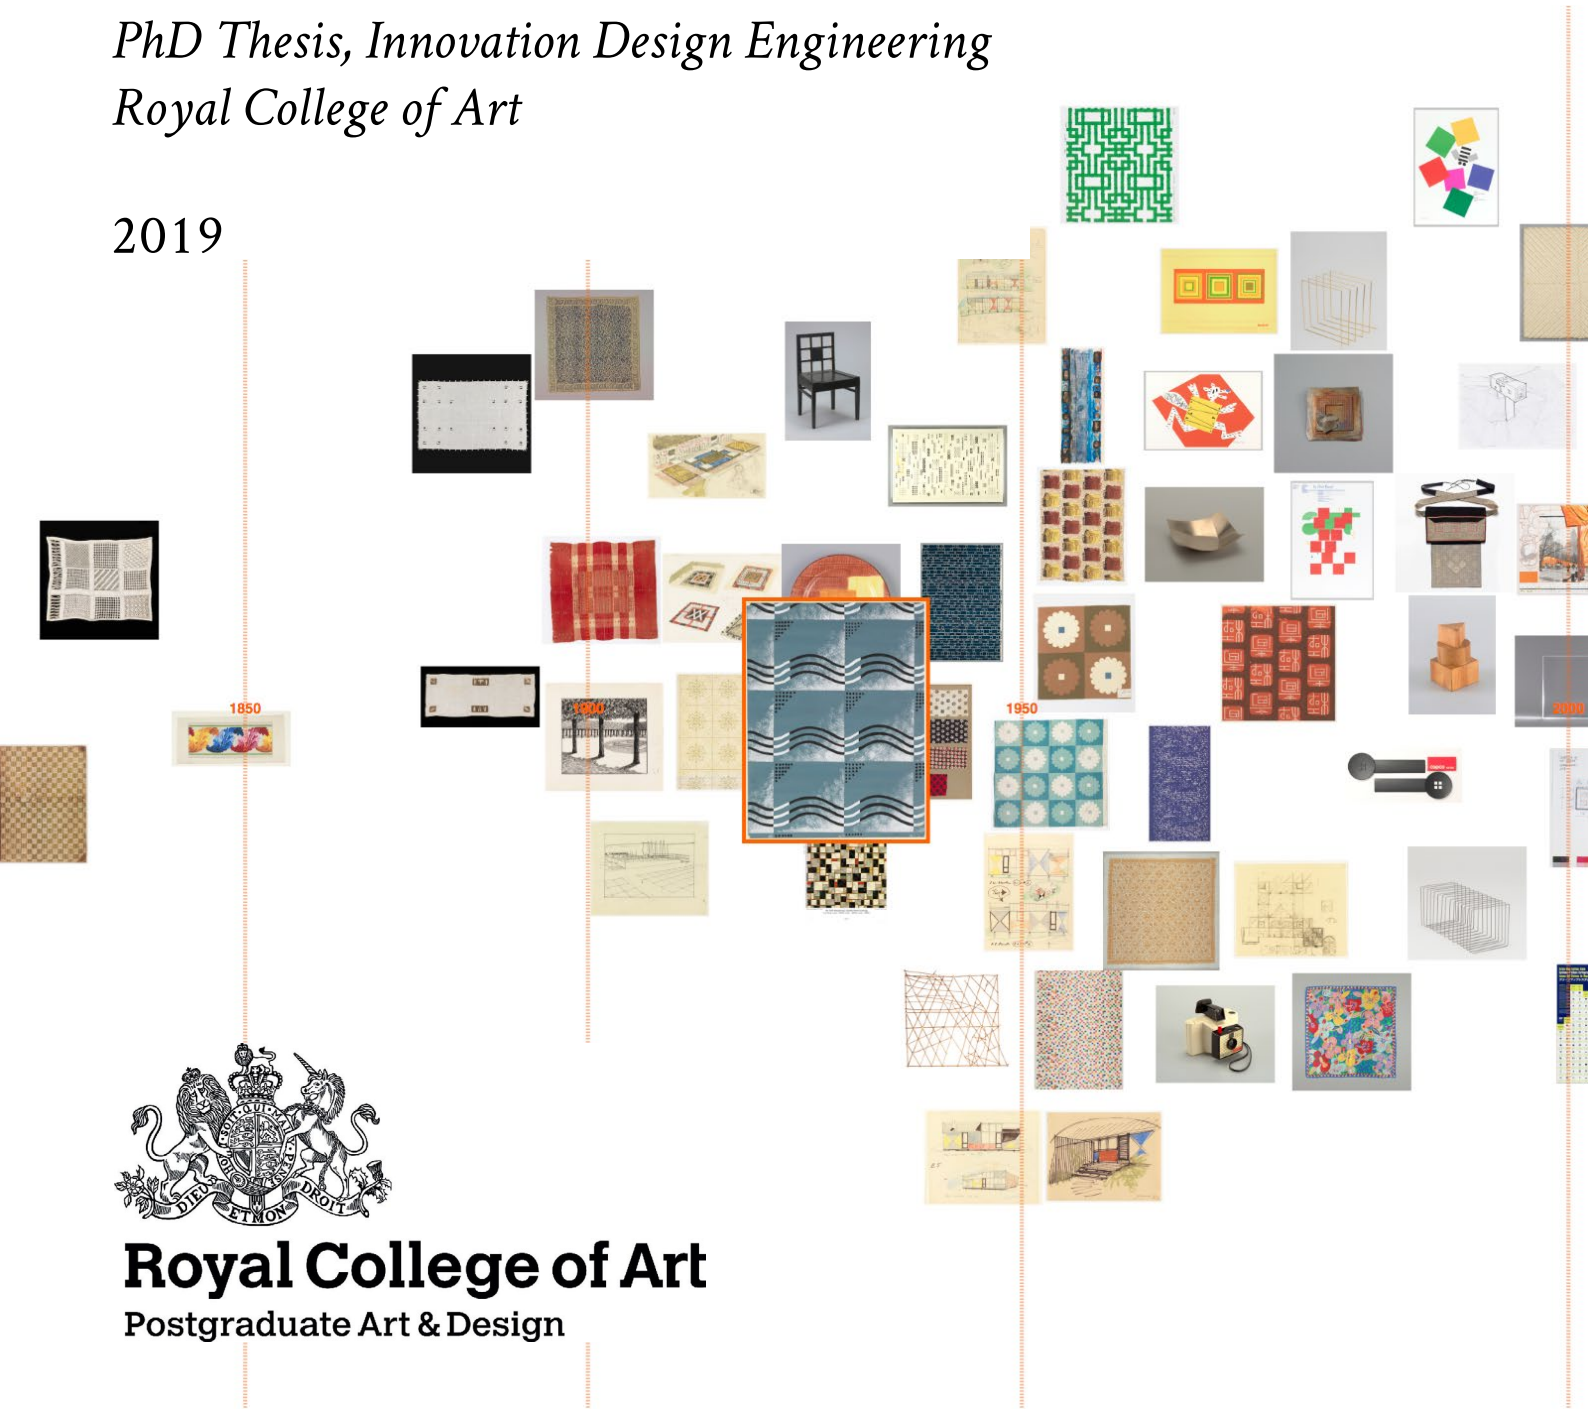


Timeline design for visualising cultural heritage data

OLIVIA VANE

*PhD Thesis, Innovation Design Engineering
Royal College of Art*

2019



Royal College of Art
Postgraduate Art & Design

Timeline design for visualising cultural heritage data

OLIVIA VANE

*PhD Thesis, Innovation Design Engineering
Royal College of Art*

2019

This thesis represents partial submission for the degree of Doctor of Philosophy at the Royal College of Art. I confirm that the work presented here is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

During the period of registered study in which this thesis was prepared the author has not been registered for any other academic award or qualification. The material included in this thesis has not been submitted wholly or in part for any academic award or qualification other than that for which it is now submitted.

Olivia Vane, 5th September 2019

Abstract

This thesis is concerned with the design of data visualisations of digitised museum, archive and library collections, in timelines. As cultural institutions digitise their collections—converting texts, objects, and artworks to electronic records—the volume of cultural data available grows. There is a growing perception, though, that we need to get more out of this data. Merely digitising does not automatically make collections accessible, discoverable and comprehensible (Terras, 2015:p.739), and standard interfaces do not necessarily support the types of interactions users wish to make.

Data visualisations—this thesis focuses on interactive visual representations of data created with software—allow us to see an overview of, observe patterns in, and showcase the richness of, digitised collections. Visualisation can support analysis, exploration and presentation of collections for different audiences: research, collection administration, and the general public. The focus here is on visualising cultural data by time: a fundamental dimension for making sense of historical data, but also one with unique strangeness.

Through cataloguing, cultural institutions define the meaning and value of items in their collections and the structure within which to make sense of them. By visualising threads in cataloguing data through time, can historical narratives be made visible? And is the data alone enough to tell the stories that people wish to tell?

The intended audience for this research is cultural heritage institutions. This work sits at the crossroads between design, cultural heritage (particularly museology), and computing—drawing on the fields of digital humanities, information visualisation and human computer-interaction which also live in these overlapping spaces.

This PhD adds clarity around the question of what cultural visualisation is (and can be) for, and highlights issues in the visualisation of qualitative or nominal data. The first chapter lays out the background, characterising cultural data and its visualisation. Chapter two walks through examples of existing cultural timeline visualisations, from the most handcrafted displays to automated approaches. At this point, the research agenda and methodology are set out. The next five chapters document a portfolio of visualisation projects, designing and building novel prototype timeline visualisations with data from the Wellcome Library and Victoria & Albert Museum, London, Cooper Hewitt Smithsonian Design Museum, New York City, and the Nordic Museum, Stockholm. In the process, a range of issues are identified for further discussion. The final chapters reflect on these projects, arguing that automated timeline visualisation can be a productive way to explore and present historical narratives in collection data, but a range of factors govern what is possible and useful. Trust in cultural data visualisation is also discussed. This research argues that visualising cultural data can add value to the data both for users and for data-holding institutions. However, that value is likely to be best achieved by customising a visualisation design to the dataset, audience and use case.

Keywords: cultural heritage data; historical data; cultural analytics; cultural informatics; humanities visualisation; generous interfaces; digital humanities; design; information design; interface design; data visualisation; information visualisation; time; timeline; history; historiography; museums; museology; archives; chronographics.

Publications, Conference Presentations and Exhibitions

Publications

Boyd Davis, S. & Vane, O. (2019) Design as Externalization: enabling research [Forthcoming]. *Information Design Journal*.

Vane, O. (2019) Cultural Visualisation and the Value of Simplicity. *EuropeanaTech Insight*, 11. Available from: <https://pro.europeana.eu/page/issue-11-generous-interfaces#cultural-visualisation-and-the-value-of-simplicity> [Accessed 19th March 2019].

Vane, O. (2018) Text Visualisation Tool for Exploring Digitised Historical Documents. In: *Proceedings of the Conference Companion Publication on Designing Interactive Systems, DIS 2018, June 2018, Hong Kong, China*. ACM. pp.153-158.

Vane, O. (2017) Using Data Visualisation to tell Stories about Cultural Collections. In: *Proceedings of the Conference Extended Abstracts on Human Factors in Computing Systems, CHI 2017, May 2017, Denver, USA*. ACM. pp.335-339.

Boyd Davis, S., Vane, O. & Kräutli, F. (2016) Using Data Visualisation to tell Stories about Collections. In: *Proceedings of the Conference on Electronic Visualisation and the Arts, EVA 2016, July 2016, London, UK*. British Computer Society. pp.221-228.

Unpublished conference presentations

Boyd Davis, S. & Vane, O. (2018) *Designers — what are they good for (in data visualisation)?* [Presentation] Information+ Conference 2018, University of Applied Sciences Potsdam (FH Potsdam), Germany, 20th October 2018.

Vane, O. (2018) *Visualising the Cooper Hewitt collection*. [Presentation] Europeana Tech Conference 2018, Rotterdam, Netherlands, 15th May 2018.

Vane, O. (2017) *Temporal Visualisation for exploring Qualitative Themes across Historical Documents*. [Poster Presentation], ACM Conference on Human Factors in Computing Systems CHI 2017, Denver, USA, 8-9th May 2017.

Boyd Davis, S., Kräutli, F., Cottrell, S., & Vane, O. (2016) *Designing Histories of Design: Critical Chronographics*. [Presentation] Design History Society annual conference DHS 2016, Middlesex University, UK, 9th September 2016.

Exhibitions

Barbican Centre. (2019) *Curator Computer Creator: Museums and AI*. [Exhibition], London, 4th July 2019.

London Design Festival. (2018) *Design Research Evolution*. [Exhibition], Royal College of Art, London, 20-23rd September 2018.

Cooper Hewitt Smithsonian Design Museum (2018) *Saturated: The Allure and Science of Color*. [Exhibition], New York City. 11th May 2018-13th Jan 2019.

Acknowledgements

My biggest thanks go to Stephen Boyd Davis who guided me through this journey. Thank you so much for your knowledge, advice, enthusiasm, and for making this an enjoyable experience!

My thanks also for the support of the Wellcome Library, Cooper Hewitt Smithsonian Design Museum, Victoria & Albert Museum, National Swedish Heritage Board and Nordic Museum. Particular thanks to Pamela Horn, Kati Price, Richard Palmer, and David Haskiya. There are many others at these institutions who generously offered their assistance, time and knowledge contributing to this project - thank you.

At RCA, I am grateful to John Stevens for his feedback in my supervisions, Diana Tanase for her valuable technical assistance, Neil Parkinson for his help, expertise, and interest in the project, and my friends in the RCA School of Design PhD studio for their cheerfulness. Thank you to Florian Krätli and Sam Cottrell whose doctoral research this PhD builds on for their kind support and help.

Thank you to Siân Lindley for a stimulating and enjoyable summer at Microsoft Research Cambridge.

Thank you also to my family for their encouragement (and sachertorte!) and to Matthew who was there all along the way.

I gratefully acknowledge this PhD was supported by AHRC award AH/L503782/1 through the London Doctoral Design Centre and AHRC International Placement Scheme award AH/P002722/1.

Table of Contents

Chapter 1

Background	9
Cultural data	9
Visualising cultural data	11
Cataloguing & classifying	16
Time and timelines	19
Timeline visualisation of cultural data	22
Summary	24

Chapter 2

Cultural timeline visualisations—Survey	25
Highly curated narrative	25
Curated groups	26
Quantitative	27
Chronological grid	28
Mapping individual items by time	29
Plots	31
Item-to-item connections	32
Sampling	33
Summary	34

Research Questions	35
---------------------------------	----

Chapter 3

Methods	36
Institutions and datasets	36
Research through design in digital humanities	37
Cultural visualisation prototyping	37
Engaging collection experts and users	38
Evaluating cultural visualisations	39
Technical details	41

Visualisations	42
-----------------------------	----

Chapter 4

'Steptext'

Medical Officer of Health reports, Wellcome Library, London	43
Exploring historical texts with visualisation	45
Prototyping	46
Visualisation design	46
Exploring the MOH reports with Steptext	49
Feedback sessions	52
Discussion	53
Summary	57

Chapter 5

'Tag-timeline'

Cooper Hewitt Smithsonian Design Museum, New York City	58
Tagging in the Cooper Hewitt collection	58
Visualising image data by themes	60
Date information in the Cooper Hewitt collection data	60
Visualisation design and prototyping	60
Selecting tags, selecting objects	64

Scale issues	64
What do the tag-timelines show?	65
Feedback sessions	72
Discussion	73
Summary	77
 <i>Chapter 6</i>	
'Dive into Color'	
Cooper Hewitt Smithsonian Design Museum, New York City	78
Previous colour visualisations of cultural data	79
Colour theory and design	79
Colour in Cooper Hewitt's collection data	80
Visualisation prototyping	82
Final interface design	88
What does 'Dive into Color' reveal?	90
Feedback sessions	96
Discussion	97
Summary	99
 <i>Chapter 7</i>	
Exploratory visualisation	
Royal Photographic Society Collection at the V&A Museum, London	100
Prototyping	101
Image similarity across the RPS collection	107
'Time-tiers' visualisation	113
Summary	118
 <i>Chapter 8</i>	
'Faces of Sweden'	
Nordic Museum, Stockholm	119
Portraits in the Nordic Museum collection	119
Visualisation design	121
Discussion	124
Summary	128
 <i>Chapter 9</i>	
Discussion	129
Is 'just the data' enough?	129
Trust	132
Quantitative and qualitative visualisation	135
Completeness	135
How many interface designs do we need?	136
 <i>Chapter 10</i>	
Reflections	137
Designing cultural timelines	137
Evaluations	139
Limitations and future work	140
 Research contribution	142
Practical research contribution	142
Theoretical research contribution	142
 References	143
List of Figures	164

Background

The core of this research project is exploring new ways of interacting with cultural data. By 'cultural data' I mean the result of digitising cultural institution holdings. Cultural institutions include museums, galleries, archives and libraries. Their holdings are diverse, consisting of different types of objects and media ranging across artworks, artefacts, specimens, books, documents, ephemera, videos, sounds, and born-digital objects (see Figure 1). I use 'cultural items' as a general term for these.


Cultural institutions are increasingly digitising their holdings (Nauta & van den Heuvel, 2015:p.4). While the primary reason institutions digitise their holdings is to broaden and enhance access to the materials, there are also assumptions that once digital, these materials can be read, viewed, experienced etc. in new and creative ways (Hughes, 2004:pp.8-17). Essentially, there are things you can do with cultural collections as digital data that you cannot do with the physical items; one such possibility is data visualisation. But before discussing cultural data visualisation, it is helpful to characterise digitised collections.

Cultural data

Typically, when a collection is digitised, the associated institutional catalogue or index for the items is converted to digital data and it is now common for institutions to enter cataloguing information directly into a database. Catalogues perform a number of roles. They are finding aids for physically locating items, provide historical descriptive information necessary for understanding the material, flag related material in the collection, and record provenance (Riley, 2017:p.5). The catalogue can also serve an administrative function; institutions use these systems to track their acquisitions, exhibits and loans. The completeness and detail of catalogue records very much varies. The institutional catalogue both helps interpret a collection and describes the relationships and connections of items to each other.

The items themselves may be digitally reproduced: the objects and artworks may be digitally photographed, the texts may be transcribed to digital text files. Cataloguing information can be dynamic as new acquisitions are made and records are fortified or revised. Many institutions have an ongoing digitisation programme, describing their digital collection as a work in progress, constantly being improved and added to.

Figure 1. Physical cultural collections: (top) National Museum of Natural History, Smithsonian, Entomology Collection, photo Chip Clark (public domain); (middle) Brooklyn Museum open storage, photo Mark Schlemmer (CC BY 2.0); (bottom) historic newspapers at the British Library, photo Luke McKernan (CC BY-SA 2.0).





```
{
  "id": "18094513",
  "tms:id": "5398",
  "accession_number": "1896-1-12",
  "title": "Fragment, ca. 1700",
  "title_raw": null,
  "url": "https://collection.cooperhewitt.org/objects/18094513/",
  "has_no_known_copyright": "1",
  "department_id": "35347501",
  "period_id": null,
  "media_id": "35410699",
  "type_id": "35254623",
  "date": "ca. 1700",
  "year_start": null,
  "year_end": null,
  "year_acquired": "1896",
  "decade": null,
  "woe:country_id": null,
  "medium": "silk, metallic yarns",
  "markings": null,
  "signed": null,
  "inscribed": null,
  "provenance": "Jarves collection; pre-1896: Hewitt collection; 1896: acquired by the Museum.",
  "dimensions": "H x W: 27.7 x 26cm (10 7/8 x 10 1/4in.)",
  "dimensions_raw": {
    "height": [
      "27.70",
      "centimeters"
    ],
    "width": [
      "26.00",
      "centimeters"
    ]
  },
  "creditline": "Gift of Eleanor and Sarah Hewitt",
  "description": "Fragment of brocaded silk damask in the 'bizarre' style. Pink, red and silver ground with details brocaded in shades of blue showing a basket of flowers and fantastic shapes.",
  "justification": null,
  "gallery_text": null,
  "label_text": null,
  "videos": null,
  "on_display": null,
  "woe:country": null, ...
}
```

Figure 2. Example digital record: digital reproduction (left) and beginning of the associated metadata (right), in JSON format, from Cooper Hewitt collection. Fragment, ca. 1700. Public domain image. Data © Cooper Hewitt, Smithsonian Design Museum.

I use 'cultural data' to refer both to the data—digital reproductions of items—and associated metadata (or data about data)—the digital catalogue—though, in practice, the distinction between these two can be blurred. Metadata includes information about an item such as: its creator, creation date, origin, publisher, material, style, artistic movement, colour, acquisition year, donor, department, description, dimensions, page number, genre, and display or storage location (see Figure 2). From metadata, cultural items gain meaning and value.

Cultural institutions can use the term 'collection' to refer to the entirety of their holdings or to describe groups of items that share some characteristic within the whole, such as a subject or the way by which items were acquired (Society of American Archivists, 2019). New digital collections can be formed of material from different physical archives. I use the term 'collections' interchangeably for these cases.

Diverse character, content and structure

Just as institutional holdings are diverse, so are digitised collections. Cataloguing practices differ between institutions, leading to varying metadata structure and content. How a collection is digitised also varies. Sometimes a digital collection consists only of the cataloguing information with no digital reproductions of the items. And while there are standards and guidelines for digitising institutional holdings (for example Lourdi & Nikolaidou, 2009), technology developments over the last few decades mean existing digital collections are sometimes the product of outdated methods and approaches: documents previously scanned in black and white would commonly now be colour photographed in much higher resolution, or the recommended data format and structure of the time has since changed. It is a challenge for cultural institutions to continually update their digitisation processes and equipment, and it is common to encounter cultural data that is the product of legacy systems.

Besides, as Whitelaw (2015) argues, there is no intrinsic way to convert cultural collections to data. When digitising documents should the page images be edited to remove the background browned paper colour? Should document text transcribed to digital text files be segmented into different

data fields, and how? What exactly constitutes a single object? If digitising a dinner service, for example, should each plate get a separate digital record, or not (see Figure 3)? Digitisation does not prescribe an exact set of actions: “our encounter with a digital object consists of a particular representation or rendering of that data; but other representations are always possible” (Whitelaw, 2015b:p77).

Decades of digitising materials has resulted in large volumes of data, and levels of digitisation across the sector continue to grow (European Commission, 2019). The volume of cultural data available is, therefore, only set to increase. To add to this, aggregation platforms (such as Europeana (2019), or the Digital Public Library of America (2019)) and web technologies for exposing connections between items in different repositories (such as linked open data/semantic web initiatives (Hyvönen, 2012)) reveal a trend towards access to more. This creates challenges for making sense of, or managing, the data created, but also opportunities to interact with and ask questions of collections that would not have been previously possible. (Though, of course, just because the volume of available cultural data grows does not mean a complete record of the past is ever achievable). At the same time, digitisation projects are expensive to undertake and maintain: there is a growing emphasis that we need to better evidence the value of digitised collections (King, Stark & Cooke, 2016:p.77).

Is this ‘data’?

Using the term cultural ‘data’ to describe digitised collections draws on two senses of the word: that it is computationally processable, and that it can hold evidential value for particular claims or arguments (Oxford English Dictionary 2012, ‘data’; Owens, 2011). Many humanities scholars, though, who use collections to support producing scholarship, do not identify with describing digitised collections as ‘data’ (Marche, 2012). The word ‘data’ has scientific associations with objectivity (Drucker, 2011), and for some its use plays down that the creation of cultural items themselves, as well as the creation of collections and their digitisation, is always authored. Drucker (2011) has even suggested an alternative word to data—‘capta’—to make this emphasis. I use the term ‘data’ here, which I think is useful for dialogue across disciplines.

