diverse and uninformed communities.

Keywords: digital art history; photographic archive; visual pattern recognition; co-design; cultural heritage

1. Introduction

The Giorgio Cini Foundation photo library in Venice preserves one of the largest historical art history photo archives in Italy and one of the most important in Europe.

In 2015, the Swiss Federal Institute of Technology in Lausanne (henceforth ETH Lausanne), the curators of the Giorgio Cini photo library and Adam Lowe of the Factum Arte Foundation began to discuss how to proceed with a massive and rapid digitization of the most substantial corpus of the photo library's photos, i.e., the index cards, about 1 million items.

It was a synergy of undoubted value and heterogeneity between the big players, since the photographic archive could boast a cultural heritage of exceptional value and eminently representative of centuries-old visual artistic memory, the ETH Lausanne embodied cutting-edge expertise in the field of computer science, in particular computer vision, capable of dealing with the historical and technical challenges posed by the corpus in question with the tools of state-of-the-art algorithms, and Factum Arte was the leading foundation in the field of facsimile reproductions of artistic objects, be they archeological, pictorial or sculptural.

This article explains the reasons for such a collaboration, describing the 'Replica' project that these institutions have co-designed and led, creating a virtuous and mutually enriching partnership, which involves digitally enhancing the Foundation's historical heritage, from digitization to valorization and online publication for communities of young scholars from all over the world.

This contribution also aims to show how this mode of co-operation can become a model for other cases where digitization of heritage, conducted jointly between a heritage

institution, a research institution and a stakeholder involved in the vanguard of digitization techniques, can follow a development pipeline from document digitization to making the document available to the public.

The Replica project also showed how all the issues raised by the digital transition of heritage can be addressed by finding effective answers in the plurality of profiles in the field and in the involvement of the student-researcher community to build an image search engine suited, in the first place, to the needs of the research world.

As a patrimonial institution, the Giorgio Cini Foundation brought the experience of archival and curatorial practice, linked to a deep tradition of expertise, and structured according to the logic of the ministerial approach and policies linked to the need for conservation and protection over the long term. In addition to this structural mission, there was also the need for transmission to the community of researchers and more generally to citizens [1]; it is no coincidence that the institution's motto was also inspired by a famous phrase by Mahler "Tradition is the handing down of the flame and not the worshipping of ashes" [Tradition ist die Weitergabe des Feuers und nicht die Anbetung der Asche].

The Giorgio Cini Foundation Photo Archive has an immense intrinsic historical value hosting different funds from the great photographers and photographic campaigns of the 20th century. Over time, the Giorgio Cini Foundation has brought together different photographic funds whose main vocation is to represent images of works of art that are representative of Venice and its historical and artistic heritage in the broadest sense of the term. The vocation of the Foundation's photo library is to provide visual encyclopedic access to Venetian cultural production, thus characterizing itself as having a very strong form of locational interest. At the same time, this characterization has always attracted scholars from all over the world who specifically seek out elements for research focused on Venetianness, at a granularity of content that more generic photo archives could hardly provide.

The ambition of the Replica project, from the point of view of the heritage institution, was to make available internationally, on a global scale, access to these sources, which are very specific documents of the Venetian heritage, in order to open a very local archive to global research, and to democratize the visual elements constituting Venetian history through the open access of the photographic corpus [2]. For this reason, the Replica project made use of the involvement of a community of young art history scholars who, through special workshops and training courses, helped to co-develop the main features of the image search engine.

As an academic institution, ETH Lausanne, and in particular the Digital Humanities Laboratory, incorporates research and teaching, the digitization project therefore reached out to the student community, training them in new critical practices of combining technology and historical questions. In fact, digitization required from the very beginning of the planned pipeline to adapt the process to the historical research questions to be addressed. The introduction of the community of researchers and students made it possible to pose digitization problems as transdisciplinary issues to which students had to find solutions. Project-related activities also made it possible to offer work placements in the heritage institution, within the framework of the project, which were useful for training new professionals.

From a research perspective, the ETH Lausanne Digital Humanities Laboratory was interested in experimenting with the possibility of manipulating a large corpus of images by analyzing the concept of visual similarity and pattern migration. In fact, the development work of the image navigation interface thematized this aspect, focusing on the analysis of visual patterns.

In addition, the collaboration with the Factum Arte Foundation has made it possible to define a new digitization system suited to the precipitousness of the corpus or works chosen. In particular, the Giorgio Cini Foundation wanted a modular scanning system in which the scanner could be a set of components that could be easily replaced and implemented to improve performance.

Factum Arte had in fact already worked with the Cini Foundation on the pioneering project of digitizing and reproducing on a 1:1 scale Paolo Veronese's Marriage at Cana, the original in the Louvre, in order to relocate the copy in the Refectory of the convent of San Giorgio Maggiore where it originally stood before its removal by Napoleon [3].

From Photo Archive to Digital, Codesigning the Project Pipeline

The initial aim of the project is to convert Cini's vast photographic archive into a digital archive. However, this is only the general concept summarizing a will, but then several aspects come into play related to the concrete elaboration of the pipeline, the organization of the work and the realization of the most interesting and useful digital enhancement system for the institution but also for researchers.

Hierarchies of roles were then decided in order to optimize the progress of the project. For each institution there was a project leader. In the case of the Cini Foundation, it was the IT manager who had to ensure a good scanning and preservation method within the parameters defined by the institution itself. In the case of the ETH Lausanne, it was a figure capable of understanding in detail the technical characteristics required and the development of digitization to answer some key research questions on the side of Digital Humanities and Art History. In the case of Factum Arte, responsible for the implementation of the fast and accurate scanning system, this was the technical manager able to understand the technical requirements of the institution but also capable of knowing the most advanced state of the art with respect to this technological challenge.

Before materially starting the operations, all partners involved had several workshops characterized by an open approach and brief to try to understand what the heritage institution's practices had been up to then and how they would like to improve their daily work and the provision of materials to the public [4]. At the same time, the ETH Lausanne recorded potential points to be developed technologically but also shared what was the state of the art with regard to computer vision applied to the management of photographic heritage collections and especially to art history. This exchange made it possible to put all actors on the same starting line of 'making it possible' [5]. Once the project's wish list of expectations had been established, the critical list was defined. These had different natures: technical logistical in the material process of digitization, optimization of the human resources involved in the various development phases of the pipeline, problems of conceptualization of the digital approach, implementation of the digital access interface, optimization of the digital interface for users.

Aware that the first attempt to define the list of critical issues could not be exhaustive and that others would arise during the operations, it was decided to increase the moments of confrontation and discussion by multiplying the debriefings for each phase of the project (Figure 1).
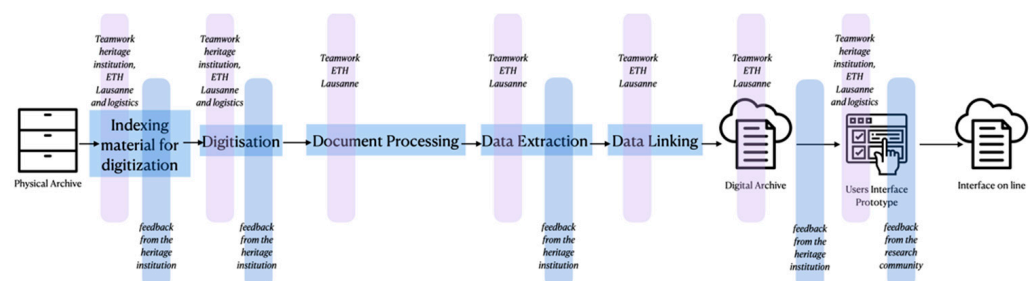

Figure 1. Replica Project General Pipeline (author: Isabella di Lenardo).

The only element that had to meet the requirement of stability, and on which there could be no turning back after operations had begun and were already well advanced, was the technical digitization protocol. To finalize it, in fact, preliminary tests were carried out before the scanner was even operational. Since the autumn of 2015, several tests have been carried out between to understand the technical characteristics of image acquisition,

the most suitable file format to store the metadata related to the archival location, the best method to index the physical items of the photographic archive, and the cardboard circumventing manual indexing.

It is very important in co-design projects to distinguish between processes that are circular and those that are one-off. The former can be characterized by multiple briefs with all parties involved because they can by nature be improved, indeed they derive their optimization from circularity. For one off, on the other hand, such as digitization, notwithstanding the fact that they could be started again, careful preliminary planning and adherence to the quality that one has set oneself very early in the process is necessary, otherwise starting everything all over again would be a great waste of time. Furthermore, in digitization projects there are phases in which quality feedback is automatic and does not require conceptual feedback, such as document processing and data linking. In-depth feedback is absolutely necessary for the final output, i.e., for the 'product' of valorizing and empowering the 'heritage capital' that the institution wants to value and the research institute to study.

This is why a lot of feedback is needed once the pipeline has reached the stage of generating the digital archive (Figure 1). First of all, there are inalienable control parameters such as the 'coverage' and 'completeness' of digitization. In other words, it has to be checked whether all analog elements have been converted, e.g., no images to be digitized have been forgotten and the metadata of the archival location have been correctly reported. Then there is the information inferred through document processing, i.e., a step following digitization, which must strive for completeness as far as possible. These checks, which are semi-automated, serve as feedback for the heritage institution to eventually understand which documents need to be re-digitized and for which a set of information can be completed manually.

Once these aspects have been assessed and told in brief, the digital collection can be considered 'finished'. It can then already be made available to the usual public with a simple 'demonstrator viewer' or disseminated through an ".api" of the institution. In our case, however, it was planned to create a specific search engine for the photographic archive that could not only be a traditional 'collection manager' but also allow searching through clusters of images that are similar to each other from a formal point of view without only classifying between metadata. For this several feedback from the heritage institution were necessary and as a complement also a series of workshops with art history researchers to present the idea and get opinions on which tools were useful for the scholarly community. This made it possible to build a circular approach in which at each debrief of these communities, not always in agreement with each other, the most relevant comments were implemented in the realization of the search engine [6]. The co-design certainly allowed a more comprehensive approach to the realization of the interface. All transversal observations were implemented to make the interface more useful for conservators but also for scholars. Usually, in fact, heritage actors, in particular photographic archives, revert to IT professionals to whom they commission, often with tight budget constraints, collection management software that is rather standardized and does not use state-of-the-art technologies, ignoring the dynamics of use of researchers and the public who should instead be involved as essential and final users of the digital archive.

2. How to Massively and Quickly Digitizing a Photo Archive

2.1. Co-Designing the Scanner

The co-definition of digitization goals allowed us to prototype an effective and sustainable pipeline. The photo library's cardboard had to be digitized at a much faster rate than was commonly practiced. But digitization was only the first step in a pipeline that also had to include, as a continuous workflow, the transmission of the images to the ETH Lausanne engineers in charge of processing them immediately in order to proceed with the extraction of the content and its study. In each part of the workflow, the skills of different professionals came into play. It needed a system capable of semi-automatically digitizing as many items

as possible per minute, knowing that each cardboard contained valuable information on the front and back and that both had to be stored. The resolution had to be defined according to current international criteria and the practices already tried and tested by the conservators of the photo library. If we had been able to digitize 1000 cardboard for five hours a day, it would already have been an extraordinary achievement, but it was still not enough, we had to reach a possible average of 1500 cardboard per day to be able to finish in the two years we had given ourselves and which could be financed by the project. In addition, the need to digitize in recto and verso posed a further challenge and demanded an ergonomic system for operators. Another problematic element in contrast to speed was the manual input of metadata, something usually practiced by conservators, and evidently difficult to apply in a system that was intended to be fully automatic or semi-automatic in which the machine would have to process all the cardboard consecutively without interruption.