Visualising cultural data

Thinking of digitised collections as data opens up the possibilities of applying data methods to these resources, including visualisation. Data visualisation (or information visualisation—I use these terms interchangeably here) can be thought of as the use of “visual representations of abstract data to amplify cognition” (Card, Mackinlay & Shneiderman, 1999:p.6). Through creating visual mappings of data “a viewer can see patterns, trends or anomalies, constancy or variation, in ways that other forms—text and tables—do not allow” (Friendly, 2008). It has a long history in analogue form pre-dating computers. Now digital visualisations permit interactivity. (When referring to visualising cultural data, I am not concerned with 2D/3D renderings of heritage sites or objects which are also sometimes described as visualisations).

There is a growing body of research concerned with cultural data visualisation, and a growing number of visualisation examples (scholarly and popular). Windhager et al. (2018) reviewed

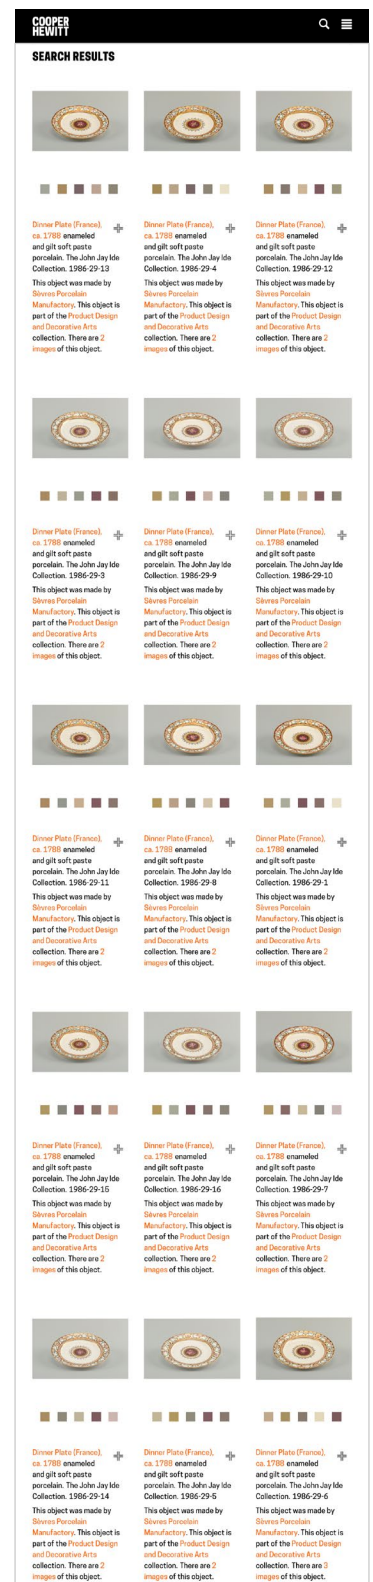


Figure 3. Cooper Hewitt museum collection website. Sèvres dinner plates set. Screenshot 16/08/19. © Cooper Hewitt, Smithsonian Design Museum.

visualisation approaches to digital cultural heritage collections, covering the wide range of visualisation types that have been employed: network diagrams, geographic maps, plots, grids, charts, timelines and also hybrids connecting multiple visualisations (see Figure 4 to Figure 7). Visualisation has been applied to datasets that represent entire collections, single items (for example, visualising the history of an item's owners and geographic locations (Cooper Hewitt Labs, 2013a)) and intermediate views between these extremes.

Digital humanities—an area of scholarly activity at the intersection of digital technologies/ computing and the humanities (though definitions remain contentious as the area is emergent (Terras, Nyhan & Vanhoutte, 2013))—has historically tended to focus on text. This skew has been reflected in visualisations created for scholarly purposes. Dörk, Pietsch and Credico (2017), however, describe a trend towards increased interest in visualising other data types, for example image, video or sound data. Only recently, Münster and Terras (2019) introduced the term 'visual digital humanities', covering research approaches that depend on using and/or making image, rather than textual, information of which, they note, we are seeing more.

What cultural data visualisations do

Data visualisations, generally, can be used to facilitate the exploration and analysis of data in order to arrive at new discoveries and/or to communicate existing insights, knowledge, or arguments (Munzner, 2014:Chap.3). In addition to supporting analytic insight, visualisations can also be used for storytelling (Segel & Heer, 2010; Kosara & Mackinlay, 2013; Lee et al., 2015), 'casual' visualisations encourage awareness of an issue (Pousman, Stasko & Mateas, 2007), and artistic examples use aesthetic means to express a point of view (Viégas & Wattenberg, 2007).

To ask 'what do/can cultural data visualisations do?', though, is also to ask 'for who?'. Visualisations in this domain have been created to serve scholars working with collections, collection managers at institutions, and also for more casual audiences (both in-gallery examples and online). They can help support research activities, collection administration and management, and casual browsing and display. Below, I walk through what visualisation can do in these different cases (though in practice there is overlap between audiences and activities, and an individual visualisation design can support multiple modes).

Analysis

Visualising cultural data can support analysis of collections by showing patterns that result from sorting, grouping, filtering and other visual rearrangements of the data to expose outliers, trends and abrupt changes (Sinclair, Ruecker, & Radzikowska, 2013). This is particularly productive for datasets too large to grasp or process otherwise and relates to Moretti's notion of 'distant reading', which I will return to. Visualisation can therefore help researchers ask questions of collections that would have been too time-consuming otherwise to explore.

Analysing cultural data through visualisation can also assist in the administration/management of collections, "allowing curators and collection managers to gain intellectual or physical control of a collection" (Bailey & Pregill, 2014:p.185). Cultural data visualisations are, for instance, good at highlighting errors in the collection data: "outliers, inconsistencies, or inaccuracies within data...can be indicative of errors in cataloging, uncontrolled taxonomies, or other metadata issues requiring correction" (Bailey & Pregill, 2014:pp.179-180). Outside of highlighting errors, visualisation can offer a window on an institution's metadata practice (Boyd Davis and Kräutli, 2015).

But pinning down the exact process by which knowledge is produced using a cultural visualisation is difficult, and an issue I discuss in more detail in 'Methods'. Visualising collection data does not necessarily provide answers: it can raise more questions. Bailey and Pregill (2014) describe the

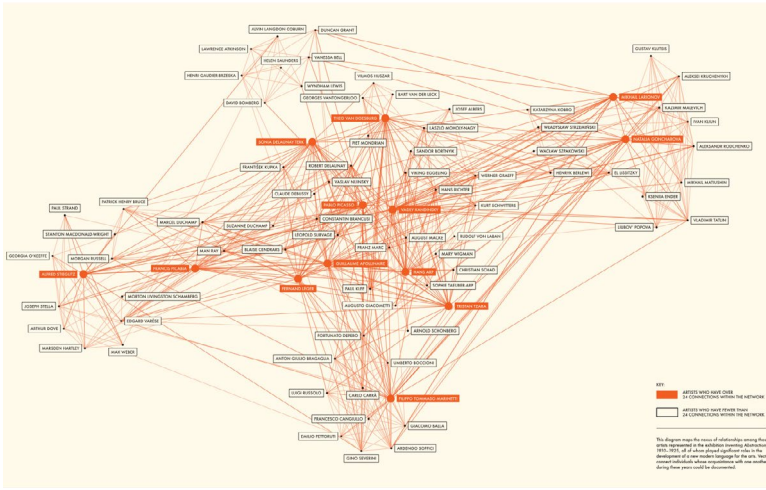


Figure 4. MOMA. 2012. 'Inventing Abstraction'. Content produced and provided by Second Story.




Figure 5. Mitchell Whitelaw. 2009. 'The Visible Archive'.

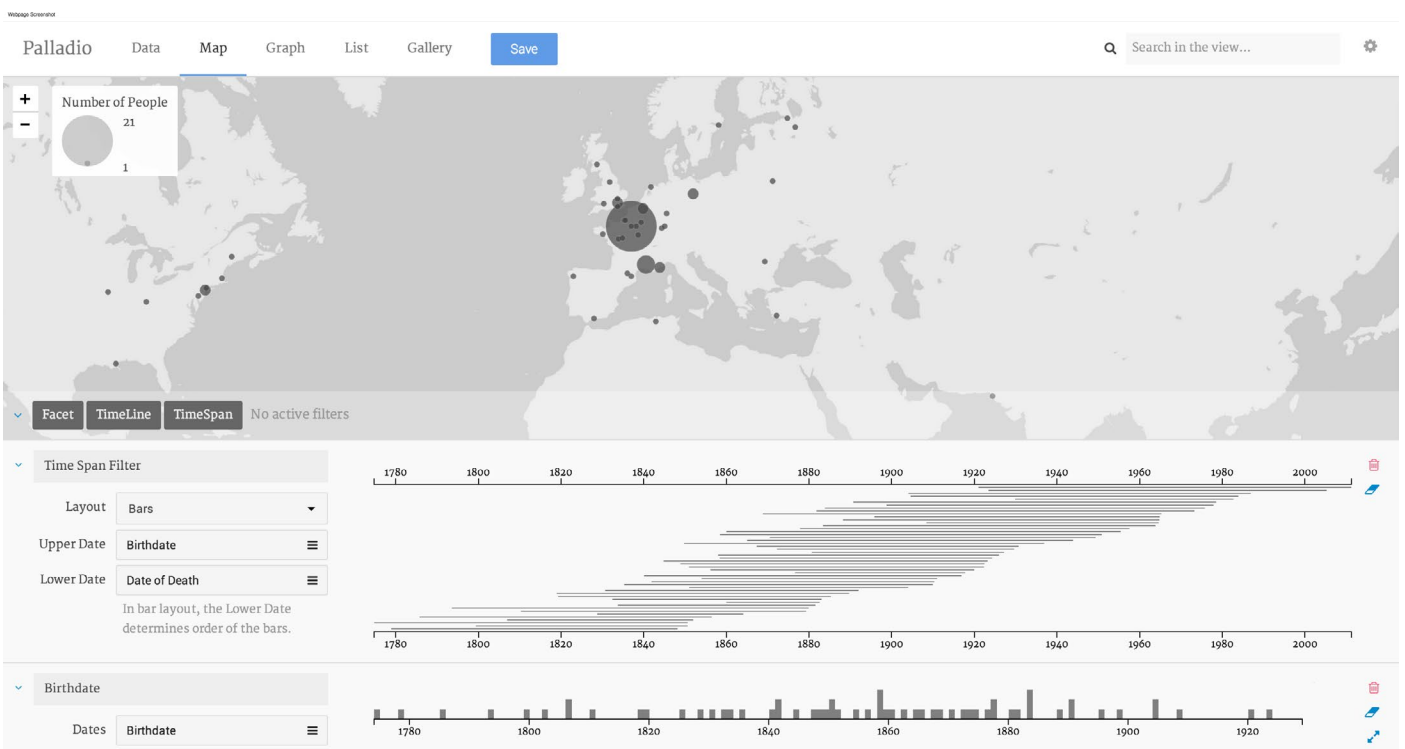


Figure 6. (above) Giovanna Ceserani, Nicole Coleman, Mark Braude, Giorgio Caviglia, Ethan Jewett. 2014. Palladio.

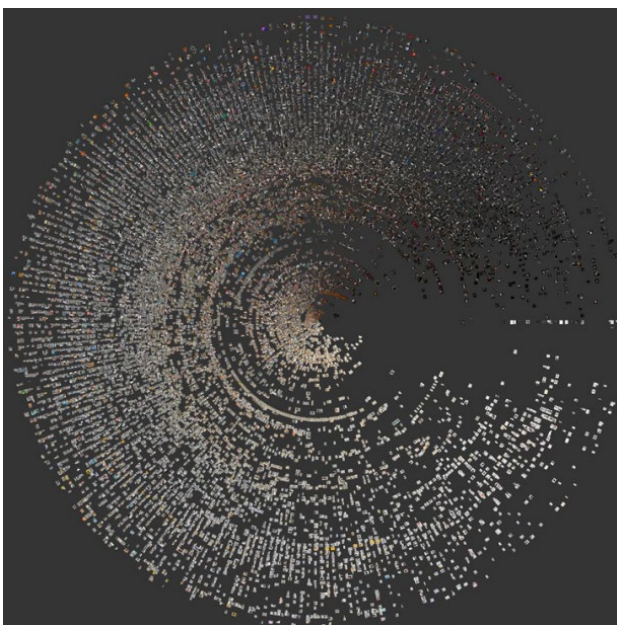


Figure 7. (left) Nadav Hochman and Lev Manovich. 2014. Visualization of the Thomas Walther Collection.

example of The Samuel H. Kress History and Conservation Database Project where visualising the collection data “facilitated a better understanding of the collection’s origins, qualities, and donation patterns” (p.185). Rather than the visualisation answering pre-existing research questions, however, trends revealed by the visualisation “served as wayfinders for further exploration” (p.186). Visualisation can reveal trends in the data without giving away their origins. In the case of the Samuel H. Kress project, visualising the provenance information of this collection led to “a suggestion, a curiosity, that [might not have emerged]... from study of the collection data itself” (Bailey & Pregill, 2014:p.186).

Further, Johanna Drucker’s (2009; 2010) work highlights the interpretive and sensory role of the visual in producing knowledge with cultural visualisation. As Drucker (2010) argues, “knowledge [in the humanities produced using visuals] is not transferred, revealed, or perceived, but is created through a dynamic process” (p.36).

Exploration

Analysis and interpretation are only part of the scholarly process and, for researchers working with cultural data, an equally important task is that of locating and assessing new material (Sinclair, Ruecker, & Radzikowska, 2013). Cultural visualisation can also help to support exploration of collection data for these kinds of activities, providing a unifying and structuring effect to surveying material, and enhancing contextual detail in collections (Bailey & Pregill, 2014). Interactivity has special potential to offer multiple views on cultural data (Sinclair, Ruecker, & Radzikowska, 2013).

There is also increased interest in designing interfaces for cultural collections which offer a richer experience for casual browsing, for instance Whitelaw’s ‘generous’ (2015) and Ruecker, Radzikowska and Sinclair’s ‘rich prospect browsing’ (2011) interfaces. Promoted as an alternative to the inadequacies of search-only access, and as engaging for non-expert users, visualisation can reveal the scale and complexity of cultural data. Visualisation can be a way to support serendipity in information seeking across cultural data (Thudt, Hinrichs & Carpendale, 2012), and can provide visually appealing and playful views on collection data.

Presentation

Visualisation of cultural data can also be used for presentation and display; there are both static and interactive examples. It can be used to communicate an argument or point of view, but also curated paths and sequences can be employed to present narratives about a collection (Boyd Davis, Vane & Kräutli 2016). Visualisation design can also be a way to communicate a statement about the collection as a whole, what Whitelaw (2015b:p.91) calls a “curatorial premise”. I encounter this function later in my practical work. Whitelaw gives the example of a visualisation from MoMA (2012)—see Figure 4—mapping documented relationships between artists featured in the museum’s ‘Inventing Abstraction’ exhibition. While the visualisation supports exploring these connections, as a whole it also embodies and promotes the thesis of the exhibition: “abstraction was not the inspiration of a solitary genius but the product of network thinking” (MOMA, 2012).

Design values for cultural visualisation

Close, distant and intermediate views

A distinction has been made between cultural visualisations that enable ‘close’ or ‘distant’ readings of cultural data, and those that combine approaches (Jänicke et al., 2015). Essentially, to what degree is the data abstracted to achieve an overview? (There are similarities with the contrast made in data visualisation/interface design between focused and contextual views of a dataset (Cockburn, Karlson & Bederson, 2008)).

The ‘distant reading’ concept derives from Franco Moretti (2005). Distant reading is not reading in the normal sense of reading a text, but taking an abstracted view on a text/s to examine global patterns, by applying statistical analysis and/or data visualisation (for example, see Figure 8). Moretti’s idea is that, by taking this approach, it is possible to study a much larger volume of texts than time would allow an individual to actually read (‘close reading’) and get different kinds of insights. Thus a far more totalising view is possible. (‘Close reading’ can mean a specific way of engaging with texts deriving from literary and cultural studies (Culler, 2010:p.20), though in the context here the term is often used loosely). Parallel ideas have been applied to visual sources and artefacts, for example ‘distant viewing’ (Arnold & Tilton, 2019) is now a term in use.

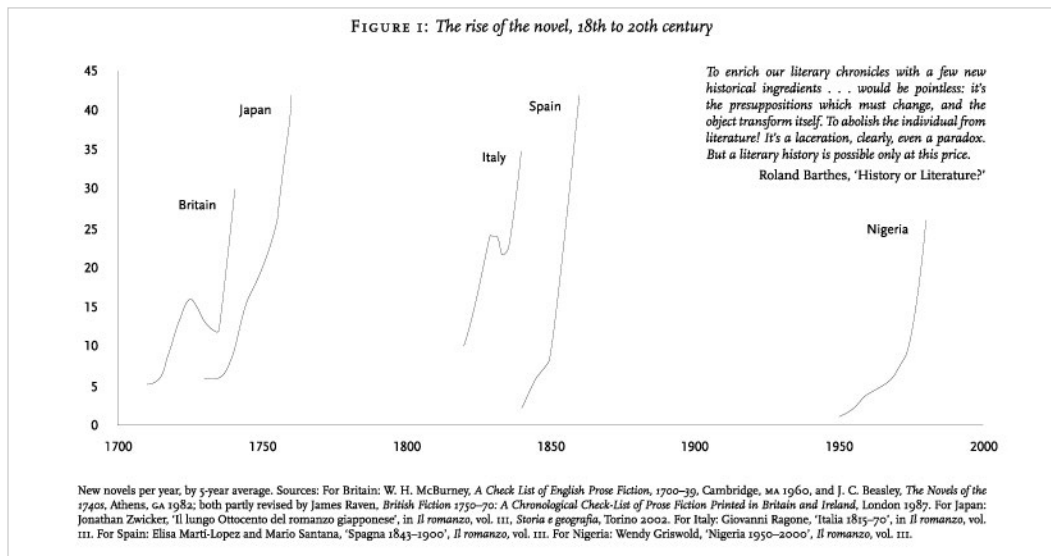


Figure 8. Franco Moretti. 2005. ‘The rise of the novel, 18th to 20th century’. From ‘Graphs, Maps, Trees: Abstract Models for Literary History’ (p.6).

Distant reading, however, has been a very controversial idea within the humanities (Marche, 2012). Quantitative forms of analysis are not a surrogate but a fundamentally different way of engaging with these items, which has not been welcomed by all. Coles, a poet and poetry scholar, articulates some of these concerns going into a collaborative visualisation project:

Because much Digital Humanities research at the time was driven by technology and the desire to extract data from or quantify texts or their features rather than to engage the kind of qualitative, aesthetic experience we considered central to our pursuits, we worried that the project would require us to subordinate our deepest values to accommodate ourselves to what the machine could already do (Coles, 2017:p.1).

There is a fear of being pushed into doing what, for the computer, is easier but not, by some views, necessarily more insightful. Winters (2017a) similarly cautions about the limits of the “decontextualised macro-level view” where complexity—one of the hallmarks of humanities data and, as Hitchcock (2014) argues, a source of the rich meaning that can be achieved through its close analysis—“is masked by a smooth curve on a graph”.

Quantitative, qualitative

Many commentators assume that data visualisation means dealing primarily with quantities and numbers: “what is humanistic about data visualisation?...a body of methods for exploring and explicating quantitative datasets” (Graham, 2017). Certainly digital humanities’ origins were quantitative: “the first wave of digital humanities work was quantitative, mobilizing the search and retrieval powers of the database, automating corpus linguistics” (Schnapp et al., 2009:p.2). In the readiness to deal with quantities, rather than other aspects, do we perhaps fail to stretch our representations to deal with more difficult, but potentially more useful things? There is also an

analogy here with the tendency for quantitative visualisation to have a much bigger literature, and more scholarly debate, than the visualisation of qualitative or nominal data.

As the field has developed, however, (interconnected with developments in computing technology), increasingly projects are “qualitative, interpretive, experiential, emotive, generative in character... [harnessing] digital toolkits in the service of the Humanities’ core methodological strengths: attention to complexity, medium specificity, historical context, analytical depth, critique and interpretation” (Schnapp et al., 2009:p.2). My cultural visualisation projects described in this thesis belong in this space. While quantitative visualisation of cultural data is productive for certain aims, there are many examples of it and it does not push the boundaries of what computational methods can do.

Cataloguing & classifying

Pertinently, the kind of data found in digitised collections, while it can be treated quantitatively, is not inherently so. It is largely nominal (as in, assigning names). As defined by Meirelles (2013:p.187), nominal data is distinguished on the basis of quality; examples include objects, names and concepts. Categorisation plays a major role in manipulating nominal data, and shared characteristics allow grouping.

How data is structured and what is recorded informs what can be done with it. Visualising data in a particular way is dependent on there being suitable attributes in the data, and the data being in a suitable form. What do cultural datasets allow or afford? Factors include: how are the objects to be catalogued divided into records (is a sketchbook assigned a single object, or every sketch within it, or both)? What is the nature of the classification system? What sorts of categories have been attended to? How is the metadata produced—manually? automated? How flexible is the dataset to change/augmentation? To be clear, cataloguing is the general process of creating metadata representing cultural items. Classification meanwhile, which may be part of the cataloguing process, involves assigning that item to a set class in a classification system.

There are two important points here: that cultural datasets are not necessarily a definitive representation, and they are not set in stone. They can be changed, reworked, enriched—and the desire to use a particular visualisation template can be the impetus for changing cataloguing approaches. (Data can also deteriorate. For example, moving electronic records from one institution to another can introduce errors because of incompatible data systems).

Classification across cultural institutions

Different types of cultural institutions have traditionally had different attitudes towards, and practices of, classification and cataloguing, leading to datasets with particular characteristics. Robinson (2014) identifies fundamental differences between libraries, archives and museums. Generally, libraries aim to provide broad access to their collection using a standardised categorisation scheme with rigid hierarchical subject definitions. These subject definitions form a “sort of epistemic cartography - mapping knowledge”(Olson, 2001:p.652). Dewey Decimal Classification is an example of such a scheme. (Although discrete digitised collections from libraries, rather than a broad view of their holdings, are often available as standalone datasets with flat data structure).

In contrast, for archives, the most important concern is “retaining the relationship between the documents and the institutional functions and activities that gave rise to them” (Robinson, 2014:p.418). The documents must be kept in the order in which they were received: the principle of provenance. This has an implication for data structure, in that archives often group documents in series (groups of similar items), rather than just indexing and describing items at item level. The

digital records for the Babbage Papers archive (relating to computing pioneer Charles Babbage), for example, follow this pattern; the data structure is a hierarchy of grouped documents (see Figure 9).

Further, good archival practice demands staying at arms length from the interpretation process. As Robinson (2014) writes: “the interpretation of collection holdings in historical or thematic contexts by archivists is actively discouraged and even regarded as antithetical to good archival practice...[thus archivists generally avoid the use of] subject or theme indexes and other finding aids common to museum collections and libraries” (p.418).