These elements were resolved when Adam Lowe thought to take inspiration from a famous image taken from the book 'Le diverse et artificiose macchine' compiled by Agostino Ramelli, an engineer at the service of the French monarchy, and published in 1588. The machine taken as an example was a writing desk at the service of humanists in libraries, capable of holding several books at once for the reader (Figure 2).
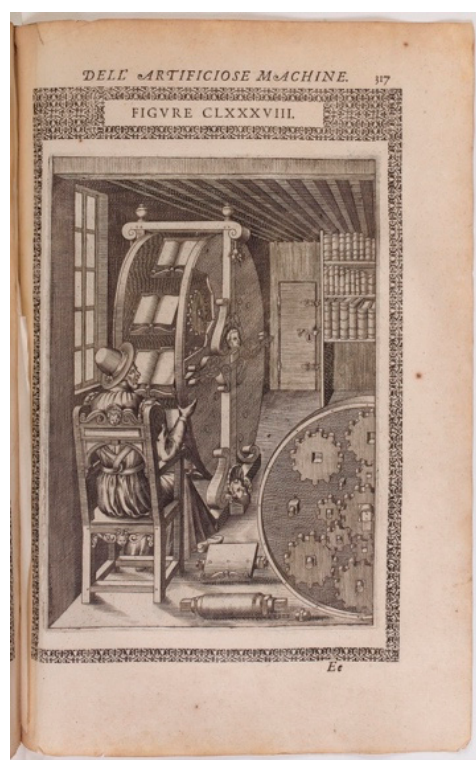

Figure 2. Agostino Ramelli, Delle artificiose Machine, 1588, p. 317. Image in CC0-Creative Commons.

A rotating system allowed multiple items to be read at the same time. The solution could therefore be a system similar to a rotating conveyer belt, capable of digitizing the same item simultaneously in recto/verso (Figure 3). A further requirement was sustainability and autonomy in customizing the device. It was necessary to produce a modular scanner whose constituent elements could be replaced by choosing better ones, corresponding to evolving performances, and which were made up of interchangeable parts available on the market without having to go through specific production companies. Last aspect: the possibility of associating the scanner with scanning software produced for the occasion and specific to our collection indexing needs.
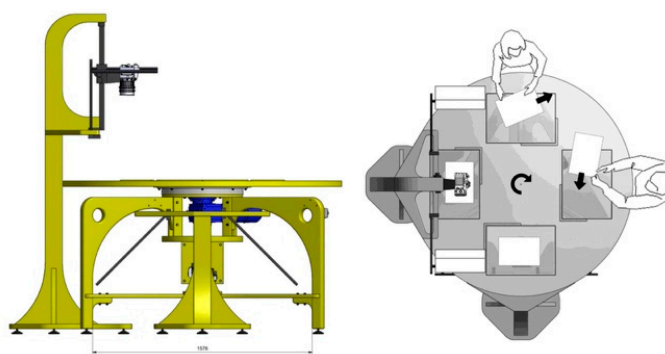
Figure 3. Left: Technical drawing of Recto/Verso scanner; Right: Birds-eye view of the Scheme 1. documents per day on average. The presence of two people, both close together, ensured better social interaction of the team, improving the quality of work that risked deteriorating due to the repetitive aspect of the operation. The design of the scanner and the components were put in open source by Factum Arte. The scanner, after quality control, obtained the European 'ISO' safety certification. The whole system was designed to make it replicable for other institutions wishing to maximise their digitisation experience by reducing costs and maintaining a high level of technical performance, without being forced to choose between companies currently on the digitisation market. Image taken with the kind permission of Factum Arte Foundation.

The scanner constructed consisted of a circular rotating table 2 m in diameter with four transparent glass plates on which to place the cards. The maximum possible format for scanning was 594 × 420 mm, an A2 format.

The rotary table is controlled by a precision motor with variable speed that allows the uninterrupted digitization of 1 image every 4 s.

A team of two people is required to carry out the scans, one of whom places the images on one of the glass plates and, at each complete turn of the table, two cameras, one positioned above and one below the table at a distance of 180 degrees from each other, first digitize the front and then the back of the article. The document continues its rotation and a second operator picks it up from the shelf. A sensor system calculates the position of the document and detects when it is placed on the glass surface. The data acquired by the cameras are immediately stored in a dedicated server. The cameras are replaceable if necessary: those prepared for the scanner allowed a document resolution of 400 ppi (5424 × 3616 pixels).The flash units were designed and engineered by Factum Arte to provide the lowest possible light level for a high-quality image, minimizing glare.

2.2. Co-Designing the Workflow Inventoring for Digitisation

The first part of the digitised collection, consisting of cardboards (Figure 4), collected in drawers, did not have an inventory. This case is not unique and that indeed it is very frequent, especially for large collections, not to have inventories or to possess inventories that one may want to re-check because one no longer considers the information given to be reliable [7].

Therefore the research activities of the ETH and the professional knowledge of the curators conservators of the Cini Foundation had to find a quick but extremely reliable method to index the images to be scanned. The curators have the archival habit of proceeding by indexing item-by-item in a semi-manual way, inserting the information considered key to their institution and/or those that meet needs defined by the need to index each item according to nationally defined criteria, for example by the Ministry of Culture. In the case of a photographic archive, there are also conservation needs as each image is in itself not only representative of the importance of the object represented, but a document of intrinsic patrimonial value [8].
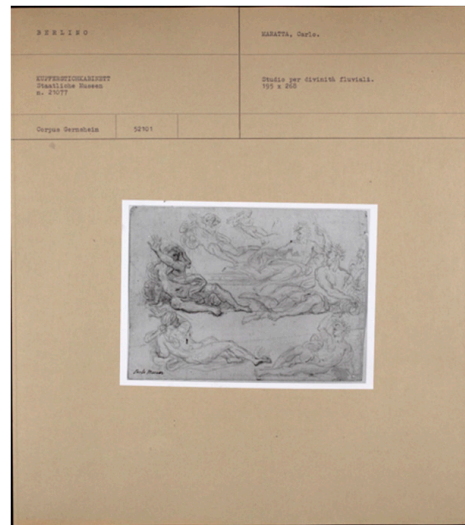
Figure 4. Exemple of Cardboard from Cini's photographic archive ©EPFL Digital Humanities Laboratory.

The card description is therefore often extremely detailed and impose a long and patient work on the curator and archivist to describe the physical and content of the object to be inventoried. In this sense, the practice of inventorying may be at odds with a massive digitization approach whose timeframe does not correspond at all to that required for such a descriptive practice [9].

Imagining then that each single item could be inventoried with an archival work of 20 min per item, a time that is already very fast and not at all likely, about 6000 items per year would have been indexed (upward estimate) by one person, and it would have taken 166 years to inventory 1M items making up the collection. It is well understood that drawing up a comprehensive inventory prior to digitization makes it effectively impossible to complete the digitization in a reasonable timeframe. The inventor with his manual skills and knowledge represents the bottleneck of the digitization chain. Therefore, if the need is to digitize in order to preserve and conserve, the production of a very detailed inventory in parallel jeopardizes the efficiency of the digitization itself, thus making it impossible to complete digital preservation. It was therefore planned to proceed by successive rounds of post-digitization inventories, gradually becoming more specific and accurate. For this reason, the pre-digitization inventory consisted of documenting the archival position of each item.

It was decided to split the digitization pipeline into two different processes in which the recto and verso of each item underwent two different content extractions (Figure 5).
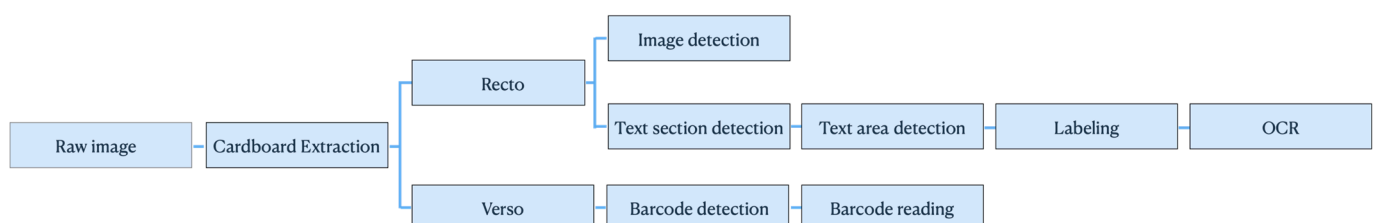

Figure 5. Document Digitisation Pipeline.

The *recto* of the image file contains two different types of information: the image and a block of text structured in cells with some descriptive fields.

The *verso* contains a barcode generated within the pre-inventory of the project, in which the archival location of the scanned item is described (Figure 6).
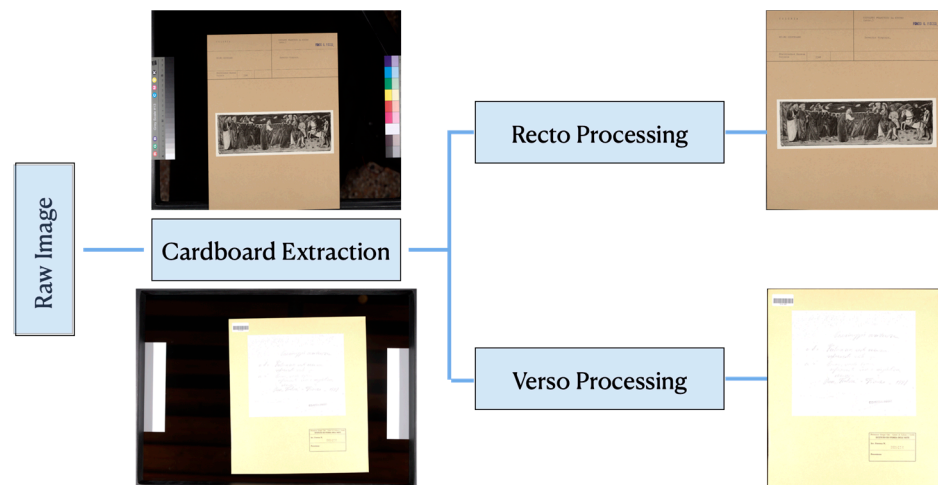
Figure 6. Schema of Recto and Verso Extraction Processing.

This is in fact the key information: the address of the item that makes it possible, in the topographical vastness of an archive, to find it and to validate the equivalence between the scanned image and its material consistency [10].

Cardboards are physically placed in drawers. It was therefore decided to assign a barcode that identified the archival position relative to the drawer and the position of the cardboard inside the drawer.

In this way, through the automatic reading of the barcode with the scanner software, a json file was generated containing the archival position data associated with the scanned image. The various operators were able to move forward quickly by applying barcodes on the *verso* of items, only having to verify the exact correspondence between the item, the archival location and the generated code (Figure 7).
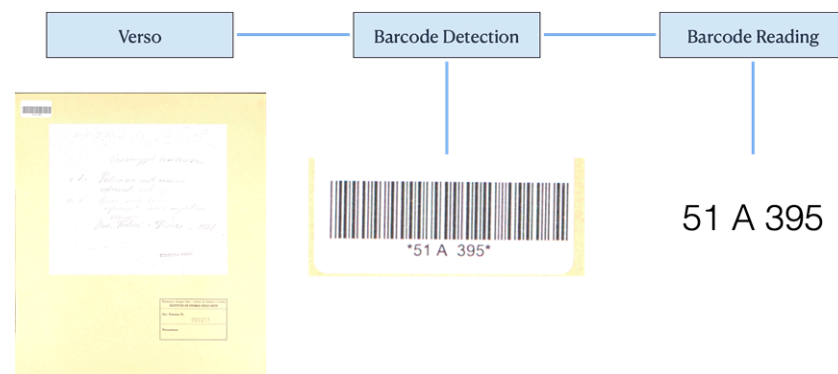

Figure 7. Schema of Cardboard *Verso*—Barcode Detection.

With regard to any content information important for indexing, the descriptive elements in each file were automatically extracted after assigning semantic labels [11].

In fact, each cardboard presented a grid in which certain information provided by the curators appeared in a structured manner: place of conservation of the object represented in the photo, photographer's name, inventory number, title and technique of the work, author (Figure 8).

This descriptive space was identified and then labelled into categories and then structured into metadata describing the content of each item.

Once digitisation was complete, the question arose of opening up access to the photographic archive to a community that was not local, such as the scholars who frequented the library, but global.
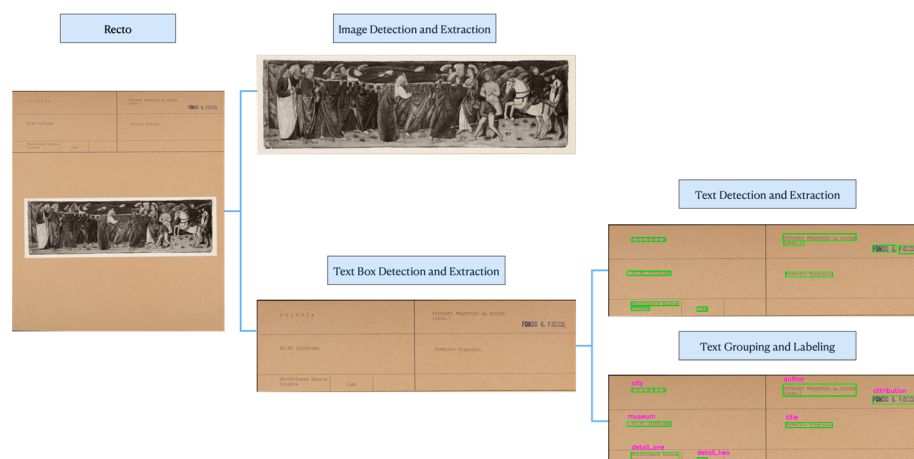
Figure 8. Schema of Cardboard's element extraction.