In contrast, museums are much more in the practice of explicitly providing context and interpretation to collection items in cataloguing. Lester (2001) argues that museums see their essential role as interpreters of collections. Consequently, in addition to standard descriptive criteria—such as manufacture date, maker, and physical measurements—museum datasets are more likely to include contextual or interpretative data, for instance sitting items in broader narrative and thematic groupings. You can find this put explicitly in software documentation. Describing the functionality of a particular data system for museum collection records: “through [this software function]...name authorities and places can be linked to objects, creating history and context” (Collections Trust, 2019). As Robinson (2014) argues: “in this process, museums are not about unmediated open access to collections” (p.422).

The dominant mode for organising digital metadata in cultural institution databases is by a faceted classification system, and this is largely what I encountered in practical work. Faceted classification systems classify each item along multiple explicit dimensions called facets (Tunkelang, 2009) which allow items to be accessed and ordered in multiple ways. In the case of cultural data, the metadata fields (for example ‘creator’, ‘creation date’, ‘material’) are the facets. Thus faceted search, the information retrieval method for accessing information organised according to this system, both reveals navigable features of the collection and allows the user to progressively narrow their choices by different dimensions (facets) (Whitelaw, 2015b).

Recently, there has been much interest in introducing linked data structures (Hyvönen, 2012) to cultural data. This is an alternative way of structuring data, whereby the relationships between entities are made explicit and items from different datasets can, in theory, be seamlessly connected by shared characteristics. This would have other implications for visualisation, but is not the focus of this PhD.

What is recorded?

What is recorded about a cultural item “is not self-evident or singular” but determined by the identity and aims of the institution (Cameron & Robinson, 2007:p.171). As a curator I interviewed explained: “if you don’t work in a museum I think there’s a common misconception that these collections are definitive and they just represent [,for example,] the history of photography and everything you might need to know about it. But they’re incredibly subjective and they’ve gone in very different directions in the past” (V&A curator, 2019).

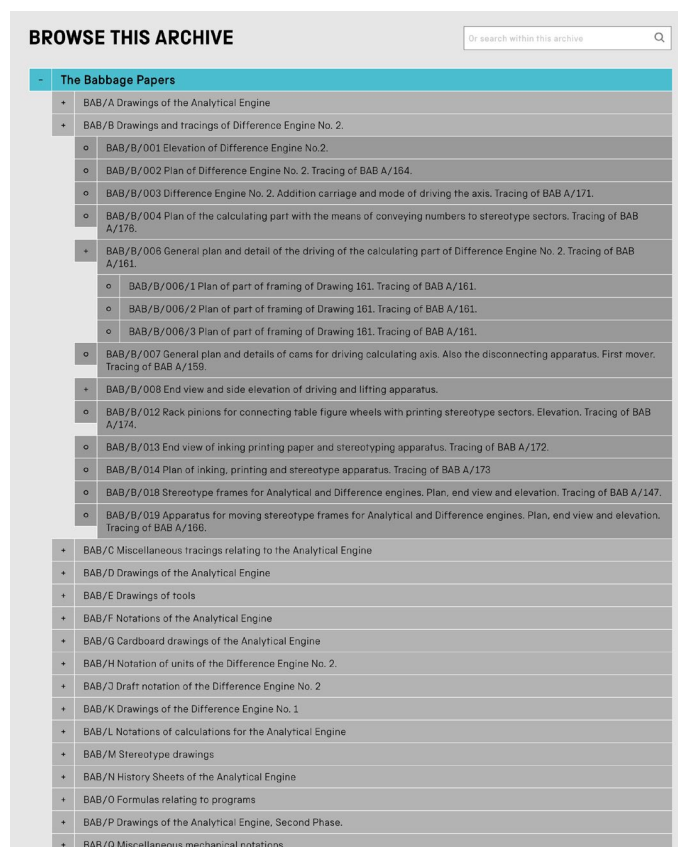


Figure 9. Science Museum Group Collection website. 2019. The Babbage Papers archive.

Classification schemes are not just culturally contingent (Cherry & Mukunda, 2015), but also institutionally contingent. Hooper Greenhill (1992:p.7), for example, identified the potential for a silver teaspoon to be classified as 'Industrial Art' in Birmingham City Museum, 'Decorative Art' at Stoke-on-Trent, 'Silver' at the Victoria and Albert Museum, and 'Industry' at Kelham Island Museum in Sheffield.

Additionally, cataloguing data is subject to change. What is recorded about items intends to meet the requirements of those accessing the collection. But those requirements can change, and cataloguers have to try to keep up. The same curator continued:

20 years ago they'd [researchers] come in and say, 'I want to see works by this artist', 'I want to see works by Constable'... Now, they're more likely to come in and say, 'I'm researching early feminism' or 'I'm researching black British history'. It's much more thematic and much more social history based perhaps. So they're looking for specific imagery that we didn't necessarily think to record at the time. We were more focused on who painted it or who took the photo (V&A curator, 2019).

Further, Cameron and Robinson (2007) raise the ways publishing collection data on the internet and a contemporary social inclusion agenda aimed at cultural institutions (pp.166-167, referencing Wallace, 2001) has impacted collection documentation. Cultural datasets published on the internet are often not identical to those on internal systems and cultural institutions' web presence is now its own research genre (for example, *Museums and the Web*, 2019). Initiatives now allow public contributions to cultural data, for example crowdsourcing (Oomen & Aroyo, 2011).

How is it recorded? Manual and automatic metadata

Beyond manually cataloguing items, it is also possible to use computational methods to automatically generate new descriptive metadata, for example from the digital reproductions of cultural items themselves. Whitelaw describes this process as 'automatic content-based metadata' and gives the example of computationally extracting a colour palette from page images of a highly illustrated newspaper collection (Whitelaw, 2015a). Other methods for automatic content-based metadata include natural language processing techniques to extract topics or sentiment analysis of texts (Hai-Jew, 2017), or by applying computer vision technologies to classify images (Pim, 2018). In these cases computation is essential; it would be prohibitively difficult or time-consuming to create these kinds of data otherwise.

Cultural datasets are changing: by who and how they are created, what they contain and can support. But are some kinds of cultural data more legitimate or important than others? In this quote filmmaker Wes Anderson reflects on the experience of co-curating an exhibition at the Kunsthistorisches Museum, Vienna. In this exhibition, Anderson and his co-curator Juman Malouf grouped objects unconventionally—by color, material and size—rather than provenance or rarity:

...one of the Kunsthistorisches Museum's most senior curators...at first failed to detect some of the, we thought, more blatant connections; and, even after we pointed most of them out, still questions their curatorial validity in, arguably, all instances (Anderson, 2018).

Time and timelines

Amongst all the possible ways of organising cultural data, time is unusual for being almost universally applicable and is the focus of this PhD. Time is certainly not the only productive way to organise cultural data, but the temporal dimension is very important for identifying meaningful connections between items in cultural collections. User feedback for collections visualisation has stressed the primacy of time (Dörk, Pietsch & Credico, 2017). In addition, while spatial visualisation such as cartography has enjoyed significant innovations in digital form, design of digital temporal visualisation has been relatively neglected (Boyd Davis, 2012).

What is time? In a general sense, time is indefinite, continuous duration (Oxford English Dictionary 2003, “time” entry). Different disciplines characterise the physical dimension of time in different ways, but the common thrust is that “time is unidirectional (arrow of time) and that time gives order to events” (Aigner et al., 2011:p.45). The standard model of time is linear, with a hierarchical structure of granularities—days, weeks, months, years, decades, centuries, millennia etc. Granularities combine to form calendar systems and cyclicity is common in time data, for example weekdays or seasons. Sometimes non-linear or uneven models of time are useful; there has been much research in Human Computer Interaction concerning an experiential view on time (Lindley et al. 2013). In the historical context, an exponential model of time may suit to foreground the recent past (Schmidt, 2018:p.158-9). A numerical model of time, as an absolute, uniform frame of reference, was notably championed by Isaac Newton. This concept was a crucial precondition for mapping events and objects to time in visualisation (Boyd Davis, Bevan & Kudikov, 2010).

Challenges around time in cultural data

As generally, there is great heterogeneity in the ways date information is recorded across and within collection data. Loukissas (2017:p.654) recorded finding at least 1,719 unique templates for date information in New York Public Library collection data. Examples included: “_ _ March, _ _ _ _”, “probably before _ _ _ _” and “[c_ _ _ _]/ _ _ _ _”. In my practical work with different datasets I encountered a variety of different templates and data formats across and within datasets.

Moreover, there are larger issues for visualisation around how time is understood in history, how it is recorded in cataloguing of cultural collections, and how this squares with the constraints and requirements for data encoding and visualisation. I encounter some of these issues later in the PhD practical work (and a number are illustrated with examples found during the practical work).

Periodisation

When describing the past there is a tendency to chop time up into discrete blocks. Categorising the past into discrete, quantified, named blocks of time is called periodisation (Rabinowitz, 2014). Different systems apply for different fields of study eg. geological, anthropological and historical.

Different periods can be synchronous or overlapping. Figure 10, for example, shows an historical coins collection mapped time against time (date against historical period). And periods are not just temporal entities, but can carry other meanings too, for example geographic (eg. Ming Dynasty). Particular decades and centuries have a particular character in popular memory—think of the 60s. Schmidt (2018:p.156) gives the example of USA categorisation of architectural styles. ‘Second Empire’ and ‘Victorian’ relate to time periods and could describe buildings from the same year. But the first implies a French-influenced architectural style, while the second British. To convert between periods and numerical time spans, PeriodO (Rabinowitz & Shaw, 2019) is a gazetteer of (historical, art-historical, and archaeological) period definitions.

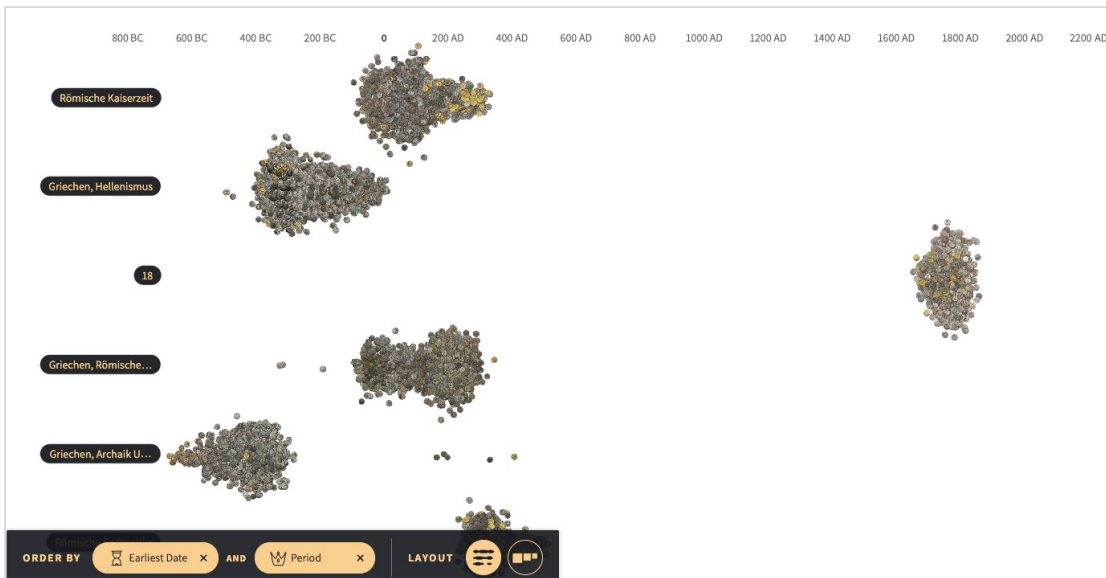


Figure 10. Flavio Gortana, Franziska von Tenspolde and Daniela Guhlmann. 2017. 'Coins – a journey through a rich cultural collection', showing data mapped by date against historical period.

Boyd Davis and Kräutli (2015) used visualisation to uncover different dating tendencies in institutional catalogues: certain institutions prefer dating items to the year, decade or 50 year period. This tendency has tangible impacts for working with digitised collections where the most dramatic changes will be visible at boundaries.

Uncertainty

Uncertainty is a property of data that refers to imperfect or unknown information. Uncertain data is endemic across cultural datasets and is a particularly common characteristic of date descriptions. Kräutli and Boyd Davis (2013) and Schmidt (2018:p.162) describe different varieties of date uncertainty in cultural data, including types of inherent imprecision in knowledge of dates, but also that it may be introduced by the constraints of how a data system allows input, and how it translates/stores that input. Uncertainty can be expressed by a low resolution of date information (for example, '18th-century'), but also by attaching ambiguous terms such as: 'about' or 'approximately'. Often date uncertainty recorded as free text is converted to a possible numerical date span following some formula.

If not explicit, date information for cultural items may be estimated with expertise. As a curator I interviewed described, concerning a prints collection,:

A lot of stuff is not dated and it's not signed. And we look at them and we try to make an educated guess about the dating...quite often it's due to the paper quality. But, corroborative information is often what colors are they using, what is the general sort of aesthetic of it (Curator, 2018).

Though this raises to what extent visualisation can challenge established verities if data creation practices render them self-fulfilling?

The Extended Date/Time Format (EDTF) (Library of Congress, 2019) is a recent data format development that aims for greater expressiveness in recording date uncertainty, but is currently rare to encounter. In any case, it still does not prescribe how to visually represent imprecise or uncertain data. Various cultural visualisation projects have explored visual representation of uncertain date information (Cottrell, 2017; Grossner & Meeks, 2013; Kräutli & Boyd Davis, 2013; Boyd Davis, Bevan & Kudikov, 2010; Drucker & Nowvskie, 2003).

Individual item histories

Records in cultural data may have multiple dates attached. These can refer to different episodes in an item's history: a production date, date it was accessioned by the institution, date it was catalogued, dates it featured in exhibitions or was loaned etc. In Swedish Open Cultural Heritage data (K-samsök, 2019a), for example, date information relating to different episodes in an item's life is recorded within a single date facet, with generic (eg. 'create') and specific (eg. 'design') categories for different episode types.

Even just focussing on an item's production date, there can be multiple dates attached. A photograph, for example, can have a different date for when the photograph was taken, and when it was printed (see Figure 11). Items can be reproductions of earlier items (see Figure 12). Historical buildings may have many dates attached capturing rebuilding, extensions, remodelling etc. (Figure 13 shows the record for a church with 46 dates attached). If mapping data by production date, which one is appropriate? These kind of problems likely need to be dealt with on a case by case basis, determined by the goals of the visualisation.



Figure 11. The Royal Photographic Society Collection. Photograph. Date: 1860s (photographed) mid 20th century (printed). © Victoria and Albert Museum, London.



Figure 12. Mariner's Astrolabe (replica) (USA), original: 1602; replica: 1963. © Cooper Hewitt, Smithsonian Design Museum.

Figure 13. (right) Bebyggelseregistret (Swedish Building Register). Rölunda church. Record has 46 dates attached for episodes in the building's history. Screenshot shows beginning of this date information in the record.

År 1100 - 1529 Nybyggnad	Åtgärd	Långhus, betydande partier bevarade i väster
År 1100 - 1529 Nybyggnad		
År 1650 - 1699 Nybyggnad - Korparti		
År 1730 - 1730 Fast inredning - altarpuppats	Upphovsman	Nils Falk (Konstnär - Bildhuggare)
År 1797 - 1797 Underhåll - takomläggning	Åtgärd	Taktäckningen ändras från spån till skiffer på underlag av nä
År 1873 - 1873 Ändring - ombyggnad, fönster	Åtgärd	Nuvarande större fönsteröppningarna tas upp.
År 1873 - 1873 Fast inredning - predikstol	Åtgärd	Flyttad från södra sidan till norra. Kraftigt om- eller nybyggd. Enligt dåtida beskrivning "En för dessa trakter vanlig 1600-talskaraktär".
År 1873 - 1873 Ändring - ombyggnad, interiör	Åtgärd	Altare med altarpuppats placeras mot östra korväggen. Tidigare hade där funnits ett sakristieutrymme.
År 1873 - 1873 Fast inredning - altare	Åtgärd	Nytt altarbord
År 1873 - 1873 Fast inredning - altarring	Åtgärd	Ny altarring
År 1873 - 1873 Fast inredning - läktare	Åtgärd	Ny dubbelt så stor läktare
År 1873 - 1873 Fast inredning - bänkinredning	Åtgärd	Troligen nya bänkar
År 1873 - 1873 Ändring - restaurering	Åtgärd	Större renovering, främst interiört.
År 1873 - 1873 Ändring - ombyggnad, port	Åtgärd	Sydingången sätts igen och litet vapenhus framför rivs. Två nya ingångar upptas. Korportal i söder och västingång.
År 1873 - 1873 Underhåll - takomläggning	Åtgärd	Skiffertaket reparerat
År 1928 - 1928 Äldre kulturhistorisk inventering	Åtgärd	Kulturhistorisk inventering av lösa föremål.
	Upphovsman	Bertil Waldén (Fotograf)
	Referens	ATA (Arkiv)
År 1935 - 1935 Nybyggnad - Sakristia	Åtgärd	Ritningar och program av arkitekt Berglund 1933. Arbetsledning på plats arkitekt Brattberg.
	Upphovsman	Allan Berglund (Arkitekt) Arthur Brattberg (Arkitekt)
År 1935 - 1935 Arkitekturbunden utsmyckning - måleri, interiör	Åtgärd	Vegetativt dekorerad taklist.
	Upphovsman	Anders Berglund (Arkitekt)
År 1935 - 1935 Underhåll - målningsarbete, interiör	Åtgärd	Bl a marmoreringar på altare, bänkgavlar och kolonner.
	Upphovsman	Olle Hellström (Konservator)
År 1935 - 1935 Ändring - restaurering, exteriör	Åtgärd	Omfattande renovering och tillbyggnad av sakristia mot norr. Ritningar och program av arkitekt Berglund 1933. Arbetsledn på plats arkitekt Brattberg.

Depiction

The Nordic Museum, Stockholm uses a data class—'tidsanda', which crudely translates as 'time spirit'. Here, 'tidsanda' (which is a sub-category under the facet 'motif') is used to record the historical time an artwork or design depicts, if this is at odds with its production date. For example, the print shown in Figure 14 has the production date '1828-1835', but depicts an historical figure from an earlier time: its 'time spirit' is recorded as first half of the 18th Century. Similarly, Figure 15 shows a V&A museum record where the item's 'style' is temporal—'20th Century'—a wider

time span than its manufacture date. Wikidata, the structured data underpinning Wikipedia and associated projects, which has recently drawn interest from cultural collections, similarly now offers a 'date depicted' data class (Wikidata, 2019). These examples evoke the creative potential for visualising date information in, for example, the history of art and design, where movements and styles often revive or reference the past (think of Pre-Raphaelite art or Gothic Revival architecture).




Figure 14. (left) Nordic Museum. Print. Production date '1828-1835', and 'tidsanda' (time spirit): '1700-talets första hälft' (first half of the 18th Century). (CC BY-NC-ND 4.0.)

Figure 15. (right) Jewellery design. 'production date': '1908-1917', 'style': '20th century'. © Victoria and Albert Museum, London.

Timeline visualisation of cultural data

Visualising cultural data temporally lends itself to a timeline design: a visualisation form mapping a sequence of events by time. (As opposed to visualising time series data, which is a sequence of data points taken at successive equally spaced points in time).

Chronology, history and visualisation

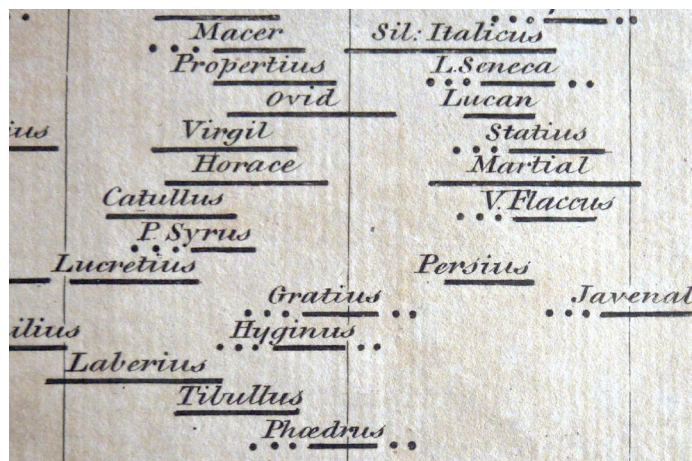
What does visualising cultural items by time do? How does it relate to chronology and history? Where does historiography factor in this? Ordering collection data by time is employing chronology to derive meaning. History, in contrast, has connotations of narrative and storytelling (Boyd Davis, Bevan & Kudikov, 2010:p.2). To shed light on these questions, Boyd Davis, Vane and Kräutli (2016) recommend an historical perspective. Prior to the mid-18th Century, there had been two traditions in the visual representation of historical time (comprehensively documented by Rosenberg and Grafton (2010)). The first is typographic, where events are ordered in lists or tables. The second used symbolic shapes like chains, trees, rivers (eg. Figure 16), the body, animals etc. to present history as a graphical story, guiding the viewer what to think. The mid-18th Century sees the emergence of a third, mechanical tradition, exemplified by Joseph Priestley's (1733–1804) 'Chart of Biography' 1765, see Figure 17. In the designs that emerge in this new tradition, events were now procedurally mapped in space often with a uniform time scale. Priestley, suspicious of rhetoric, argued that this stripped-back approach allowed viewers to gain an understanding of the information almost automatically. By observing the resulting visualisation, viewers can see ordering, what was simultaneous, intervals between events, and an overall sense of scale: the data is left to 'speak for itself'.

The tradition of using symbolic graphics, however, did not die away with the emergence of this new mechanical approach: "[their] very neutrality...seemed to some...to diminish the ability to 'tell a story'" (Boyd Davis, Vane & Kräutli, 2016:p.221). Debates around the value and ethics of rhetoric in visualisation, for example Tufte's (1983:pp.106-121) rejection of "chartjunk", still show no sign of abating today (Boyd Davis, 2017:p.4; Bateman et al., 2010).



Figure 16. (left) Girolamo Andrea Martignoni. 1718. 'Imago Romani Imperii'. Detail. Collection: Institut Cartogràfic i Geològic de Catalunya. (Public domain).

Figure 17. (right) Joseph Priestley. 1765. 'Chart of Biography'. Detail. Collection: Chetham's Library, Manchester. Photo: Stephen Boyd Davis.



Boyd Davis, Vane and Kräutli (2016) further question the relationship between historiography and cultural data visualisation: “[What of the] contentious and disputed nature of history? For a museum collection or an archive, what multiple, competing and conflicting narratives can be discovered – or created – by historians, researchers, archivists or non-experts? And what should such visualisations look like?” (p.223).

These early designs illustrate contrasting approaches to presenting historical time through visualisation: overtly presenting a story using handcrafted rhetorical graphics, or procedurally mapping the data and leaving it to ‘speak for itself’. Storytelling with visualisation (‘narrative visualisation’) has previously been explored building on journalistic techniques (Segel & Heer, 2010; Hullman & Diakopoulos, 2011). In this research, though, I am also interested in more subtle forms of visual narrative. In a simple sense, narrative can be thought of as an account of a collection of items given in order and with the establishing of connections between them (Oxford English Dictionary 2003, “narrative” entry). I find a useful analogy here in how exhibitions can function as visual narratives (this is especially relevant as the projects in this PhD predominantly visualise digitised museum collections). Literature in museology stresses how the selection and spatial arrangement of exhibits around chosen principles creates its own forms of implicit meaning-making through juxtaposition, proximity, sequence and succession (Albano, 2014:pp.535-536). This idea of visual narrative is more subtle and implicit than overt storytelling, and perhaps is closer to how meaning may be formed in the ‘only the data’ visualisation approach.