This is why we opted for the creation of an IIIF server that would allow access to the images and files in an international format [12]. To do this, we hosted the images in an IIIF format by generating a manifest for each physical drawer that collects a number of corresponding documents. Thanks to the manifest, any IIIF vewer can navigate through the original cardboards in the same order in which they are physically located in the physical archive. In digitisation, it is indeed important to respect not only the arborescence of the collection holdings but also the physical 'topography' of the digitised objects. Their location, in fact, is not random and follows a curatorial, or historical, order imposed by the logic of the nature of the items preserved. By respecting the visualisation of the actual location in the archive, the organisation of the collection as a historical result and the manifest of the extracted photograph are better understood, effectively linking the image to its primary source in the original archive. Photographs are thus related to each other by their archival position, and this position can only be grasped by visualising the items through the manifest of each drawer, otherwise all images in the collection would be presented at the same logical level, as if it were a flat indexing system.

3. Improving Curators' Knowledge on the Collection

Aligning Entities Extracted for Enriching Information

Extracting information from the text box by labeling it is already a very important achievement, but it still does not allow us to really identify the content of the text unless we link the extracted information to entities from a defined dictionary.

More precisely, all we have transcribed is raw text, words that in themselves have no meaningful meaning and do not tell us much about what has been digitized. Curators should be able to know which photographic collections are present, which artists are represented in the collection and how many photographs they have. They should be able to filter the images of geography or by century. In order to enrich the cognitive value of the collection, we agreed that, first of all, the 'Author' field was the most interesting to examine. Indeed, even if it is only the name of an artist, for a curator or art historian this same name is automatically associated with the corresponding period as well as the city in which the artist was active, already providing an approximate context on the creation of the artwork. To extract this information we therefore need an external knowledge base that indexes the artists and identifies them on existing references. To solve this record-linking problem we relied on the Getty Union List of Artist Names [13], which among the Linked Open Data on the web represents a comprehensive reference of artistic identities. The alignment resulted in a satisfactory coverage of 73.8% of the collection. The remaining percentage corresponds to a different use of the author field. In addition to trivial OCR errors, there appear many century mentions i.e., "16th century", multiple attributions, and many "Anonymous" which make a realignment to a known identity impossible [14].

We have discussed the alignment method and the effectiveness of its results in detail in specific contributions, and mention here only those elements that are useful for understanding in which aspects the alignment has been able to empower conservators and scholars of the Cini Foundation collection [12].

Alignment already makes it possible on a global scale to visualize the geographical distribution of the artists included in the photographic archive (Figure 9) and their chronological prominence. On a finer analysis including the places of activity of the mentioned artists, for example, by realigning the entities on the Wikidata knowledge [15], it is also possible to study the places of artistic production most represented in the collection and the geographical parabolas of the Venetian artists in the corpus in a comparative manner. The alignment of the names of the sub-collections that make up the corpus of 1M photos then provides a quantitative image of the composition of the collection, i.e., the names of the funds and photographers that make up specific corpuses, homogeneous in terms of authorship and often also in terms of subjects. In these sub-collections, the chronologies of the works represented can also be studied to get a precise idea of the chronological distribution (Figure 9). It was discovered, for example, that in the "Gernsheim" corpus there are no medieval items, which characterize the "Alinari" fund instead (Figure 10).
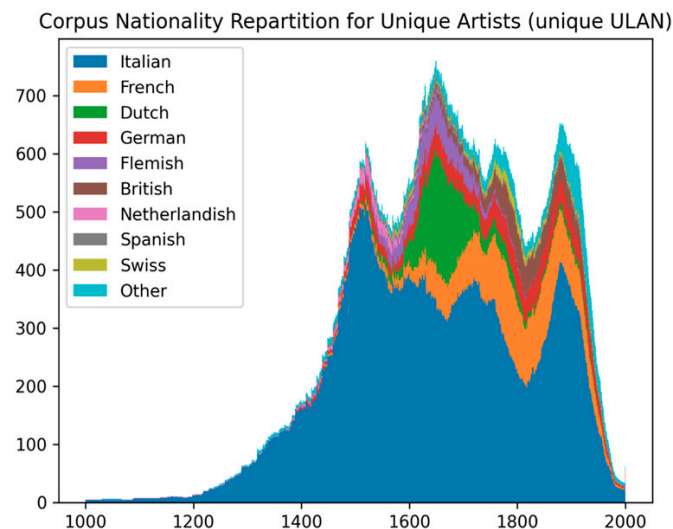

Figure 9. Corpus Nationality Repartition for Unique Artists based on ULAN realignment. (author: Paul Guhennec).
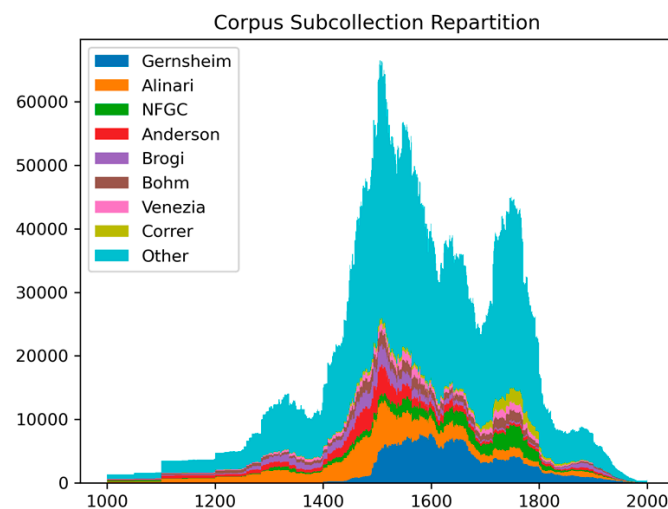

Figure 10. Corpus Temporal Subcollection Repartition (author: Paul Guhennec).

The alignment allowed for a distant view of the collection, while retaining the possibility of a granular analysis in order to grasp the general trends then questioning them qualitatively down to the individual item. This ensemble analysis allows the totality to be grasped independently of the curators' individual skills in analyzing the collections cardboard by cardboard in a partial manner, without ever being able to manipulate the totality of the information spread over the almost 1 million items. Which photos were collected over time by this institution and with what content logic, these are some of the questions that can be answered thanks to the alignment, allowing the history of the photo archive to be revealed, something that is difficult to reconstruct without the direct testimony of previous curators. The analysis of the evolution of the logic of acquisition of photographic funds, applied in general, would thus allow each photographic archive to know its history over time, with the undoubted enhancement of the heritage.

4. Searching for Visual Pattern Propagation: An Interface and Learning System to Drive a Collective Research Effort

The Morphograph: An Interface for Sharing and Updating Knowledge

From the research point of view, the ETH's interests were to design an efficient information extraction pipeline that respects the specificity of the photographic archive and to build a digital tool that would allow images to be searched not only with textual keys, but by comparing them with other images. Searching for images with images means revealing visual connections between works of art [16]. The definition of visual connections, e.g., by analyzing the propagation of certain specific visual motifs, patterns, is extremely important for art historians: for example, it substantiates knowledge of the relationships between artists, the fortunes of a particular work and its collecting interest. The navigation tool to be constructed therefore had to search for this propagation of patterns in the photos of the works of art in the collection [17]. The corresponding search engine then had to reflect, for example, the need for invariance with respect to the style of the represented artwork or its medium.

To succeed in programming such a search engine, the problem of similarity, and pattern propagation, had to be formalized from a computational point of view. We cannot go into the technical aspects of this approach and the results in detail in this paper, so we refer to the reference publications that explain the details.

The fact that two works of art appear to be related to each other is in fact not only historical visuals, but can also be discretized using a descriptive approach. Similarity is in fact dictated by an articulation between local and global equality: we can find the reuse of a global composition in several works, while the local elements do not match and vice versa there can be local elements that are reused in several works, while changing the rest of the composition. We could therefore define the visual relationship between two works of art as a mixture of global and local elements [18]. Currently there is still no ontology describing visual connections or making explicit standardized classes of annotations. In order to decode the different cases of visual relationships of similar elements on a local (detail) or global (composition) scale, it was first necessary to define the relevant elements, from an art history perspective, to be classified.

More precisely, we were interested in documenting the propagation of visual forms that were sufficiently identifiable that we could guarantee that there was a physical chain of actions linking the works in which they appeared. It was not a question of documenting styles or manners but of following the invention and propagation of singular forms [15].

In order to proceed with the identification of these connections, it was therefore decided to create a tool that would allow the art history research community to effectively annotate bilateral relationships between elements. The most effective way was to arrange the similarities between the images in the form of a graph to be annotated, which we called the "Morphograph" (Figure 11) [16].

Figure 11. Morphograph of Titian's artworks in the Fondazione Giorgio Cini dataset.

The use of a graph is particularly effective because it is a standard data structure for computers and easy for humans to manipulate. Art historians can easily edit, modify and visualize it. The most important aspect of the Morphograph is that every historian can annotate the images he or she knows, those that are part of his or her preparation, and that represents a part of the whole, in the specific case of the Cini photo archive, and more generally of the whole set of images of works of art. All this partial knowledge can flow into the Morphograph, which in this way becomes a mirror of the collective knowledge of a community of researchers all traditionally focused on partial aspects of the entire history of art. The graph itself thus becomes an object of knowledge [17]. The training of art historians traditionally makes them focus only on certain centuries, certain specific authors, even certain precise iconographic themes. This, however, prevents a global and completely shared understanding of the history of image transmission, which must instead be thought of as a space-time continuum.

The Morphograph allows the researcher to contribute to a system that goes beyond his own knowledge and the community to create a common and entire knowledge, even revisable if desired, from partial knowledge.

We therefore brought in art history researchers who, for various research reasons, were more interested in mapping pattern propagation across broad chronologies and artistic geography. We drew on their method of identifying significant details and migrating patterns in a work. We focused on the community of practice that identified their method and knowledge [18]. Identifying similar patterns in images can be an extremely generic and inclusive concept, but what interested in our study, characterized by the need to answer this question from an art-historical point of view, was to research the relevance of the pattern propagation found (Figure 12). This required a community informed in the same way, sharing the same methods and results, the same practices, common goals and common language [19,20]. Therefore, several workshops were conducted to enable scholars to annotate works in a relevant manner and to safeguard these annotations to train the system to recognize what was a visual similarity and propagate the concept of relevant information.

The method used is as follows: researchers annotate visual links between images in the Morphograph, which is then used to train a visual similarity system based on a deep learning model. This similarity function is exploited by an interface/annotation system, which, by improving the links between images through annotation, allows annotators to find other elements to add to the graph [21].

Figure 12. Exemple of Art historian annotation work on pattern propagation (author: Carlotta Striolo).

This organization creates a positive feedback loop because the more connections that are identified in the graph, the better the visual model will be due to the increase in data provided to the training, the better the performance of the search engine will be, which will then allow even more visual links to be found (Figure 13).
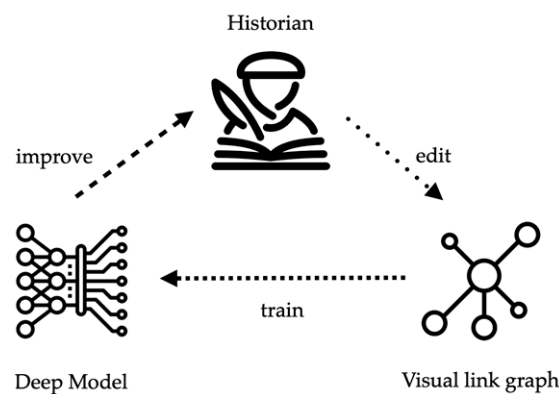

Figure 13. The loop of the annotation system.