Context and completeness

Is there enough context in the data alone? An historical record may have a production date attached, but without other information we are in the dark about whether that item was typical of the time, or if it was pioneering. Was the item everyday for the period, or the finest example? And wider historical context can have implications for visualisation design. In an interview I conducted with a museum curator discussing visualisation, they recommended splitting the artwork dataset under discussion geographically; the artwork from one region, the curator described, had been influenced by that from another, but with a temporal delay. Otherwise, “the timeline will look a bit strange” (Nordic Museum curator, 2018). Context can be important both in the design and the interpretation of cultural visualisations.

Another important issue is that of completeness. One of the criticisms leveled at applying data methods in the humanities is that cultural datasets are often not complete in the ways other kinds of datasets can be: “literature is terminally incomplete. You can record every baseball statistic. You can record every [stock] trade over the course of a year...But you cannot know even *most* of literature, even English literature. Huge swaths of the tradition are absent or in ruins” (March, 2012). Big numbers can create a false sense of completeness. What has been collected and recorded is defined by institutional history and identity. Klein (2013) highlights the importance of recognising “archival silence” in visualising collections.

In addition to assessing the completeness of a dataset against some criteria, there is a related question concerning how important it is to include ‘all’ the data in a visualisation. Completeness in this sense was one of the big claims from the makers of early mechanical visualisations of historical time: through the ability to overview, it was claimed these visualisations offered the possibility to see a ‘complete’ view (Boyd Davis & Kräutli, 2014:p.5). This question is connected to the different ways audiences use collections. Collections can serve as institutional memory, but they are also used by various audiences whose interests overlap with the collection but who may not necessarily be interested in the ‘big picture’ of the collection in itself. It also may not be desirable to show everything in a cultural dataset. This is a big issue (and active area of discussion, see Museum Computer Group (2019)) for cultural institutions putting their collections online if, as is common, not all the data is high quality. Should quality or quantity trump? There is a weigh-up between practicality and principle.

Summary

Time is a fundamental dimension for making sense of digitised cultural heritage collections, and timeline visualisation can support analysis, exploration and presentation of these datasets. Early digital humanities projects focussed on a quantitative approach, and some assume data visualisation necessarily means a quantified view; how can timeline visualisation instead support qualitative, interpretive, experiential engagement? Collection cataloguing provides an interpretive structure for visualising these collections, though different datasets will have different metadata. Museums, libraries and archives have traditionally had different attitudes to cataloguing which is reflected in the resulting datasets. Further, publishing collections on the internet, an accessibility push, and automated cataloguing are changing what cultural datasets look like. How does the structure, content, origins, and intention of cultural data inform its visualisation and for who? To what degree is the metadata available enough for different audiences to make sense of collections and for historical narratives to be made visible? Is there enough context in the data alone?

Cultural timeline visualisations—Survey

How has timeline visualisation been used to trace characteristics in cultural data through time? In this space, Brehmer et al. (2017a) proposed a timeline design space for storytelling. Building on this, Microsoft Timeline Storyteller (Brehmer et al., 2017b) aims to bridge the gap between exploratory and narrating modes. Kräutli (2016b:pp.92-129) reviewed cultural timelines as tools for analysis. I am concerned with the extent to which a timeline design is authored, and review the following examples using Whitelaw’s (2015b) spectrum for collection representations between ‘curated’ (the data selection and its arrangement is handcrafted, with additional annotation/narration) and ‘procedural’ (generated through a formal, usually computational, process operating on the collection data). At the heart of this is the tension between database and narrative, previously explored by Manovich (1999): is the database a democratising force (especially relevant at a time when there are calls for an opening out of perspectives on cultural collections (Cameron & Robinson, 2007)), or a jumble difficult to make sense of without expert guidance?

Highly curated narrative

Museums have long employed timelines in analogue form for exhibition storytelling, for instance the Eames’ history wall for the 1971 exhibition ‘A Computer Perspective’ (Eames Office, 2014). The ultimate ‘curated’ digital timeline involves ordering a handcrafted, limited selection of digital media

along a timeline—through a single, linear sequence—with additional custom text explaining the significance of items and the connections from one to the next in a clear, explicit narrative (eg. Mucha Foundation, 2012; Anne Frank House, 2010; McCloskey & Wei, 2014; Darwin Correspondence Project, 2017a—see Figure 18 and Figure 19). These designs are handcrafted, and tailored to display and storytelling rather than free exploration of collection data.

The Metropolitan Museum of Art’s (2000) ‘Heilbrunn Timeline of Art History’ organises content by time period and geography with the aim of “telling the story of art and global culture through the Museum’s collection”; it combines essays, chronologies, groups of collection items and

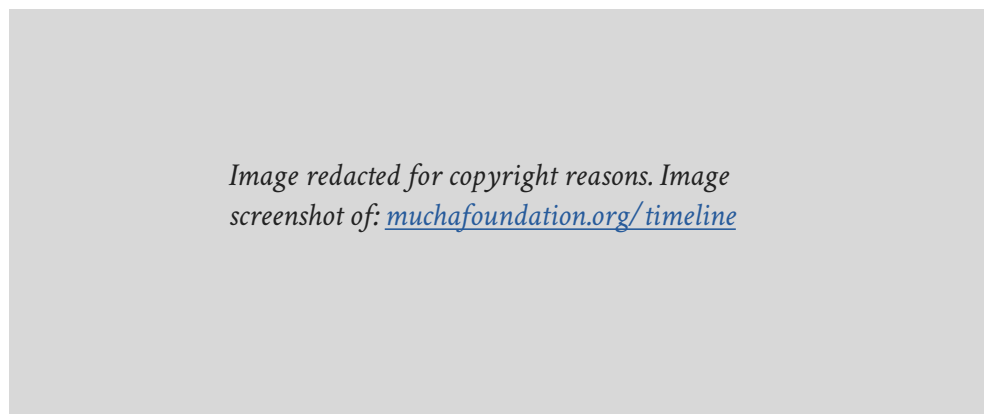


Figure 18. Mucha Foundation. 2012. Timeline.

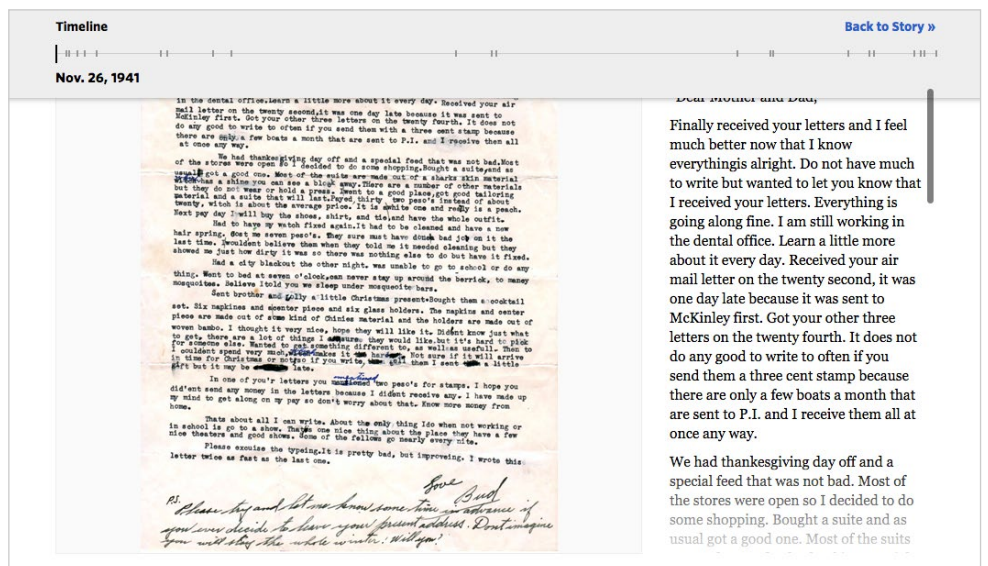


Figure 19. ProPublica. 2014. ‘Bud’s Story, from the Records’ (McCloskey & Wei, 2014).

timelines. Timelines here are not used to map items, but to contextualise them by visualising relevant historical periods (both political and from art history)—see Figure 20. This is a highly authored, encyclopedia-like platform rather than a search/browse interface, where the timelines provide context to curated groups of items.




Figure 20. Metropolitan Museum of Art. 2000. 'Heilbrunn Timeline of Art History'. Screenshot 16/08/19.

Curated groups

Alternatively, curated groups of items (around some characteristic/theme) can be mapped or arranged by time. The British Museum & Google Cultural Institute's (2015) 'Museum of the World' timeline, for example, plots a selection of items organised by curated themes and geography (see Figure 21). Items are represented as abstract dots; lines connect data points of the same theme. But it is difficult to get a sense of how each theme plays out over time, even quantitatively, due to the 3D orientation of the time dimension. Users have to click on a dot to reveal what the object is and to retrieve more information about it. Only one object can be viewed at a time so it is difficult to make comparisons or detect temporal trends.

Cleveland Museum of Art's (2014) 'ArtLens Wall' (see Figure 22) also visualises curated, cross-collection groups of items by themes presented chronologically (Alexander, Barton & Goeser, 2013). Items are represented as images ordered by time so direct comparisons can be made. Exact placement in time is not important here, a collage-like layout is employed; this is a design for presentation and exploration, and rough temporal ordering is enough for the casual observation of visual trends and contrasts.

In these cases, themes and data selection are curated by the museum, and only a limited selection of the collection is available. Users have more freedom to explore than a prescribed linear narrative though: skipping between different themes, choosing their own route through the collection items.

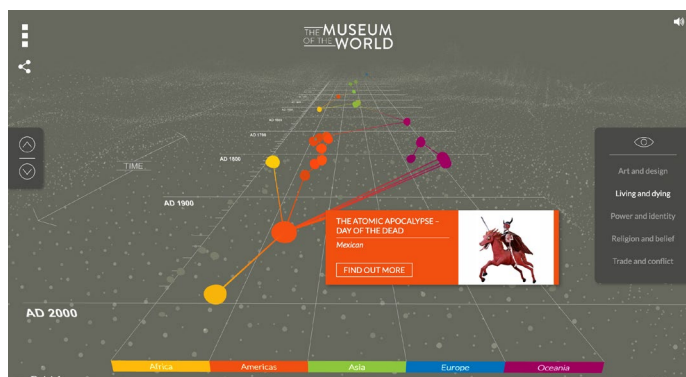


Figure 21. British Museum & Google Cultural Institute. 2015. 'The Museum of the World'.




Figure 22. Cleveland Museum of Art. 2014. Gallery One: 'Love and Lust' timeline.

Quantitative

Timeline visualisation can offer a quantitative overview of cultural data through time, and temporal filtering through interactivity. Quantitative overview can be coupled with other ways to filter data, including: free-text search with snippet views through to texts (Figure 23: The Darwin Correspondence project, 2017b), thematic groupings (Figure 24: Ashmolean: Eastern Art Online, 2013), and filtering by other facets, including automated ones such as colour (Figure 26: Hinchcliffe, 2016; Whitelaw & Hinchcliffe, 2013). And it can be combined with other visualisation strategies for overview, such as tag clouds (Figure 25: Dörk, Pietsch and Credico, 2017).

The timelines may be paired with annotations for context (eg. milestones in Charles Darwin's life represented as dots below the timeline graph, Figure 23: The Darwin Correspondence project, 2017b). This, however, assumes viewing the content against a particular narrative. How useful this is will depend on the users and use, but also the collection scope; it is perhaps less productive for larger, more heterogeneous collections.

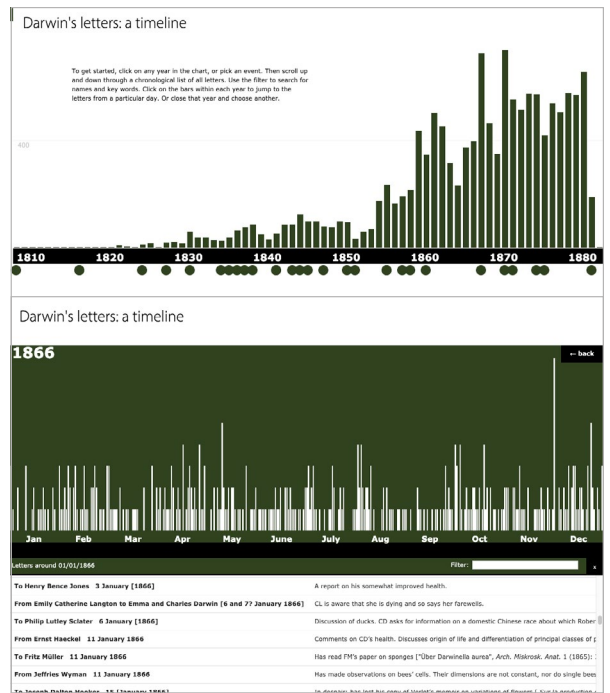


Figure 23. 'Darwin's letters: a timeline', 2017. (Darwin Correspondence Project, 2017b). Created by Surface Impression Ltd.

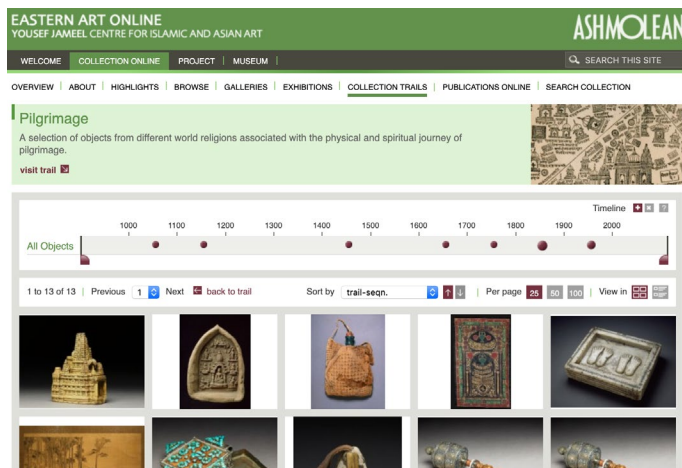


Figure 24. 'Pilgrimage' timeline, Ashmolean: Eastern Art Online (2013). © Ashmolean Museum, University of Oxford.

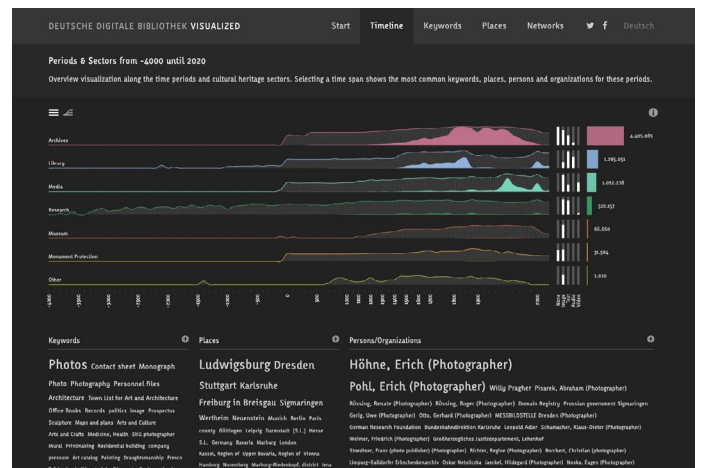


Figure 25. Christian Bernhardt, Gabriel Credico, Christopher Pietsch, Marian Dörk. 2014. 'Deutsche Digitale Bibliothek visualized'.

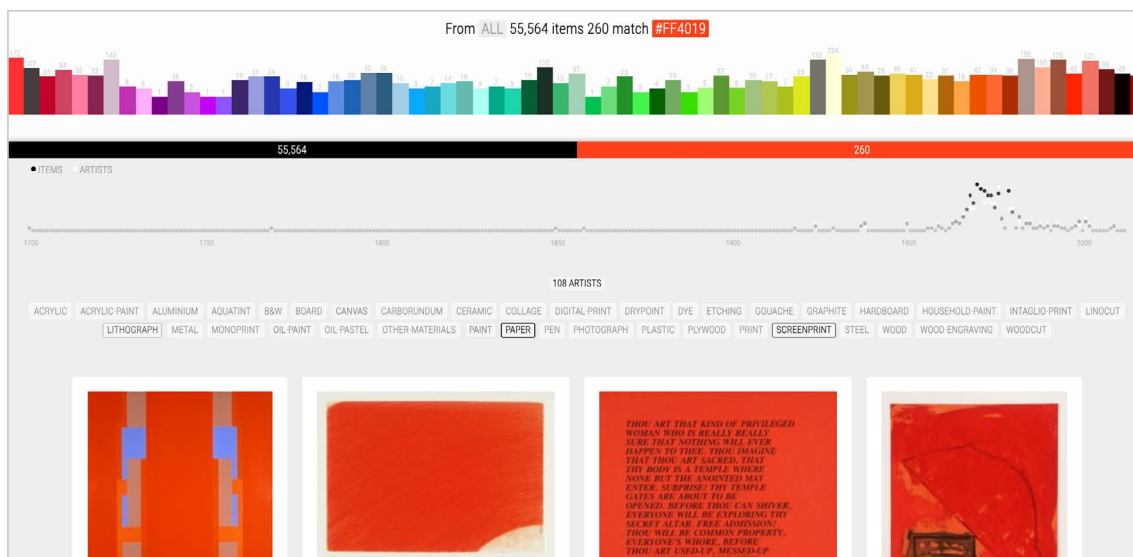


Figure 26. 'Tate Explorer' (Hinchcliffe, 2016).

Chronological grid

Pioneered by Manovich (2011b), ‘media’ visualisation involves plotting images from an image dataset as data points: “keep[ing] enough of the details from the original images. . .to enable the study of the subtle patterns in the data” (p.4). Manovich’s work leans towards visual collections, though this approach can be applied more generally—‘direct’ visualisation (Manovich, 2011a)—whereby whatever the media (text, images, video etc.) elements of the media itself are used as data points, rather than substitute abstract shapes. Individual images are a more straightforward case. Taking this approach with texts, for example, requires some kind of transformation into a more compact form, such as representing a text as a tag cloud.

Examples of Manovich’s visualisations include chronological layouts, such as Figure 27 which organises all 4,535 covers of Time magazine 1923–2009 in a grid. It is a particularly fitting technique for a publication published at a uniform rate, and arranging data in rows plays to the yearly cycle. This visual overview of covers reveals a variety of patterns through time, for example shifts in the colour (black and white to colour, trends of periods preferring particular hues), cover medium (painting, photography), changing content of cover images.

Foo (2016) uses a similar technique to overview New York Public Library data (see Figure 28), but chunking the data into century blocks. Clearly, this approach preferences visual information about a collection dataset over other kinds. Moving between focus and overview is possible, for instance by zooming, but the limits of perception constrain what is possible to discern at overview.

Also employing a chronological grid layout (combined with a quantitative time graph and faceted filtering), ‘Below the Surface’ (City of Amsterdam, 2018) displays a digitised collection of archeological finds (see Figure 29). Here, items are ordered from new to old vertically; vertical scrolling evokes the sense of depth as time in archeological excavation. Metaphor is used to underline what and where these items came from.




Figure 27. Lev Manovich and Jeremy Douglass. 2009. Covers of every Time magazine issue published from 1923 to 2009, arranged in order of publication.

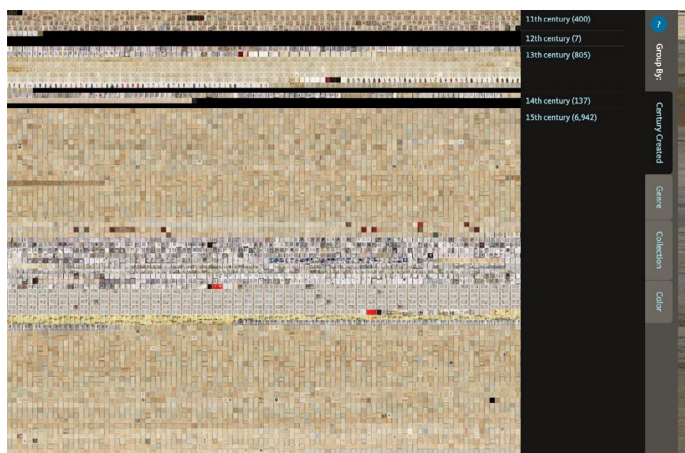


Figure 28. Brian Foo. 2016. ‘NYPL Public Domain Visualization’.

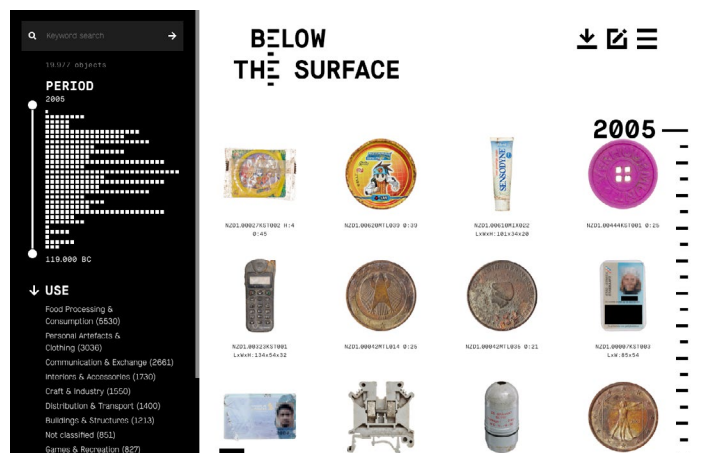


Figure 29. City of Amsterdam, Monuments and Archaeology, Belowthesurface.amsterdam. 2018.

Mapping individual items by time

Timeline visualisation can be used to map entire collections as individual items—items represented as images (Pietsch, 2017 see Figure 30), image crops (Krug, Lerner & Wucher, 2018 see Figure 31), or with different degrees of abstraction depending on the zoom level (Kräutli, 2016c:p.175 see Figure 32)—along a linear time axis. These examples adopt various forms—columns, a stream, a rectangle etc.—and also allow items to be highlighted, filtered or grouped by categories in the data. The ‘VIKUS viewer’ timeline visualisation template (Pietsch, 2017) allows data to be filtered by custom keywords, and text annotations along the time axis contextualise the data.




Figure 30. Pietsch. 2017. ‘VIKUS Viewer’ visualising image data from a collection of drawings by Frederick William IV of Prussia (1795–1861) (Glinka, Pietsch & Dörk, 2017).