The deep learning model uses Morphograph links to modify the similarity space between images using a method called "triplet learning" [22]. During the learning process, the space is progressively deformed to ensure that images that are linked by the Morphograph are closer than those that are not. It is important to note that this similarity based on visual pattern propagation in art history is very different from other measures used in computer vision. Indeed, some pairs of images, such as, for instance, the preparatory drawing of a painting and the final painting, may share an identifiable visual pattern, while having nothing in common with respect to textures or colors. Visually the common visual pattern occupies only a small place in the image and yet it is this element that is crucial to recognize. For this reason, the deep learning model must learn to ignore some aspects that connect the two images to focus only on the presence or absence of patterns that are the reason of their connection in the Morphograph [22].

Based on the similarity space of the deep learning model, the search engine can propose, for each new search, suggestions that may contain similar motives to the work used as a search key. Exploring hundreds of thousands of works, the researcher can thus find works that he or she does not know. When the search identifies works with similar patterns but not yet documented in the Morphograph, the interface allows registering them and thus to train the whole system again, recalculating the similarities of all the works on these new bases (Figure 14). Thus, the searches and discoveries of each user are progressively integrated not only in an explicit knowledge structure, the Morphograph, but also in the structure of the similarity space used by the search engine. The tool increases the work of each researcher by integrating the work of all.
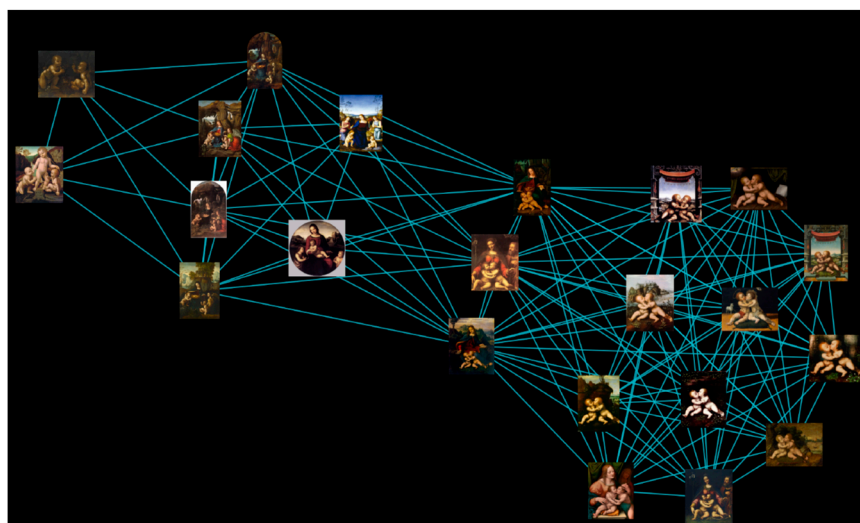
Figure 14. Visual connections of pattern propagation in the Morphograph from the Virgin of the Rocks (author: Benoit Seguin).

In less than two years, thanks to the work of the art historians and the dynamics of the system, more 6000 visual links were documented in the Morphograph, studying the Cini corpus. Initially, historians reported in annotation the knowledge they had already acquired based on the specific historiography with respect to this topic but then the Graph was extended based on the suggestions of the search engine.

It is important to realize that building a structure such as the Morphograph that aims to record all the visual links between the works of art contained in it is an impossible task for an isolated researcher to achieve. Not only is it unlikely that a researcher could have seen and remembered more than a few tens of thousands of works, while the number of works photographed in the Fondation Cini's photo library exceeds one million, but also doing a systematic search for common patterns is practically impossible for a researcher for these orders of magnitude of documents. The Replica system provides a tool that allows each researcher to judge as if he or she had access to the integrated knowledge of all his or her colleagues, a system whose performance improves each time someone uses it to document a new visual connection between two works. Each new connection is useful for oneself but also for others [23].

5. Conclusions

The trajectory of the Replica project illustrates several levels of interdisciplinary collaboration with the objective of designing a new research method in digital art history.

In the first phase, the construction of a new scanner and a new method to transform a physical archive into a digital database, researchers, each expert in their field, with different but complementary backgrounds and experiences, combine their knowledge to produce, through a form of co-design, an innovative approach.

In the second phase, the construction of a search engine fed by the explicit knowledge developed by the art history researchers allows a collective co-construction of a superior knowledge. A feedback loop offers the opportunity for everyone to continuously improve the system.

The last phase, not developed in this article, should extend this virtuous dynamic beyond the restricted community of art history researchers. The space of visual similarity progressively transformed by the information of the Morphograph can be explored by algorithms which try to extract groups of works automatically, connected by visual links, not yet documented by the Morphograph. It would not be a visual search engine, but an automatic system to "discover" pairs of linked works. These discoveries must then be validated by verifying the presence or absence of visual links in the identified works.

This is where the general public can get involved. If it takes a great knowledge of art history to successfully query a visual search engine, pursuing a specific line of inquiry, the verification of the presence of a migration of form between two works is potentially more immediate. The last step of the project will be to include them in the project through an extended collaboration between the machine and a large number of communities who are curious and motivated to progress in solving art historical questions. This proves that the construction of specific tools that combine a discipline-specific approach to a more general versatility makes it possible to answer research questions that are potentially transversal to the communities that traditionally study them with their own methods. It is foremost a matter of opening up and making concepts and tools available to suggest answers in which users from different communities benefit, who learn to sharpen their eye for the specific problem, and the community of art historians who benefit from a global expertise on which they can always comment later on the relevance of their selections.

The co-design of the Replica project was characterized by many discussion phases, especially in the design of the final prototype. Finding an initial convergence point on the general pipeline was not difficult, but converging on a material prototype of the final realization required much discussion and rejection of certain interface development hypotheses. In the case of the Replica project, moreover, it was not possible to start from an already existing prototype precisely because it had to be packaged over time thanks to the feedback of the curators' and researchers' community, and it was precisely in these latter aspects that its intrinsic value lay. Moreover, in this case, the curators of the photographic archive would probably have welcomed as a non-critical imposition the presence of a prototype, even if only a conceptual one, which could then be perfected by adapting it to the specific case. In fact, an element not to be underestimated in co-design projects for cultural heritage is that all participants wish to feel protagonists and believe they have a specific legitimacy, by authority or knowledge, that cannot be avoided. The development of co-design between institutions that are profoundly different in their dynamic conception of heritage has also made it possible to bypass approaches whereby heritage is conserved as "an unchanging monument of the past" [24]. and the presence of the community of end users, which in our case are mostly art history researchers, in other projects it is the general public of citizens, has made it possible to respond to a dynamic and functional vision of this heritage [25]. Indeed, it should not be forgotten that the only way to preserve cultural heritage is to find current and interesting forms for its transmission.