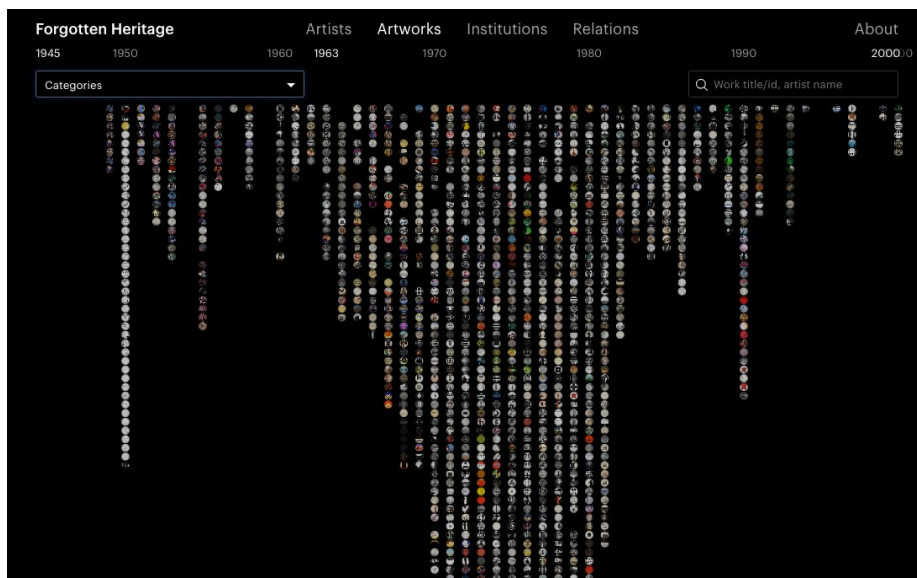


Figure 31. Kilian Krug, Markus Lerner and Severin Wucher. 2018. ‘Forgotten Heritage’: a project led by Arton Foundation and realised together with Luca School of Art, UzF Office for Photography and Kumu Art Museum. Supported by Creative Europe. © 2020 Plural, The Visual Archive: visualarchive.de www.forgottenheritage.eu

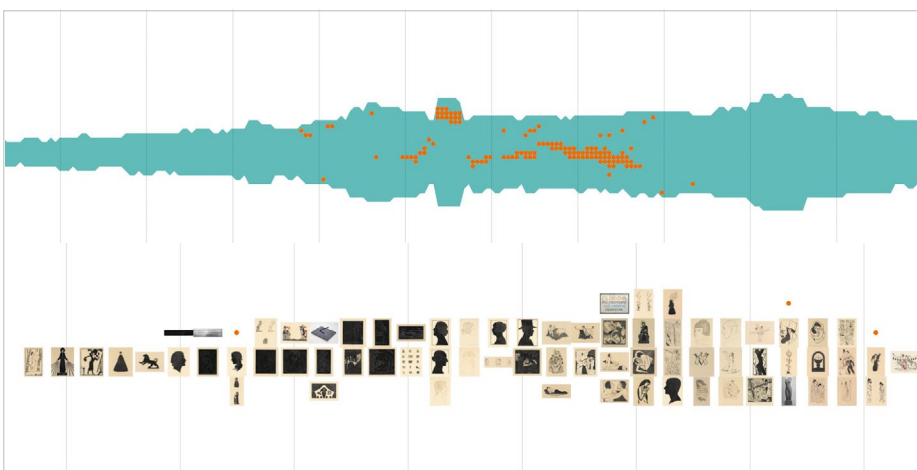


Figure 32. Florian Kräutli. 2016. ‘Timeline Tool’ visualising the Tate’s collection. (Top) colouring orange artworks created by Eric Gill and (bottom) the Eric Gill artworks separated from the remaining dataset (Kräutli, 2016b: p.178).

Jänicke introduced Time-based Impact Mosaics (2018a), and associated Timages (2018b), visualising items through time as an image mosaic (see Figure 33 and Figure 34). Here, individual items are scaled according to a chosen characteristic. This characteristic can be anything, so long as it is quantifiable: Figure 34, for example, shows musicians' portraits scaled by 'popularity' measured by "the total number of references about the musician from online and print media" (Jänicke, 2018a:p.5). Precision in date positioning is waived in favour of aesthetically-pleasing tessellated images.




Figure 33. Stefan Jänicke. 2018. Timages. Visualising portraits of musicians from the MusiX database. (Jänicke, 2018b:fig.5).

Figure 34. Stefan Jänicke. 2018. Time-based Impact Mosaics. 211 paintings with a horizontal golden ratio format. (Jänicke, 2018a:fig.3a).

In 'Coins' (Gortana et al., 2018), images of an historical coins collection are mapped in the visualisation. The data points are arranged and animated as if they are casually dropped piles of physical coins (see Figure 35). The visual layout and animation design evokes the experience of handling the physical items. What the data physically represents can matter for designing its visualisation.

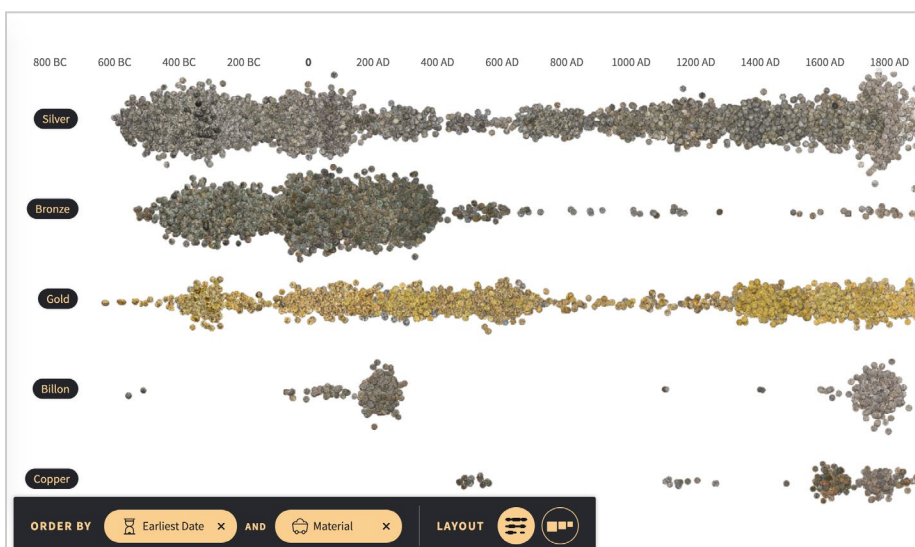


Figure 35. Flavio Gortana, Franziska von Tenspolde and Daniela Guhlmann. 2018. 'Coins – a journey through a rich cultural collection'. Data mapped by earliest date and grouped by material.

These examples are highly interactive, giving the user a high-degree of control in exploring and analysing the data: allowing overview and inspection of details, combined with filtering on chosen data facets and/or search. They are not strictly designed for displaying visual narratives in a collection, but filtering data by particular characteristics may align with patterns that can be interpreted this way.

Plots

If a data facet corresponds to a quantity, then it can be plotted. For example, Manovich, Douglass & Huber (2010) arranged Time magazine covers (1923-2009) horizontally by time, vertically plotted by colour saturation (see Figure 36) or brightness (Figure 37). These timeline plots shows that Time covers' colour saturation increases over time peaking in 1968, though this trend reverses from the beginning of the twenty-first century. It is also possible to spot outlying data points; a number of the saturated outliers (in a strong red) are covers relating to communism (Williford, 2011).

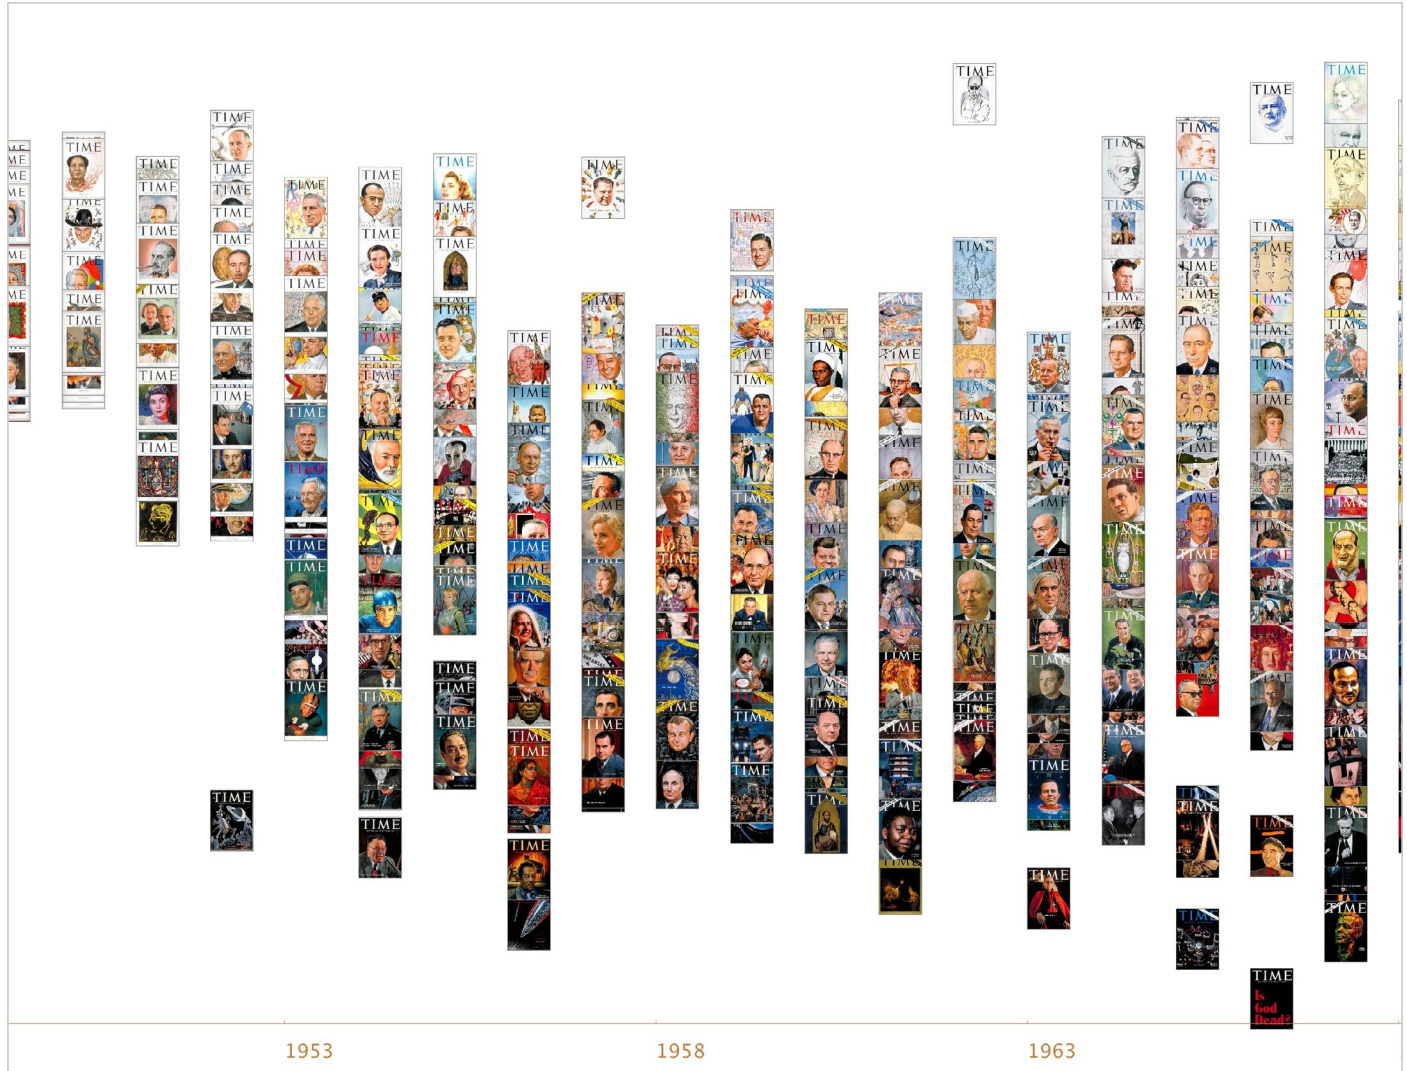


Figure 36. Manovich, Douglass & Huber. 2010. Time magazines cover visualisation. Detail. X-axis: publication date. Y-axis: saturation (softwarestudies.com, 2009).




Figure 37. Manovich, Douglass & Huber. 2010. Time magazines cover visualisation. X-axis: publication date. Y-axis: brightness (visualising data, 2011).

Item-to-item connections

In an unusual example, the Knoll Archive timeline (Gist, 2018) visualises item-to-item connections in the dataset through time. The underlying data here is not visible, but there are similarities with concepts in linked data. The visualised data is presented as a collage of images arranged by time. (Though the dating strategy is a little confusing. ‘Person’ records, for example, have a single year assigned to them but no information about where that year comes from and what it corresponds to).

Selecting an item draws curved lines to other ‘connected’ items along the timeline, fading out unconnected items (see Figure 38). Setting a filter on the item removes all unconnected items from the timeline, redrawing the decade markers and contracting the timeline width (see Figure 39) so that the selection can more easily be overviewed in one view. Only rough temporal ordering is important here. This timeline design is only possible because of the unusual data structure. Explicit connections between items have been made in the data, ie. ‘was mentored by’, ‘collaborated with’, ‘inspired’, ‘designed’ or ambiguously (and most commonly) just ‘related to’. The drawn lines, however, do not distinguish between the different kinds of relationship, and the user must do the work reading through narrative text for each item to figure out what the relationship is.

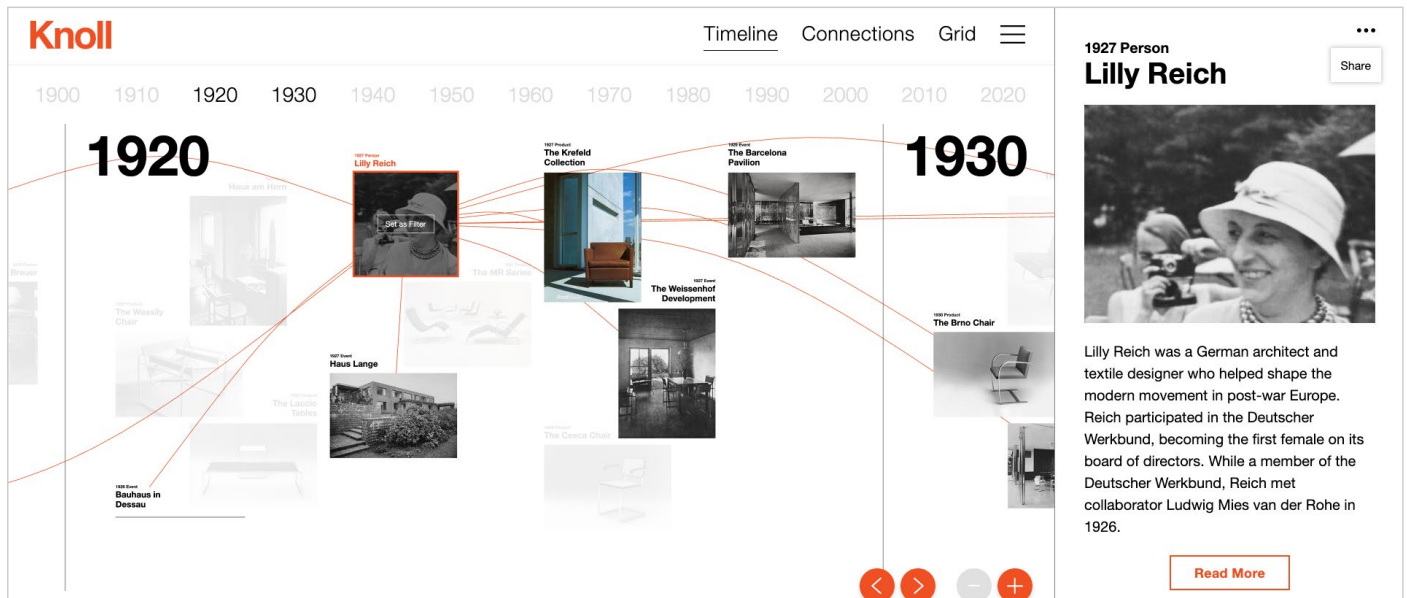


Figure 38. Knoll Archive Timeline (Gist, 2018). Lines and opacity showing items related to Lilly Reich.

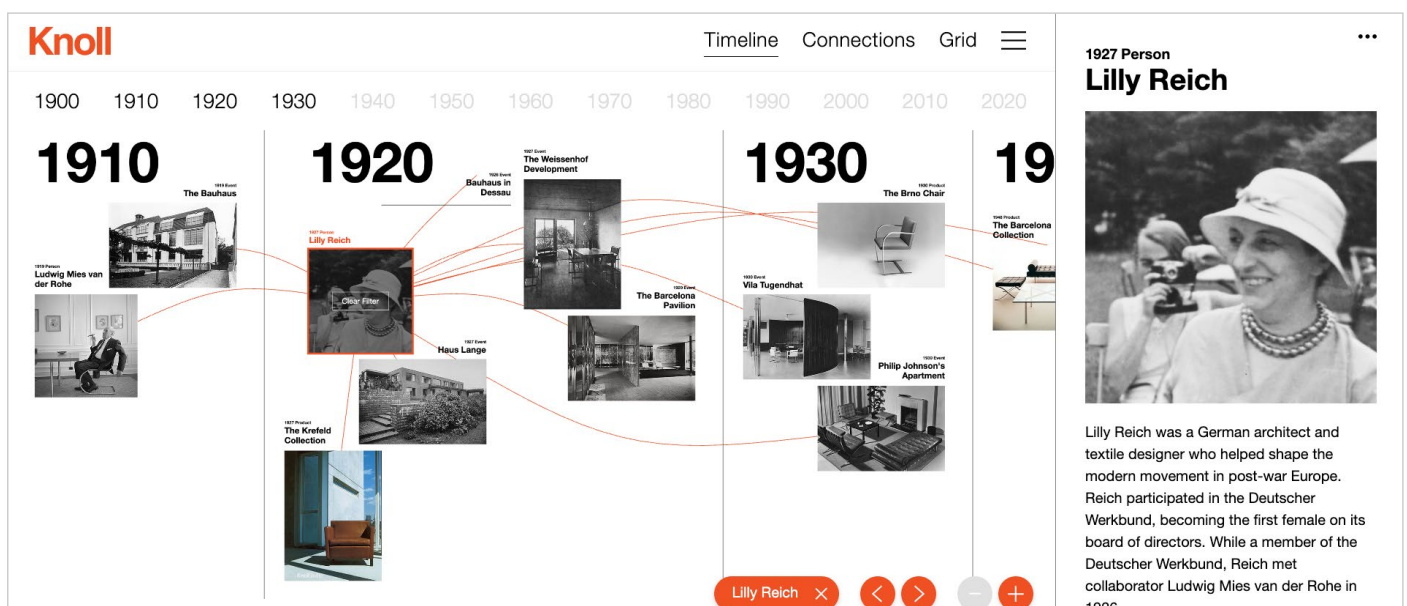


Figure 39. Knoll Archive Timeline (Gist, 2018). By selecting Lilly Reich, only related items are visualised and the date marker arrangement redraws.

Sampling

One way of combining focus and overview in cultural data visualisation is through sampling. DigitaltMuseum, the web platform for Swedish and Norwegian museum collections, offers a timeline view for the results of any search (DigitaltMuseum, 2017), see Figure 40. In this view, time runs linearly but is chunked into historical periods or centuries. Selecting a century further segments that data into decades. Each time segment has a results count and a small sample of items: images with titles presented as if a loosely stacked card pile. This design is a flexible template across scales, but only a very small sample of items can be previewed for each time period.

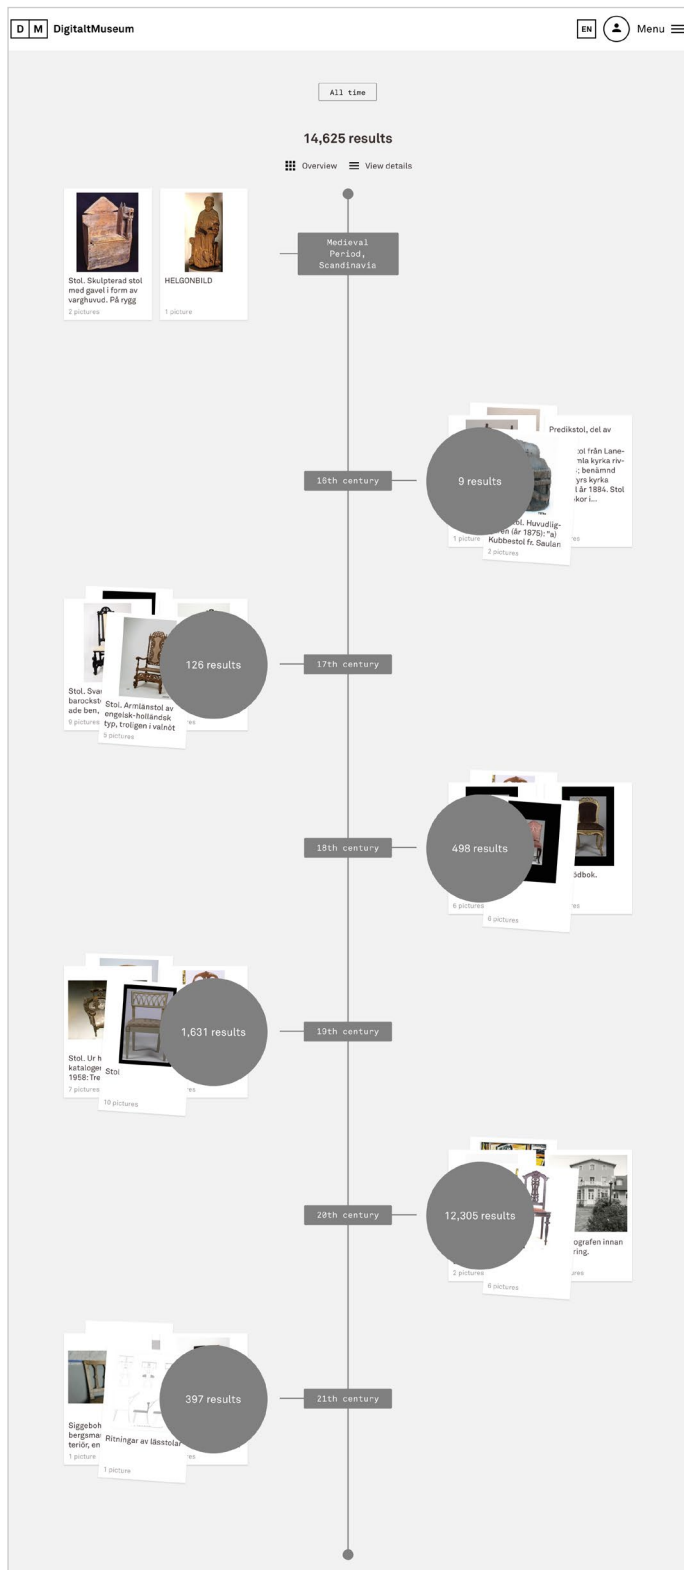


Figure 40. DigitaltMuseum. 2017 Timeline view. (DigitaltMuseum, 2019c).

The ‘Timeline of Modern Art’ at Tate Modern (see Figure 41) also uses sampling as a way to give an overview of collection data but with animation: random items appear and disappear automatically at intervals, inviting serendipitous discovery. It does not use strict mapping of time. This is also an interesting example as, in addition to visualising collection items, it also visualises artist names grouped into art movements, providing art history context to the data items.

Both these examples gesture to the scale of the collection, providing tastes of what it contains, while its quantitative shape is obscured or concealed.

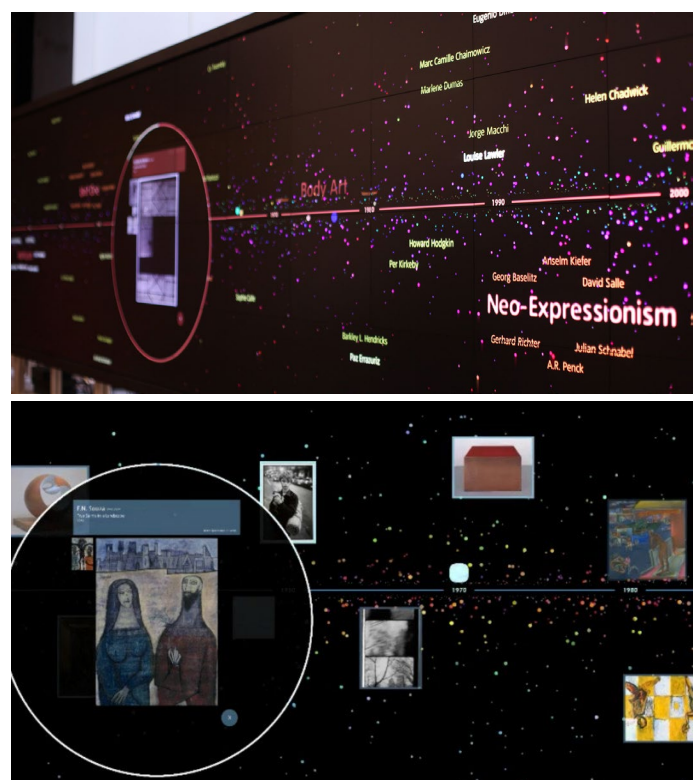


Figure 41. Tate ‘Timeline of Modern Art’. 2015. Images courtesy of Framestore, Tate.

Summary

Timelines can be employed in various ways to visualise cultural data: quantitatively, visualising individual items in aggregation, and sampled or curated groups offering overview or storytelling. Different examples look more like a graph, or a visual narrative, or something in between. The data structuring and content define what is possible in terms of visualisation, and the character and origins of the collection represented may inform the visualisation design—these factors have a bearing on the generalisability of designs. These examples show a spectrum between designs where explicit narratives are presented, to those where implicit ones may be visible. Essentially, to what extent is the display explained/contextualised?

Timeline visualisation can offer a qualitative view on cultural data in a number of ways. ‘Direct’ visualisation—plotting images for image data, or text for text data—spotlights the items the data represents and allows direct comparison. If categories have been employed in the data, they can be used for grouping, filtering, and—if they can be expressed numerically—plotting or scaling data points. Visual and interaction design can be employed metaphorically to evoke the experience of physically handling items or the collection’s origins. Designs for more casual audiences may preference factors such as pleasing aesthetics and visual overview in one glance over precise positioning of items by date, time axis granularity, and a uniform scale.

Research Questions

This thesis is concerned with timeline visualisation design for exploring cultural data. Through cataloguing, cultural institutions create value and meaning around the items in their collections. By plotting these items by time, using the categories with which they have been described, historical narratives may be made visible. This PhD asks: **is 'just the data' enough?**

By 'just the data' I mean using a procedural (formal, computational) process to visualise the cultural dataset, leveraging the categories that exist in the data, rather than curating/handcrafting the data selection and arrangement. The procedural approach has the benefit that a greater scale of data can be visualised than time would allow manually. For some, like Joseph Priestley, there is a moral benefit: the meaning is left to emerge for itself. But is it enough? To what extent is 'just the data' sufficient to visually tell the stories that people (who?) find valuable, versus the desire (whose?) for authorial intervention? And what factors have a bearing on sufficiency and success here?

Of course, what 'just the data' is depends on the collection in question, and cultural datasets are both diverse and changing. This gives rise to a series of subsidiary questions. How does the structure, content, origins and intention of cultural data inform its visualisation and for who? To what degree is the data available enough for different audiences to make sense of? Is there enough context in the data alone? How does completeness, of collections, datasets, and visualisations, play a role here?

Methods

“How can I tell what I think till I see what I can make or do?” (Frayling 1993:p.5).

This PhD takes a Research Through Design approach: I addressed my research questions by designing/making prototype timeline visualisations of real cultural datasets, in dialogue with collection experts. I evaluated designs with prospective users through semi-structured interviews.

Institutions and datasets

I developed a varied portfolio of visualisation prototypes with datasets from different cultural heritage institutions:

- the Medical Officer of Health reports at the Wellcome Library, London
- the Cooper Hewitt Smithsonian Design Museum collection, New York City
- the Royal Photographic Society collection at the Victoria & Albert Museum, London
- historical portraits at the Nordic Museum, Stockholm from Swedish Open Cultural Heritage data

These datasets cover a breadth of topics (public health to photography), collection scales (around 5,000 to 200,000 items) and item types (texts, design objects, photographs, artworks). While I do not pretend this is a universal investigation, working with diverse datasets allowed me to extend my investigation domain and to engage with a range of stakeholders/users.

Many of the reasons for choosing these particular datasets were pragmatic: encounters through conferences or events, interested individuals at institutions etc. The Smithsonian is a partner in the AHRC's (UK Arts and Humanities Research Council) International Placement Scheme, from which I received an award, enabling me to spend time at Cooper Hewitt.

The practical arrangements for these projects also differed. For some I was working remotely, in contact with institutional partners by email or video call. For others, I was embedded within the institution, attached to a team there. Again, the reasons for these different arrangements were largely pragmatic: was my relationship to the institution in the form of a fellowship/placement, or not? From these experiences, though, I learnt that there are significant benefits to being embedded within an institution for this type of project. Meetings with collection experts within the institution and users are easier and quicker to arrange. It is easier to get technical help with the collection data and serendipitous opportunities arise as there is greater awareness internally of the project.

All the datasets I worked with in this PhD belong to individual institutions, rather than e.g. thematic collections on aggregator platforms like Europeana which bridge many institutions. While there is considerable interest in linking up different collections (sometimes described as data silos), this choice made it easier to encourage engagement from partner institutions. Also, there can be additional technical work visualising data from different collections if the data structures are different. In any case, the visualisation techniques created in this PhD could easily be adapted to data bridging collections.

Research through design in digital humanities

Before describing my process in the visualisation projects, it is helpful to give some context to design as a research tool in digital humanities.

Since the introduction of computational methods in humanities research at the midpoint of the 20th century (Hockey, 2007), it is increasingly common for digital humanities projects to develop software including graphical user interfaces or visualisations. Traditionally the purpose of such software was seen as enabling humanities scholars to carry out their research more quickly and easily: research outputs were the insights made by scholars applying traditional humanities methods of inquiry through digital objects/tools. As digital humanities activities and outputs have developed, though, so have the boundaries of what counts as scholarly output. It is now argued that *making* software, can be a research activity in and of itself (Ramsay & Rockwell, 2012; Galey & Ruecker, 2010). The questions that arise and insights made during the design and development stage of archival interfaces/visualisations are valuable in their own right (Kräutli & Boyd Davis, 2016; Schofield, Whitelaw & Kirk, 2017). Evaluation (and institutional recognition) of digital outputs in the humanities, however, still remains a tricky issue (Smithies, 2012).

With increasing numbers of interfaces and visualisations developed in digital humanities projects, design—visual, interaction, information design etc.—has a growing role in this space, though presently often in the absence of designers (Burdick, 2009). More attention is now being paid to the role that designing can and does play in making digital humanities tools (Caviglia, Coleman, & Ciuccarelli, 2012; Lorber-Kasunic & Sweetapple, 2015).

Further, there is growing interest in connecting the “epistemology of building” (Ramsay & Rockwell, 2012) in digital humanities with Research through Design (RtD) (Owens, 2011; Schofield, Whitelaw & Kirk, 2017; Kräutli & Boyd Davis, 2016) which has its roots in Art and Design schools. RtD is an approach to investigating an issue/area/question through the creation of designs. The origins of RtD are connected with the Royal College of Art and, particularly, the work of Bruce Archer. Archer (1995:p.11) characterised RtD, writing: “there are circumstances where the best or only way to shed light on a proposition, a principle, a material, a process or a function is to attempt to construct something, or to enact something, calculated to explore, embody or test it”.

Cultural visualisation prototyping

My approach to designing cultural visualisations is characterised by iterative prototyping, and I begin building prototypes early on in a project. (I use the term prototype loosely, meaning mock-ups, examples, static and interactive artefacts from low- to high-fidelity, demonstrating and testing ideas (Rudd, Stern & Isensee, 1996)). Designers use prototypes as a way to externalise their ideas in visible (and, in this domain, interactive) form (Arrigoni & Schofield, 2015).

Prototypes serve a number of roles in progressing the design process. They augment the designer’s ability to think (Norman, 1991) and tentative prototypes can imply possible directions in which a design could be developed (Goldschmidt, 1991). They serve as a form of feedback and can drive dialogue and interactions between the designer and project partners/users; they communicate, demonstrate and can be used to test ideas.

Returning to the ideas of RtD, prototyping also serves as a way to progress research, as creating designs helps in defining (and possibly re-defining) the problem/brief. There is now growing recognition of design as a research tool in the cultural visualisation context. Boyd Davis and Kräutli (2015; 2016) describe how the development and evaluation of prototypes not only

helps to provide evidence for the original focus of an investigation, but can also raise new research questions. Hinrichs, Forlini, & Moynihan (2018) discuss the role of early prototyping, as they phrase it: “research thinking through visualization”. And, drawing on digital humanities visualisation examples, McCurdy, Dykes and Meyer (2016) discuss design as a research method in applied visualisation research.

Starting a project with making, possibly with a target user group not yet determined, may initially seem the wrong way round. And it contrasts with taking a more scientific outlook, for instance characterising known issues and seeking to address those or seeking to facilitate a predefined task. But early prototyping is valuable for pinning down or improving the research question being investigated; it also provokes and inspires new questions that may not have been considered.

Prototyping is also particularly productive in the cultural data context for better understanding a dataset’s qualities and scoping out what is feasible. As Whitelaw (2015a) describes:

Notably in order to understand the features of a [cultural] collection that might be represented, we must first represent the collection: the riches and voids in each collection are only evident through a process of exploratory visualisation...For a designer this process presents a challenge in that preconceived visual or information structures rarely survive contact with the collection data. At the same time, exploratory visualisations can lead to further work with the metadata, which in turn transforms the potential representations.

I experienced this myself in an early project, at the beginning of this PhD, with a dataset of historical theatre adverts. I began the project with clear ideas about a visualisation I would like to build. When I started working with the data, however, it became clear that my planned visualisation was not feasible. My intended design required extracting text data from the adverts, which was not in the existing metadata, with Optical Character Recognition (OCR). The OCR performance, however, turned out to be too poor (as a result of idiosyncratic typefaces). After this experience, I shifted to using early, exploratory prototyping with the dataset as a way to scope what might be viable and suited to the dataset’s qualities.

In this PhD, I designed/made cultural visualisation prototypes as a way to:

- familiarise myself with the content and character of the dataset, and to provoke ideas for possible design directions
- shed light on practical and intellectual issues in visualising real cultural data (through encountering problems and challenges first-hand)
- reflect on and refine my research questions
- shape my developing designs through iterative prototyping
- as a locus for feedback in dialogue with collection experts, negotiating with their perception of what is possible, and provoking ideas for new possible directions in which to develop the design
- demonstrate and test ideas with users

Engaging collection experts and users

In a cultural visualisation project, working with collection experts (curators, cataloguers, collection managers, archivists, digital/technical staff etc. internal to an institution) is critical for understanding the dataset being visualised. D’Ignazio and Klein (2019) and Loukissas (2019) stress that understanding the context of data, including the organisational motivations for its collection and use, matter deeply for using it effectively. Kräutli (2016b) has written about the biases,

patterns and strategies that hide behind cataloguing in cultural datasets. Collection experts can provide knowledge about the technical, historical, practical and intellectual forces that have shaped a cultural dataset: essentially, why the data is the way it is and what it means. (In larger institutions collection management responsibilities and roles are more compartmentalised. I often use the general term ‘curator’ in these cases.)

Collection experts are often also responsible for delivering digitised collections to audiences. They have knowledge about who these audiences are and their interests and requirements. Collection experts can, therefore, help set a visualisation brief by pointing to potentially fruitful approaches.

This PhD follows an approach used by Kräutli (2016b), who demonstrated that dialogue with collection experts while prototyping is a productive method for exploring the ideas at the root of a visualisation design. (Other projects have run co-creation workshops (Glinka, Pietsch & Dörk, 2017; Chen, Dörk & Dade-Robertson, 2014) or cross-disciplinary collaborative projects (Hinrichs et al., 2017; McCurdy et al., 2015)).

I used frequent contact with collection experts in the initial, exploratory prototyping stage. Once my design was more settled and prospective user group/s were clearer, I arranged user evaluation sessions. User groups for cultural visualisations include:

- collection managers
- scholars/researchers
- educators
- casual users (for instance museum visitors)

During the course of a project, as my visualisation design took clearer shape, I sometimes identified a narrower prospective audience. For example, visualising Cooper Hewitt’s data by colour introduced historical colour experts as new, unexpected audiences to my design. I sometimes found that my early prototype visualisations could be relevant to multiple audiences. Since, however, these different audiences have different requirements and expectations, developing a design further would require preferencing a user group. My later projects are more focussed on casual audiences, reflecting the priorities of the institutions I was working with.

Evaluating cultural visualisations

The focus of this PhD is in exploring what timeline visualisations of cultural data can show/do (for various audiences); the approach is exploratory and qualitative, drawing out questions and subjective insights. This PhD is not a scientific study evaluating usability or other task-based assessments of cultural interfaces with users, for instance Morse et al. (2019) and Speakman, Hall & Walsh (2018).

Evaluating cultural visualisations continues to be a sticky issue: how should this work be evaluated and according to what disciplinary standards? Coles (2016) writes, “do results have to be replicable? Does subjective experience actually ‘count’, and if so how? How do we value and judge aesthetics? How does creativity come into play?”(p.1). Coles goes on to strongly argue that, since digital humanities is a field serving the humanities, humanities-based values and methods should be incorporated into the entire research process “from design to evaluation” (p.4). Coles recommends “asking—and trusting answers—about user perceptions of their research experience, and using domain experts to assess the quality, originality, and persuasiveness of the arguments and other research products arising from these experiences”(p.4).

During prototyping, I carefully reflected on my design decisions and engaged in dialogue with collections experts. Once a prototype was sufficiently developed, I used semi-structured interviews with prospective users (relevant to the prototype in question) to capture rich impressions, interpretations, and feedback. Kräutli (2016:pp.37-39; *ibid.*:pp.156-158), Loukissas (2014) and McCurdy et al. (2015) used interviewing as a productive method for exploring the ideas behind a cultural visualisation with users and ensuring feedback is not solely about basic functionality.

Insights, engagement, aesthetics

Evaluating what a cultural visualisation *does* can be tricky. Jänicke (2016:p.4) describes that “for a humanities scholar, usually, a visualization rather provides a novel perspective on cultural heritage data triggering unfamiliar thought processes than that it represents *the solution* for her research question”. Similarly, McCurdy et al. (2015:p.9) argue that since cultural visualisations can be validly interpreted in different ways, comparing against ground truths is unproductive. Furthermore it is useful to consider not only if a visualisation supports making new insights, but if new *kinds* of insights are supported and how the insight gathering process may have changed with the visualisation.

Recently there has been growing interest in the role that aesthetics play in cultural visualisation. What is meant by aesthetics in this domain? While aesthetics is often discussed in terms of immediate visual appeal, here it is productive to also consider the aesthetics of interaction and that aesthetic impression is contingent on the context of use (Petersen et al., 2004). McCurdy et al. (2015) argue “*pleasure and enjoyment* are productive research outcomes [in cultural visualisation]... Creating a visually pleasurable research environment goes beyond general aesthetics, encouraging richer exploration and greatly increasing the overall efficacy of the tool” (p.9). Aesthetic experience is also important for bridging between disciplines; Lamqaddam et al. (2018) recommend, in the Art History context, that attracting more interest in visualisation there requires “putting aside the aesthetic standards used in the software world to propose visualizations users will be satisfied to engage with” (p.3). This point is illustrated in detail by Coles (2016), here discussing a poetry visualisation research project:

...the scientists assumed that any user of such a tool would be primarily seeking information, grounded in data, while we poets assumed that users would be interested in an experience in which aesthetic pleasure plays a major part...What none of the team members foresaw at the time was that the aesthetic qualities embodied in the visualizations—which bear a striking (and, to the poets, perhaps comforting) resemblance to a hand-created mark-up of a poem, would help provide not only aesthetically-enriched experience but also insight (p.2).

As Coles sums up of the poets using the visualisation tool: “we are after engagement over results” (Coles, 2016:p.3). Solely asking if the visualisation supported a scholar learning something new may not be the best question to ask.

Interview evaluations

To explore impressions of the prototypes developed in this PhD, I conducted semi-structured interviews with relevant collection audiences either one-on-one or in small groups. Interviews were audio-recorded, transcribed, and analysed for emerging themes. In each session, I introduced and demonstrated the prototype visualisation design. Where possible, I invited the interviewee to try the prototype out themselves, following their own interests.

I developed an interview script in my first project which I used as a base to adapt in later project evaluations. This interview script combined general questions about the interviewees’ impression

of the interface with specific questions about particular design features. While my focus in evaluating designs was not primarily basic functionality and usability, these concerns are still relevant. Improved user experience can increase a prototype's persuasive power, as users are better able to see its potential. Usability improvements also help smooth future feedback sessions, by focussing attention back to the intellectual issues.

I reflect on my methods/evaluation choices in the 'Discussion' chapter at the end of this thesis.

Technical details

All the interactive prototypes I built in this PhD are web interfaces built in JavaScript programming language, using the JavaScript visualisation library D3 (Data-Driven Documents, 2019). Using web technologies comes with many benefits, including access to a wealth of open-source resources and straightforward interfacing with online collection APIs (Application Programming Interface). It is also fitting as publishing collection data on the web is shaping and changing it.

For projects where the dataset was not prohibitively large, I loaded the data directly into the browser's memory. Otherwise (or if full-text search was needed), I used Elasticsearch (2019a) as a backend. I used Python programming language for data processing (for example converting colour data for collection objects to a different colour system), web scraping and building Elasticsearch indices prior to visualisation.

This suite of tools enabled me to work with different institutional datasets, where different data formats (JSON, JSON-LD, XML etc.) were used. In these projects, I worked with static data sources as well as accessing collection data over the internet using APIs.

I had some basic experience using web technologies before starting this PhD but, largely, I learnt these tools through my early practical work. I do not think this learning process had a noticeable impact on my design ambitions, though my prototyping speed increased a lot through the PhD (both because of my improved skill and knowledge, and because I could re-use code snippets in later projects). I also made visual mockups and animated demonstrations using design and video editing software.

D3 is known as a flexible tool for making custom visualisation designs and, largely, I feel it allowed me to fully realise what I wanted to do intellectually. As with all tools, it still comes with constraints in terms of what is feasible and what can be accomplished quickly. For instance, D3 offers a number of templates for chart types/visual design forms and methods which speed up implementing certain types of visualisations/interactions. I have tried to acknowledge, in my project chapters, where my design decisions were influenced by these sorts of factors.

I mostly found the tools I used transparent in their workings; I have tried to acknowledge where this was not the case and its impact on my designs. On the whole, I found the frequent lack of transparency in the intellectual and technical decisions shaping the creation of cultural datasets I worked with more of an issue, which I reflect on in my final 'Discussion' chapter.

Visualisations

The following chapters walk through the portfolio of visualisation prototypes developed in this PhD. In each, I introduce the collection/institution I worked with, explain my design motivations, record prototyping, the visualisation design, and user evaluations. Prototypes are presented in chronological order.




Chapter 4

'Steptext'
Medical Officer of Health reports,
Wellcome Library




Chapter 5

'Tag-timeline'
Cooper Hewitt Smithsonian Design Museum



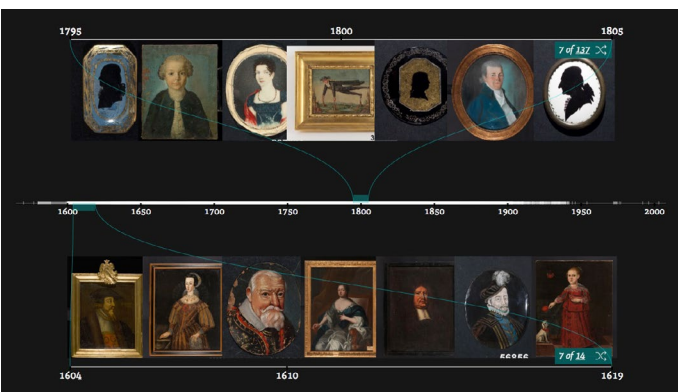
Chapter 6

'Dive into Color'
Cooper Hewitt Smithsonian Design Museum



Chapter 7

Exploratory visualisation
Royal Photographic Society collection,
Victoria and Albert Museum, London



Chapter 8

'Faces of Sweden'
Nordic Museum, Stockholm

‘Stext’

Medical Officer of Health reports,
Wellcome Library

Figure 42. ‘Stext’ visualisation of MOH reports data: ‘nurse’.

In this first project, I worked with the Medical Officer of Health (MOH) reports dataset. This is a collection of London public health reports dating between 1848-1972 digitised by the Wellcome Library. I was made aware of the dataset through a meeting with the Wellcome Library and decided to work with this digitised collection because it included high quality, complete transcribed text data for every document and it was sufficiently large (roughly 5,500 documents) to address my research questions, but not so large as to be unwieldy to work with.

The Wellcome Library digitised the document collection by photographing the reports cover-to-cover (see Figure 43), applying Optical Character Recognition (OCR) and post-OCR manual text proofing. The resulting text data, therefore, has a high degree of accuracy. The full text corpus is available to download as raw text (Wellcome Library, 2019c), and photographs of all the page images from each document are available on the Wellcome Library (2019a) website. I worked with the London reports, which were available at the time. At the time of writing, the Wellcome Library is in the process of digitising the full national MOH reports set.

The documents themselves are largely narrative text. The reports were produced annually by the MOH for each London borough and consist of a mixture of narrative text concerning the state of

public health in that district (including food safety, housing conditions and child welfare) and tables of statistics (e.g. incidence of diseases and infant mortality rates) (Wellcome Library, 2019b). The number of MOH reports per year varies, with a greater proportion in the 20th Century (see Figure 44). They also vary somewhat in structure, content, and in length (from over 100 pages, to fewer than 20). The raw text download for the reports was not semantically chunked, so it is not easy to, say, ignore anything in a table, or to search only headings. In addition to the raw text, the only other metadata available for the reports is the London Borough name and the publication year.