References

1. Cuno, J. Open Content, An Idea Whose Time Has Come | Getty Iris. Available online: https://blogs.getty.edu/iris/open-content-an-idea-whose-time-has-come/ (accessed on 13 June 2022).
2. Edwards, E. Aux Portes du «Trésor de la Connaissance». Photographies, Histoire Locale et Bibliothèques Publiques, Transbordeur: Photographie Histoire Société, No. 1. 2017. Available online: https://transbordeur.ch/fr/2017/photographies-histoire-locale-et-bibliotheques-publiques/ (accessed on 13 June 2022).

3. Latour, B.; Lowe, A. *The Migration of the Aura, or How to Explore the Original through Its Facsimiles, Switching Codes*; University of Chicago Press: Chicago, IL, USA, 2011; Available online: https://www.degruyter.com/document/doi/10.7208/9780226038322-017/html (accessed on 13 June 2022).

4. Ciolfi, L.; Avram, G.; Maye, L.; Dulake, N.; Marshall, M.T.; van Dijk, D.; McDermott, F. Articulating Co-Design in Museums: Reflections on Two Participatory Processes. In Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, San Francisco, CA, USA, 27 February–2 March 2016; pp. 13–25.

5. Avram, G.; Ciolfi, L.; Maye, L. Creating Tangible Interactions with Cultural Heritage: Lessons Learned from a Large Scale, Long Term Co-Design Project. *CoDesign* **2020**, *16*, 251–266. [CrossRef]

6. Bossen, C.; Dindler, C.; Iversen, O.S. Impediments to User Gains: Experiences from a Critical Participatory Design Project. In Proceedings of the 12th Participatory Design Conference: Research Papers, Roskildem, Denmarkm, 12–16 August 2012; Association for Computing Machinery: New York, NY, USA, 2012; Volume 1, PDC '12. pp. 31–40. [CrossRef]

7. Estelle Blaschke Rarity vs. Ubiquity in Photographic Collections. In *On Alinari. Archive in Transition*; A+MBookstore: Milano, Italy, 2021; pp. 128–136.

8. Caraffa, C. Photo Archives and the Photographic Memory of Art History, Munich/Berlin 2011, Summary. Available online: https://www.academia.edu/29908716/Photo_Archives_and_the_Photographic_Memory_of_Art_History_Munich_Berlin_2011_Summary (accessed on 9 June 2022).

9. Bülow, A.; Ahmon, J.; Spencer, R. *Preparing Collections for Digitization*; Facet Publishing: London, UK, 2011.

10. Conway, P. Modes of Seeing: Digitized Photographic Archives and the Experienced User. *Am. Arch.* **2010**, *73*, 425–462. [CrossRef]

11. Oliveira, S.A.; Seguin, B.L.A.; Kaplan, F. DhSegment: A Generic Deep-Learning Approach for Document Segmentation, Infoscience, 6 August 2018. Available online: http://infoscience.epfl.ch/record/263065 (accessed on 1 June 2022).

12. Presentation API 2.0. Available online: https://iiif.io/api/presentation/2.0/#manifest (accessed on 13 June 2022).

13. The Getty Research Institute. Available online: https://www.getty.edu/research/tools/vocabularies/ulan/ (accessed on 1 June 2022).

14. Seguin, B.; Costiner, L.; di Lenardo, I. (Eds.) Extracting And Aligning Artist Names in Digitized Art Historical Archives. In *Book of Abstracts of Digital Humanities Conference 2018 Puentes-Bridges*; Red de Humanidades Digitales AC: Mexico City, Mexico, 2018.

15. Wikidata Dictionary. Available online: https://www.wikidata.org/wiki/Wikidata:Main_Page (accessed on 13 June 2022).

16. di Lenardo, I.; Seguin, B.L.A.; Kaplan, F. (Eds.) *Visual Patterns Discovery in Large Databases of Paintings*; Infoscience: Lausanne, Switzerland, 2016.

17. The Link to the Replica Search Engine for Images. Available online: https://diamond.timemachine.eu (accessed on 13 June 2022).

18. Seguin, B.; Striolo, C.; di Lenardo, I.; Kaplan, F. Visual Link Retrieval in a Database of Paintings. In *Computer Vision—ECCV 2016 Workshops*; Lecture Notes in Computer Science; Hua, G., Jégou, H., Eds.; Springer International Publishing: Cham, Switzerland, 2016; pp. 753–767. [CrossRef]

19. Graven, M.; Lerman, S.; Wenger, E. Communities of Practice: Learning, Meaning and Identity. *J. Math. Teach. Educ.* **1998**, *6*, 185–194. [CrossRef]

20. Cattaneo, C. Community of Practices. In *Encyclopedia of Sustainable Management*; Idowu, S., Ed.; Springer International Publishing: Cham, Switzerland, 2019; pp. 1–10. [CrossRef]

21. Seguin, B.; di Lenardo, I.; Kaplan, F. Tracking Transmission of Details in Paintings, in DH. 2017. Available online: https://www.semanticscholar.org/paper/Tracking-Transmission-of-Details-in-Paintings-Seguin-diLenardo/258692fa5bf09de1284e4b5835dde63f942cafc6 (accessed on 13 June 2022).

22. Seguin, B. *Making Large Art Historical Photo Archives Searchable*; EPFL: Lausanne, Switzerland, 2018. [CrossRef]

23. Danny, P. *Wallace, Knowledge Management: Historical and Cross-Disciplinary Themes*; Libraries Unlimited: Exeter, UK, 2007.

24. Smith, L. *Uses of Heritage*; Routledge: London, UK, 2006.

25. Claisse, C.; Ciolfi, L.; Petrelli, D. Containers of Stories: Using Co-Design and Digital Augmentation to Empower the Museum Community and Create Novel Experiences of Heritage at a House Museum. *Des. J.* **2017**, *20*, S2906–S2918. [CrossRef]