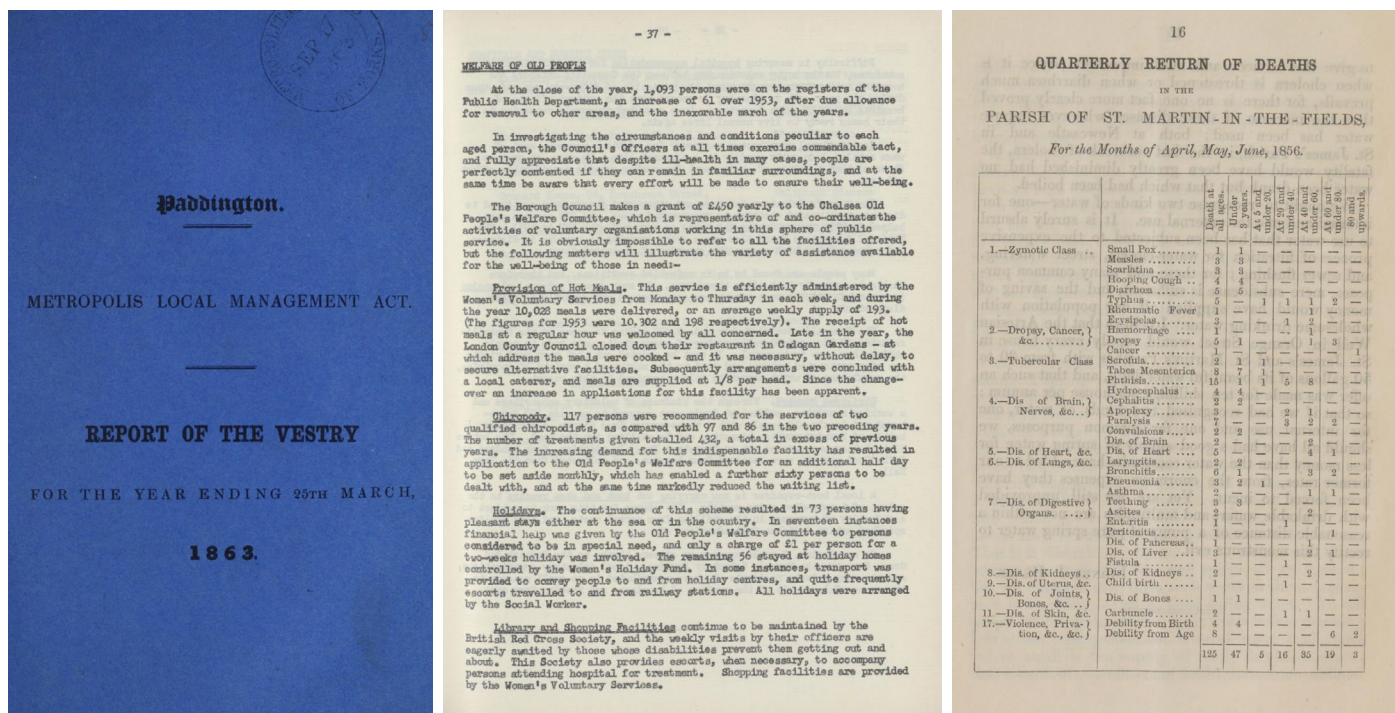


Figure 43. Example cover and internal pages from Medical Officer of Health reports, from the Wellcome Library website (2019). Left to right: Paddington 1862, cover; Chelsea 1954, p.37; St Martin-in-the-Fields 1856, p.16.




Figure 44. Line graph showing number of MOH reports per year.

This collection is primarily of interest to historical researchers. In addition to a record of local public health, the reports offer rich insights into changes and developments in everyday life in London, and the wider country (Wellcome Library, 2019d). In this project, I was interested to design a tool that would support these researcher users.

Exploring historical texts with visualisation

Tracing commentary through time is an important activity in historical research, particularly when surveying primary sources at the beginning of a research project. Dialogue with humanities scholars, in casual conversations at the beginning of this project and at the Wellcome Library's (2016) 'Data Week', confirmed to me that being able to carry this out more fluently with the bodies of digitised material already available would be very valuable.

How can timeline visualisation support exploring commentary across these texts? This collection has minimal metadata—just the year and London borough of publication—so this question really hinges on what is possible with the text data itself, and what manner of engagement with the text data is desirable for researchers. Text mining/analysis techniques, for example, can be used for 'distant reading' to explore and analyse collections of documents. Text analysis programs, such as AntConc (Anthony, 2019), use word counts, frequencies, collocations, and other methods to find patterns and make comparisons between texts at scale. (Text mining methods have previously been applied to the MOH reports (Toon, Timmermann & Worboys, 2016; Thompson et al., 2016)). Tools such as Voyant (Sinclair & Rockwell, 2019) can be used to visualise text analysis data in a variety of ways, including by time. These tools, which come from the corpus linguistics tradition, could be said to offer linguistic analysis of texts focussing on use of language and grammatical structure. An historical view on text data, by contrast, is interested in language as semantic content; there is, therefore, a tension in using these tools for historical work.

Topic modelling is an approach to statistically identify thematic 'topics' (recurring patterns of co-occurring words) across a corpus (Brett, 2012). Various visualisation techniques use topic modelling to provide a macro-view of the multiple themes present across texts (Collins, Viégas & Wattenberg, 2009; Havre, Hetzler, & Nowell, 2000; Cui et al., 2011; Cui et al., 2014; Liu et al. 2012), including in the digital humanities context (Guldi & Johnson-Roberson, 2019; Crymble, 2013). There are, however, calls for caution concerning the application of topic modelling in the digital humanities. As Schmidt (2012) argues, 'topics' are ambiguous and can be misleading. The most frequent words in a topic are not necessarily the ones that create the meaning and the statistical model may not be sensitive to linguistic and semantic shifts over time.

In contrast with statistical techniques, there are comparatively simple approaches to visualising text collections, for instance building off the keyword-in-context technique (Fischer, 1966) which is the standard way to return search matches in digital archives. In the 'Trading Consequences' project, Hinrichs et al. (2015) argue for the importance in humanities visualisation of providing context for visualised data. Providing context, for instance giving access to text snippets from the individual documents, aids fluid exploration of texts and helps users avoid misinterpretations. Visualisation techniques such as Word Tree (Wattenberg & Viégas, 2008) build on keyword-in-context to offer an interactive visualisation for exploring repetition in the contextual words following a phrase across texts. In contrast with topic modelling approaches, this technique allows the user to query the documents themselves: to direct exploration to the themes they are interested in. While Word Tree allows users to see the different contexts within which a word has been used across texts, it does not make visible the temporal dimension. Keyword-in-context has also been used to provide a "poetic" interface to digitised historical documents (Post, 2018), though again without the temporal dimension.

Keyword-in-context may be a crude stand-in for exploring commentary in texts around themes, but it is a simple and interpretable approach and allows a user to pose their own queries of the texts. My goal in this project was to build on keyword-in-context to design a timeline visualisation tool supporting exploration of the texts.

Prototyping

My starting point for prototyping in this project was experimenting with plotting text snippets around a matching keyword by date. By using snippets for each occurrence, instead of abstract dots, the user can glimpse, in situ, the context of keyword use.

My goal here was to fit as many time-ordered data points into the display as possible, to support an overview, while also giving a sense of transparent access to the content of a large number of reports. I trialled spreading out snippets at random vertical positions, horizontally mapped by date using D3's force layout (D3 API, 2019a; see Figure 45). The result is a mess and difficult to read or discern structure, so I tried other layouts, including arranging the snippets vertically by date: plotted vertically by year (see Figure 46); and ordered vertically by year (see Figure 47). This second approach enforces chronology on the data, and the snippets are still legible. (Indentation has previously been used in data journalism for temporal visualisation of texts (Groeger & Currier, 2012), though not building on keyword-in-context).



Figure 45. Prototype visualisation of MOH reports data: 'typhoid carrier'.

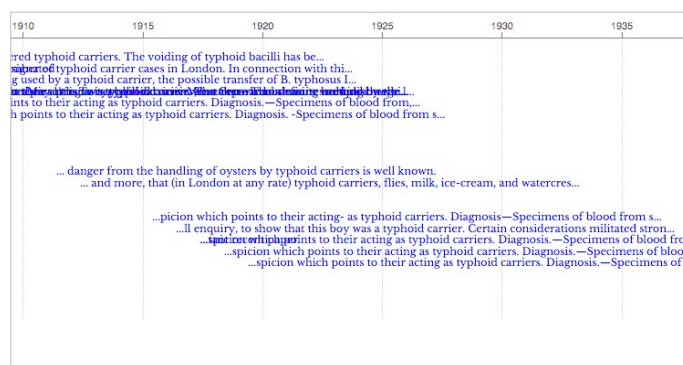


Figure 46. Prototype visualisation of MOH reports data: 'typhoid carrier'.

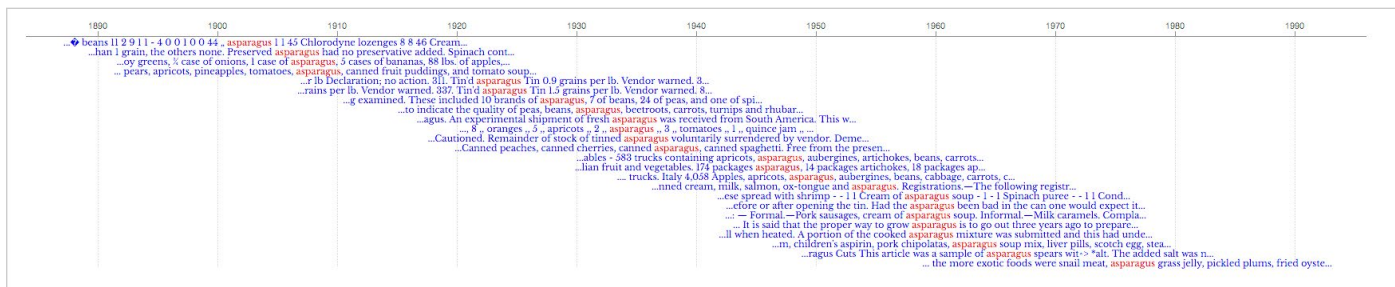


Figure 47. Prototype visualisation of MOH reports data: 'asparagus'.

Visualisation design

The final prototype, sufficiently developed to test with users—'Stepext' (see Figure 48)—visualises keyword instances in the MOH reports mapped by time, with sufficient context to get a sense of how the keyword is being used. It is a simple, minimal visualisation technique for exploring texts, emphasising time relations.

Keywords can be queried via a search box. Figure 48 shows the top section of the resulting visualisation for searching for 'nurse' across the MOH reports. (The visualisation can be scrolled through vertically, allowing the user to see later snippets). A date slider can be used to limit the time span searched.

Each text snippet, visualised at legible size, is horizontally centred on its date. Snippets are vertically ordered from oldest to newest. By using a fixed-width typeface, the alignment and

relative indentation of snippets visually reinforces their temporal relation. Interview evaluations with historians confirmed this snippet length (100 characters) is sufficient to get a sense of how the term is used in the text.



Figure 48. Detail from visualisation of MOH reports data: ‘nurse’, showing tooltip on hover.

The search term can be either a single word or phrase (for brevity, I shall just say keyword). The keyword search returns exact spelling matches and plurals are not automatically included. Keyword search is not case-sensitive. When searching for multiple words, the interface also returns instances separated by punctuation. For example, searching for ‘ice cream’ returns results including hyphenated ‘ice-cream’. For this prototype no natural language preprocessing, such as tokenisation or part-of-speech tagging, was implemented, but for future versions it would be possible to make this modification. (The back-end is an Elasticsearch (2019a) index of the MOH reports raw text; Elasticsearch functionality drives the full-text search and snippet creation).

A horizontal date axis runs across the top of the visualisation; the date extent of this axis is redrawn each time a new query is made to match the temporal extent of returned snippets. The user can scroll horizontally (along the time-axis), or vertically. (Though feedback sessions suggested also allowing click and drag with the cursor would likely be more user friendly for moving horizontally along the visualisation). The time axis remains fixed and visible at the top of the page during vertical scrolling so the date markers are always visible.

The keyword instance/s within a snippet have a colour highlight and generally appear in the centre of the snippet (thus the horizontal centre of the highlighted keyword generally marks the snippet

date). If more than one keyword instance occurs in a snippet, the highlighted keywords are spaced out evenly and may not coincide with the date. (This sort of complication, thrown up by real data, demonstrates the value of using real, rather than idealised data for developing designs).

The text snippets are far more densely packed than in conventional search interfaces for digital archives, and the interface is minimal. The snippets are not initially presented with any additional metadata or other visual distractions. The report's date and London borough can be revealed in a tooltip (a message which appears when the cursor is positioned over a text snippet in the visualisation).

A grey bracket to the left of the snippets (see Figure 49) indicates when multiple snippets from the same document are visualised. (In this prototype, if there are multiple matching keywords within a text, they are all visualised). In addition to a tooltip appearing, hovering over a snippet produces a coloured underline on it; hovering over one of multiple snippets from the same document causes coloured underlines to appear on all these snippets, confirming they are from the same document. To check the wider context of a snippet, clicking on it opens the full report on the Wellcome Library website in a separate browser tab. In the prototype, the user is taken to the report's front page but in future versions, more fluid exploration would be achieved if they were taken to the snippet location within the report.

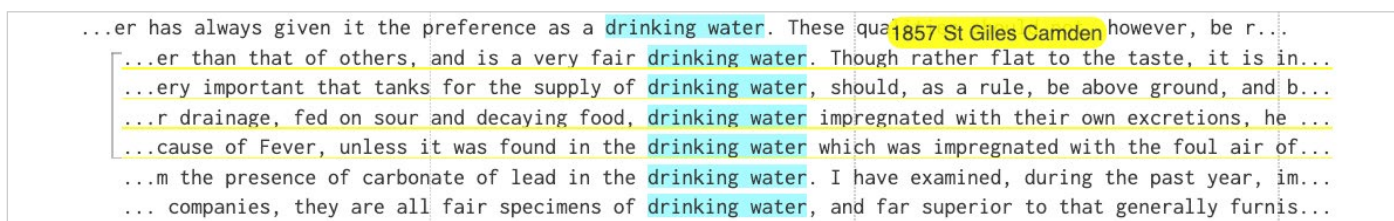


Figure 49. Detail from visualisation of MOH reports data: 'drinking water' showing bracket, tooltip and underline indicating multiple snippets are from the same document.

The snippets generally arrange in a slope, which is steeper for a higher occurrence rate. The focus of this tool is not primarily quantitative (and if a quantitative view of keyword distribution over the collection was the primary aim, this would be better achieved with a line graph, with additional information about the text volume distribution through time). But any kinks, gaps and gradient changes of the slope indicate changing frequencies of occurrence across these reports at particular times. This serves to keep the temporal relationships between visualised snippets in mind when reviewing them.

The raw text includes sections corresponding to tables. While it is not possible to easily isolate or remove the tables, it is often possible to discern in the visualisation if a snippet corresponds to part of a list or table entry from the surrounding string of punctuation or numbers (see Figure 50). The visualisation also highlights reused or repeated text across reports which visually matches (see Figure 51), for instance where a template is used. Though not visible in Figure 51, the tooltip function can be quickly used to confirm all these reports are from the same London Borough.




Figure 50. Detail from visualisation of MOH reports data: 'typhoid' showing snippets corresponding to table entries.




Figure 51. Detail from visualisation of MOH reports data: 'opiates'.

Exploring the MOH reports with Steptext

The Steptext prototype can support tracing commentary through time across texts. It is possible to see the context of what is being said around the keyword at a given time, and to quickly compare contexts through time. The tool is designed for exploration rather than presentation, so in the examples below I prioritise questions prompted by using the visualisation.

Querying the tool with 'blitz', for example, produces a visualisation with distinct shape. There are sparse pre-WWII instances of the term (names, and an instrument for stunning animals). Only in the strong column of snippets from 1940 does 'blitz' refer to the German air raids on Britain starting that year (see Figure 52). Reading through this column of snippets reveals the gradual adoption of the word 'blitz' into accepted language. Although the visualisation cannot explain the word's origins, it shows the first instances of it in the MOH reports occur within quotation marks.



Figure 52. Visualisation of MOH reports data: 'blitz'.

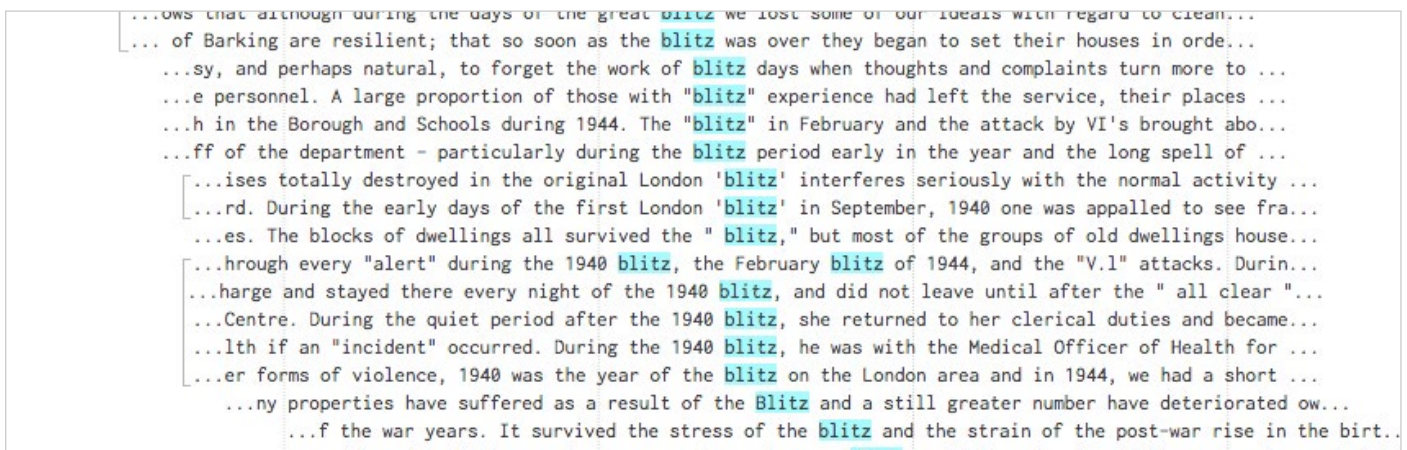


Figure 53. Detail from visualisation of MOH reports data: 'blitz'.

Over time, instances without quotation marks increase and become the norm (see Figure 53). At the bottom of the slope, by the 50s and 60s, the term is even used metaphorically: “what has been described as the ‘bed bug blitz’ took place” (Medical Officer of Health for Kensington Borough, 1957:p.13). Are we seeing the adoption of this new word into the English language?

As another example, searching for the keyword ‘heroin’ features two pre-1940 snippets and then a surge in instances of the term from 1960 onwards, shown by the almost vertical column to the right (see Figure 54). Reading the snippets reveals a disconnect in the language used between these two sides. The pre-1940 occurrences of ‘heroin’ are reporting on the drug, but in the context of regulation and a surgical technique. The post-1960 snippets instead reveal that the narrative has shifted to drug addiction/abuse (see Figure 55).

Is the shape indicating a moral panic around drug use at this time? Possibly. Use of other opioids in the UK pre-dates heroin. Querying ‘laudanum’ (an alcoholic solution containing morphine widely prescribed over the period of this collection (London, 2005:pp.98)), for example, reveals routine recording of death by laudanum poisoning/overdose from the earliest MOH reports (1850s)—but with no mentions of addiction (see Figure 56).



Figure 54. Visualisation of MOH reports data: ‘heroin’.

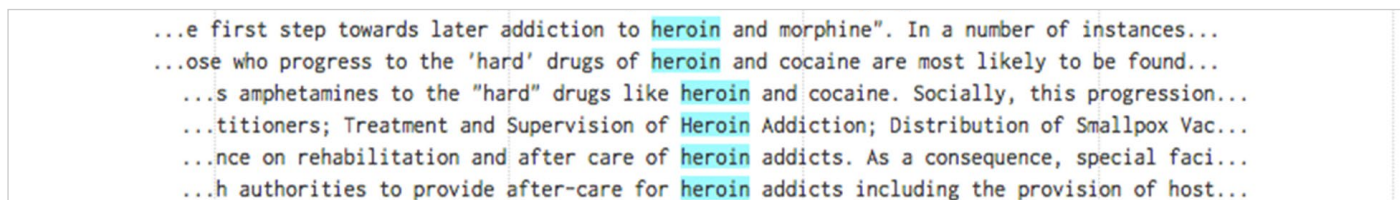


Figure 55. Detail from visualisation of MOH reports data: ‘heroin’.

Query ‘addiction’ or ‘addict’ in the tool and there are no results pre-1934, with most snippets dating from the 1950s onwards. Reading the snippets reveals that at first, addiction is largely discussed in the context of alcohol and tobacco (see Figure 57); from the 1960s the commentary shifts largely to drug addiction (see Figure 58), with some MOHs devoting greater volumes of content in individual reports to discussing drug addiction. The commentary shows increasing concern from authors: “[the] problem of drug addiction has now reached national proportions” (Medical Officer

of Health for Islington Metropolitan Borough, 1967). As always working with historical sources, checking the wider context of the snippets text in the full report helps avoid misinterpretation.

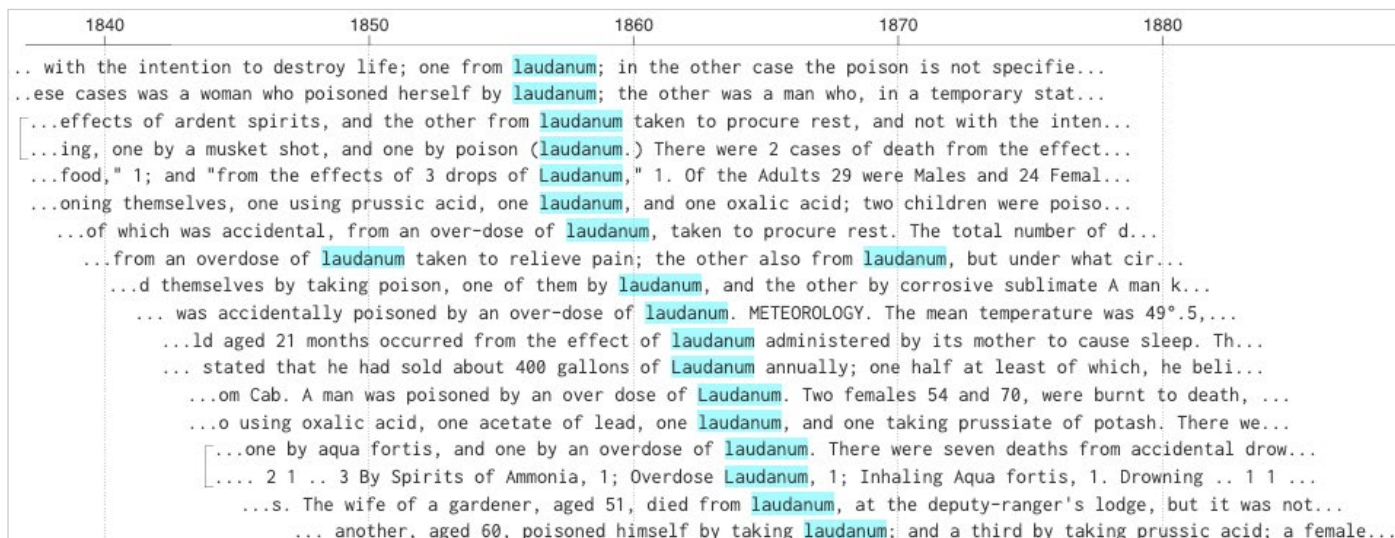


Figure 56. Detail from visualisation of MOH reports data: 'laudanum', 1857-68.

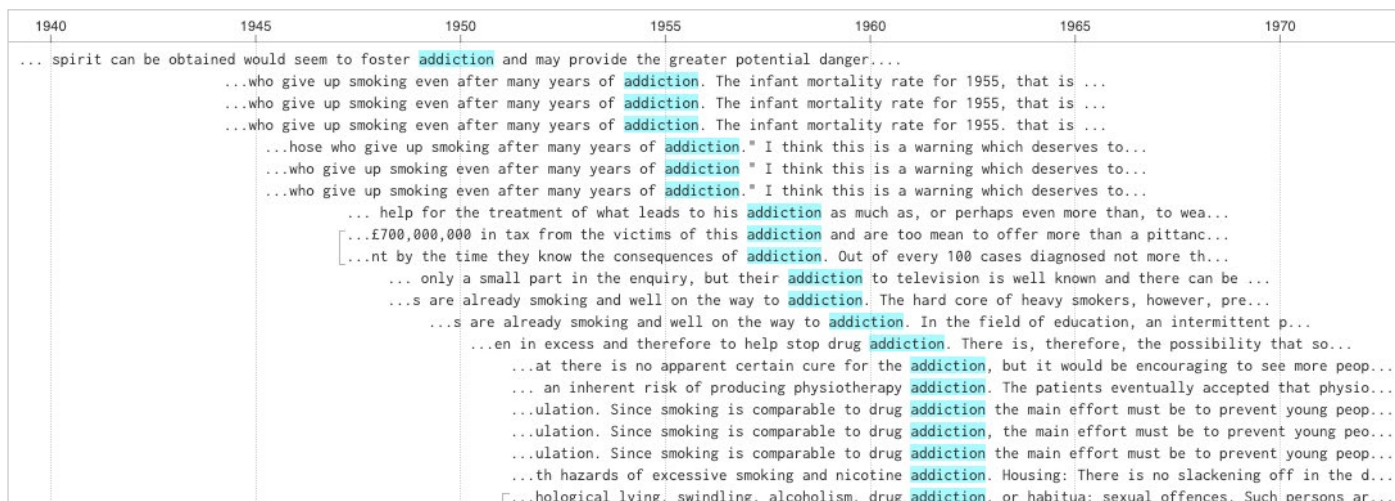


Figure 57. Detail from visualisation of MOH reports data: 'addiction', 1950-62.

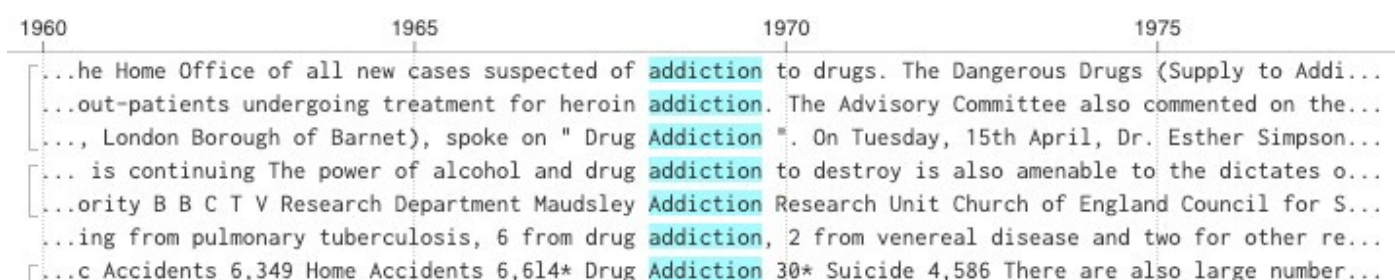


Figure 58. Detail from visualisation of MOH reports data: 'addiction', 1969.

This is a fairly superficial reading of the patterns highlighted by the visualisation, but referring to literature around these topics would suggest they align with wider narratives: rapid government drug policy reconsideration and change starting from the 1950s, though in fact in response to a small increase in illegal drug use (Bennett, 1988:p299); and the emergence of the addiction concept in medicine (London, 2005:pp.99-103). Exploring the MOH reports with Steptext can help highlight these kinds of temporal shifts, though the researcher still has to do the work of interpretation.

I am not suggesting the questions raised by the visualisation discussed above would not have been possible without the tool, rather that by showing the data in this alternative way—an enforced

temporal view of the results, where it is possible to scan many contexts of keyword use in one view—such patterns can be more apparent. The visualisation does not provide unequivocal answers, but it supplies evidence and stimulus for further investigation.

As the prototype is set up, it is less productive when there are large numbers of results and there is little temporal variation (or little apparent temporal variation in the screen view because of results volume). ‘Typhoid’, for example, appears in over 4,000 of the MOH reports, often multiple times within a report and frequently as part of tables. Over 21,000 snippets for ‘typhoid’ are generated in total (there are 461 snippets generated for ‘typhoid’ in the reports from 1896 alone). Without options to reduce snippet numbers by, say, just showing the first matching snippet from a text or options to filter results by additional query terms or London borough, the resulting visualisation is overwhelming. And within the screen view, the snippets simply become a chronological list.



Figure 59. (left) Detail from visualisation of MOH reports data: ‘typhoid’, showing that when there are many hits the snippets become a list.

Figure 60. (right) Zoomed out view of visualisation of MOH reports data: ‘typhoid’.

Feedback sessions

To capture impressions of this timeline tool design from researchers (and explore how it can better serve them), I conducted a series of semi-structured interviews with historians. These feedback sessions were an opportunity to capture impressions of the visualisation design and solicit suggestions for improvements and refinements.

I conducted interviews with seven historians: each interview was conducted in person and was roughly 30 minutes in duration. Interviews were audio-recorded and later transcribed. The participant group was a mixture of historians with an interest in, or familiarity with, this collection of historical texts, and digital historians (researchers with expertise in using computational methods, such as text mining or visualisation, in their research). The interviewees were a mixture of ages, genders, positions of seniority, and degree of familiarity and comfort using digital tools or coding. Participants were a professor, three lecturers, two freelance scholars, and a PhD student. They had different institutional affiliations and two of these individuals also freelanced as curators.

Some had a specialism in history of medicine or pharmacy, and they all generally researched time periods post-1600. One interviewee specialised in historical geography. They were recruited through contacts at the Wellcome Library and the Institute of Historical Research.

At each interview, the scholar was introduced to the digitised texts in the current Wellcome Library web interface, useful as the current interface design is a very typical contemporary example for digital collections. Keyword search is available across all the reports or within a single document in the current web interface. This was demonstrated to all interviewees at the beginning of the interview.

I demonstrated the Steptext prototype with a few prepared example queries generating visualisations, and briefly explained the goals behind the design. I then invited interviewees to use the tool freely themselves. The historians often used the interface to search for keywords relating to their current research topics, or anticipated keywords they predicted might produce interesting results. I also prepared some further example keywords to suggest for the occasions when an interviewee struggled to come up with their own.

As outlined in my 'Methods' chapter, rather than using task-based protocols, the interview structure included broad questions about their impression of the tool, and questions about the value of specific features and choices made in the design. Interviewees were asked to discuss how they might use the tool in their own research.

Discussion

While the tool was demonstrated on the MOH reports, interviewees generally discussed its potential for exploring text data of their own choosing, including mixing digitised sources from different collections.

Value and role of the visualisation design in the historical research process

I wish that I'd had this...this would have been massively useful [for my PhD research].

Interviewees described the various activities and "constant iterative process" involved in exploring texts at the beginning of the research project: surveying texts from across many archives and libraries, organising and filtering down what is relevant, looking at new material. The tool could play a role supporting some of these activities.

Interviewees described how the tool would support checking to see what keyword matches there are in the documents, what the commentary around them suggests, if there are temporal patterns in use and occurrence, and seeing the temporal range of results. Further, this would support filtering and identifying relevant source material to look at in more detail.

The visualisation design was described as a compression of steps commonly taken with standard interfaces to digitised archives. One interviewee described it as like being able to "look through all of them [the reports]" at the same time. Indeed, one interviewee described their frustration with making this sort of survey with current interfaces: "I really like it. It's a really useful way of seeing information together...because of how it [current standard user interface for searching over all the documents] brings you to the documents you're unable to even get a sense of how the terms are used until you are in the search [separate browser window for an individual document] and then you have to go back".

One historian offered a rich, step-by-step view into how they expected this tool would help them in their current research project, which focuses on lodging houses in the 18th and 19th Centuries:

I would use it [Steptext] to highlight and then look at in more detail the attitudes to lodgers and whether they...needed regulation,..I would use it by putting in the term 'lodger'... firstly I'd be interested in the patterns, I'd want to see where [are] the most tallies...I'd use it numerically, so I could say this legislation came in in 1855 and they're all reporting: are they saying the same thing?...does what they say about lodgings proceed actual legislation... you know, who's driving it? Is the evidence from these reporters...being used to put legislation in place?

“But I'd also be interested in the descriptions...this would give me a body of material which I could then play with to find...text which would give me a qualitative feel for what people were saying about lodgers, how they were perceiving them and what were the concerns, and the descriptions of actual conditions are interesting as well.

The visualisation was described as providing structure to the process of surveying sources, which can feel “a very disjointed process sometimes”. This argument—for the unifying and structuring effect of chronology—has been made historically: “[without chronology,] the Series of Time... is really nothing more, than a Bundle of Detached Fragments” (Blair, 1754 as referenced by Boyd Davis & Kräutli, 2014).

Interviewees suggested the visualisation shape draws attention to temporal patterns that might not otherwise be noticed and provides a different view on the material. It also serves to stress temporal engagement with the material; as one interviewee described it “the in-your-face chronology is...really valuable”. The following quote captures the enthusiasm of one interviewee, as the visualisation shape for the keyword 'lascar' (a term for south-Asian seamen employed on European ships until the mid-20th century) aligned with their knowledge of framing events:

Look at how much more they talk about it [lascars] in the thirties. Cool! This makes sense right. It's just so cool to see it laid out like this, because right around 1900 they are experiencing plague in India, so sailors from that region become of much greater interest to people because before...dock officers were not particularly concerned about welfare of lascars. But then it becomes a threat. You get one reference before the 1900s and then it's like 'bam!' [a column in the visualisation shape]

Importance of trust in the interface and control in the search process

Trust in the interface was an important concern for users. This topic came up in interviews when discussing strategies to manage scale in the visualisation.

Depending on the keyword searched for, the volume of results in the visualisation can render the display overwhelming. During the interviews, I showed historians different versions of the tool where either a filtering had been implemented to limit visualised results to the top 100 documents ranked by 'relevance' (I implemented the standard relevance scoring algorithm used in Elasticsearch, TF/IDF: term frequency/inverse document frequency) or the visualisation displayed every single matched instance across all documents. I also showed interviewees an iteration of the tool in which a user is able to filter results by occurrences of nearby secondary keywords (see Figure 61). Eager that the text snippets always be legible in this visualisation, I did not explore zooming techniques as the snippets quickly become illegible below a certain size.

Opinion diverged amongst the historians interviewed about whether it would be appropriate to implement a relevancy ranking (even before considering if this ranking algorithm would be the most appropriate here). Some felt it was valuable to avoid users being put off or being unable to

make sense of too many snippets. Others, however, felt a general algorithm could never anticipate the specificity of their search and would be unhelpful. There was agreement it was very important to be explicit about removing any results if it has been implemented, and there ought always to be an option to see all the results if a user wishes it. The general opinion was that showing everything and then offering the user an option to filter results on London borough, time-span and further keywords was preferred. Being able to limit visualised results to one snippet per document would be a small help with this issue.

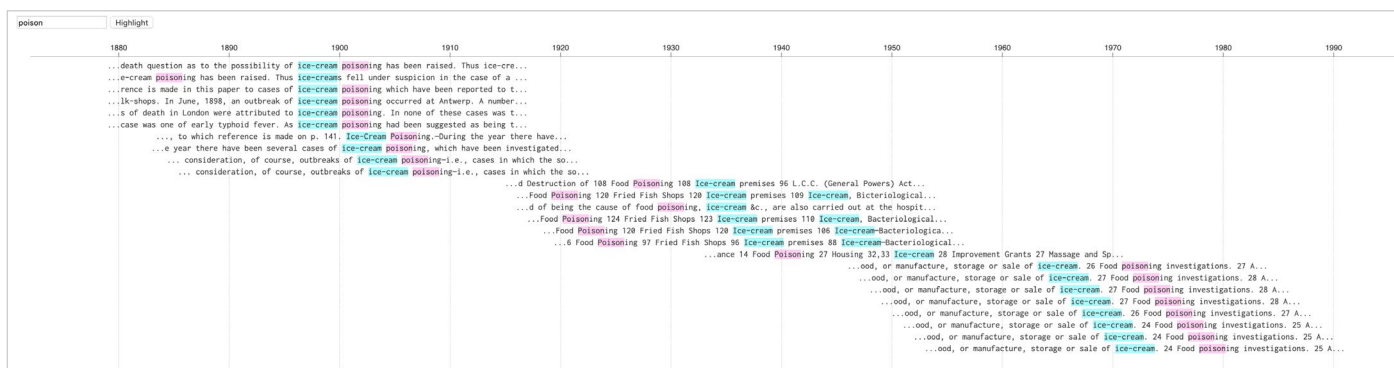


Figure 61. Prototype visualisation of MOH reports data: ‘ice cream’ & ‘poison’.

Trust and transparency in the processes by which results are returned was very important for these users: “when it come to academic research ... what I conclude from your tool feeds into my reputation”. A number of interviewees were uncomfortable with the use of a ranking algorithm because they considered it a black box. This concurs with research in search interface preferences for historians, demonstrating they value control in aspects of searching and browsing (Crymble, 2016). One historian commented that the simple aesthetics of the visualisation appealed to them also for reasons of trust: “I really like the simplicity of your visualisation actually. You haven’t tried to be flashy and sparkly; it’s just ‘this is the information’”.

Mayr et al. (2019) define trust in data visualisation as “the user’s *implicit or explicit tendency to rely on a visualization and to build on the information displayed*” (p.25). They argue that trust in this context depends on the trustworthiness of the underlying data and of the visualisation design, and also the user’s trust perception. In the case of the MOH reports, the data could be said to be trustworthy in the sense that the OCR text transcriptions are very high quality. But, on the other hand, a researcher does not read these texts strictly at face value. It is necessary for a scholar to understand the larger context of the text collection: who authored these sources, why and what may be missing? This is something this visualisation does not help with. An argument could be made that the uniform and minimal visual design of the text snippets displayed as one unit perhaps encourages a user to overlook that these are different documents created in different circumstances by different authors with different motivations and values.

Qualitative engagement with cultural data

I also showed interviewees a mock-up where sections of page images, corresponding to text snippets, had been integrated into the design (see Figure 62). I was interested in what kind of close engagement with the texts was desirable for this audience.

Interviewees suggested this alternate view is not valuable (“I’d be looking for the information rather than what the source actually looks like”, “that doesn’t add anything”) especially as in this case you can get the full page images on the Wellcome Library website anyway. It was even described as actively unhelpful as the text takes longer to scan and less text is visible in the same space. It was suggested showing page images would be more relevant for texts containing images, or if the OCR

is known to be unreliable. For a different audience however—say for an exhibition— page image sections may be more appropriate to help establish the character of the source and reinforce that the text comes from a real, printed object. Exactly what kind of close engagement is desirable for cultural data in visualisation is highly contingent on the user and the use.

The image shows a search interface for 'inoculate' with 27 out of 5388 document matches. Below is a table representing the search results, with the word 'inoculate' highlighted in yellow in the original image.

docid	content	borough	year
10	inoculate these same people with Small Pox, and he found th	Acton	1901
1025	equipment had been sent to 52 banks, insurance houses and commercial in order to inoculate staffs at their places of employment.	CityofLondon	1958
530	practitioners so that they may inoculate children whose practitioners wish their own doctor to give the injection, the	Bexley	1937
954	mission to inoculate some sheep, and he observed	CityofLondon	1864
2367	The protection takes some time to develop, so the practice is to skin-test all 13 year-old pupils, whose parents consent and to inoculate those found to have no	Heston	1957
2607	test certain children who had previously been inoculated against Scarlet Fever and Diphtheria, and also to inoculate new admissions, for whom	Ilford	1942
2241	vaccine available to inoculate some nine million school children in the inoculate school children of the appropriate ages. In the other	Harrow	1956
5218	There was not enough vaccine in May and June to inoculate all these children, and the Minister	Willesden	1956
2609	and Dick testing of certain children who had previously been inoculated against Scarlet Fever and Diphtheria, and also to inoculate new admissions,	Ilford	1944
4437	to be invaded by rats, and by the fleas which inoculate the	StokeNewingtonMetropolitanBorough	1910
3193	present in every membrane of diphtheria, and in identifying it ally present there," and that Dr. Klein proceeded to inoculate	LondonCountyCouncil	1892
2608	testing of certain children who had previously been inoculated against Scarlet Fever and Diphtheria, and also to inoculate new admissions, for whom consent forms had already been obtained. In addition	Ilford	1943
5045	The plan is to inoculate a large number of children and there-	Wembley	1946
4724	Research Council, sufficient to inoculate all children who are under	Tottenham	1950
3720	Owing to the imminent sailing of the vessel the Ship's surgeon had obtained a supply of (Anti-typhoid vaccine) and proposed to inoculate the whole of the crew and as many as possible of	PortofLondon	1957

Figure 62. Mockup visualisation of MOH reports data including page photo crops: ‘inoculate’.

Suggested features

Mirroring findings in the Trading Consequences project (Hinrichs et al., 2015), interviewees recommended that integrating access to full report texts within a single interface would aid more fluid exploration. This would also help give a sense of where text snippets appear in the narrative flow of the wider report. Integrating access to the full texts would likely be even more important for other kinds of texts, such as newspapers, which are chunked in many different sections.

As expected, interviewees expressed a wish for more control/sophisticated keyword search system: functionality for searching for multiple keywords, and approximate keyword matching (include matches where the keyword has possibly been misspelt/OCR error). Some also wished for some way to output—to save, or be able to print—visualisations for queries.

Beyond keyword search

When a historian is tracing commentary through texts, they are really interested in a concept rather than a keyword or set of keywords. Keyword search interfaces will miss instances in the text relevant to that concept where different vocabulary is used. I discussed topic modelling as a means to reveal thematic structure in text, and its critiques, earlier. But there are other ways to try and get beyond the keyword. Suggestions from interviewees including: using an historical thesaurus to create keyword groups to search over, keyword stemming in search (enable searching for a group

of related words that share a common base form, like ‘nurse’ ‘nursing’ ‘nursed’), or potentially incorporating semantic analysis of text (Alexander et al. 2015; Thompson et al. 2016). However, there is a tension here between sophistication of the system and transparency of the results.

Note on evaluation session

Occasionally in these interviews, historians would evaluate the tool by making comparisons with researching using physical documents. All interviews began by demonstrating the currently available, standard web search interface for these digitised texts and all participants had used digital archives previously. Nevertheless comparisons, particularly about speed and scale of search, were sometimes made against researching with physical texts. This tendency is something to be aware of in the future for interpreting interview statements from historians in this context.

Summary

This chapter introduced ‘Steptext’: a timeline visualisation tool for exploring texts. Designed around the Medical Officer of Health reports collection, with minimal cataloguing metadata, this tool visualises snippets from the texts themselves—allowing researchers to direct their own exploration while closely engaging with what is said in the texts. The design combines “in-your-face” chronology with glimpsing, in situ, the context of keyword use—highlighting patterns in commentary across the reports. Interviews with historians raised the issue of trustworthiness, and the value of simple visuals to promote transparency in cultural visualisation design. These interviews also highlighted that the design space for ‘close reading’ is not without options, and that different uses and users may require substantially different interface designs. Tracing keywords is not a perfect surrogate for tracing concepts through texts, but it is transparent and interpretable.

‘Tag-timeline’

Cooper Hewitt Smithsonian Design Museum




Figure 63. ‘Tag-timeline’ visualisation of Cooper Hewitt collection data: ‘Coffee and tea drinking’

Following the MOH reports project, I was awarded a fellowship at the Cooper Hewitt Smithsonian Design Museum, New York City. Cooper Hewitt is a museum of historic and contemporary design, with a collection of more than 200,000 design objects (Cooper Hewitt, 2019a), digitised and made available online through their collection website (2019b) and an API (2019c). The museum collection is diverse, covering objects of “product design, decorative arts, works on paper, graphic design, textiles, wallcoverings, and digital materials” (Cooper Hewitt, 2019a) spanning from circa 1100 BC (Cooper Hewitt, 2014; Cooper Hewitt 2019g) to the present.

Cooper Hewitt’s permanent collection was initially, in its first realisation in the late 19th Century, made up of decorative arts objects collected throughout the US and Europe to serve as a teaching resource for design students (Smithsonian Institution Archives, 2019; Cooper Hewitt, 2016). The institutional mission and collecting focus is now more expansive though, embracing all the disciplines of contemporary design (Cooper Hewitt, 2019a).

I travelled to New York City for this fellowship where I spent 5 months internally at the museum, placed with the Digital and Publishing teams. During the fellowship I enjoyed a freedom to direct my own work and there was no requirement to deliver a working product. I had an internal supervisor to discuss developing ideas and ask questions of the collection, and who could help me make other contacts within the museum.

Tagging in the Cooper Hewitt collection

I began this project by familiarising myself with the collection data and trialling the various methods available via the API. My interest was caught by tags. Tags are keywords that “identify and categorize” (Trant, 2009) records, typically in an informal way. Inspired by web-based resources, ‘tagging’ has gained increased interest in the cultural heritage sector as a way to open

up access to online collections. Typically tags use non-specialist language, and are seen as a way to close the ‘semantic gap’ between cultural heritage professionals and the public who may use different terminology (Trant, 2009). Several institutions now invite ‘social tagging’ of their collections data (Cairns, 2013) where members of the public can contribute tagging data towards a ‘folksonomy’ or folk-derived taxonomy (Vander Wal, 2007).

At Cooper Hewitt, tags are applied internally at the museum (typically by curatorial staff or Masters students/interns), rather than by the public. Brenner and Adang (2016) describe the strategy used at Cooper Hewitt: staff choose tags relating to ‘motif’ (describing the “formal visual qualities of an object”) and ‘user’ (relating to “the object’s applications and intended audience”). The intention is to avoid replicating information already present in the museum collection data (Brenner & Adang, 2016). The tags support searching/browsing and drive recommendation systems on the museum collection website; they are also used to power interactives in the galleries.

Examples of tags in the Cooper Hewitt collection data include: ‘birds’, ‘black and white’, ‘symmetry’, ‘overlapping’, ‘coffee and tea drinking’, and ‘communication’. There is no controlled vocabulary. Tagging is a work in progress at Cooper Hewitt and, at the time of my fellowship, approximately 9,000 objects had been tagged (about 5% of the collection). Typically objects are tagged when they are included in an exhibition or if they are highlighted for some reason, for example in an article on the museum website. The number of tags an object has, if any, ranges from 1 to around 20 (see Figure 64 and Figure 65 for examples). Only a small proportion of tags had more than 100 objects associated with them: 76 of the 3199 tags (2%), calculated 25th July 2017.




Figure 64. Poster, Miniwatt/Philips Radio; Client: Philips; lithograph on paper. Tags: *graphic design, electronics, commercial poster.* (20/03/19). © Cooper Hewitt, Smithsonian Design Museum.



Figure 65. Pair Of Slippers (USA), 1889; soldered tin. Tags: *interior, decoration, personal, display, gifts, women, commemorative, ribbons, figurative, walking, decorative, celebration, metal, bows, pair, fashion accessory, fashion, metalwork, shoes, footwear, buckles.* (19/03/19). © Cooper Hewitt, Smithsonian Design Museum.

I was interested by tagging in the Cooper Hewitt collection data because tags represent themes/connections (sometimes abstract) through the data, of the type often not explicit in cultural data. They connect diverse objects of different mediums from across the museum’s departments, often over large time spans. Unlike grouping objects by type or artistic period, the resulting groups are visually cohesive yet diverse. For individual objects, the tags provide a variety of perspectives highlighting different characteristics, and situating each object in a multi-dimensional set of contexts. The tags are also highly subjective. In addition to reflecting individual curators’ preferences, the choice of tags often reflects past exhibition themes. I was interested to find out: how I might design a timeline visualisation appropriate for the subjective and visual qualities of tagging, what timeline visualisation of the tagging data would show, what such timelines would support for different audiences (expert, casual), and how the work-in-progress status of the tagging data would impact its visualisation.

Visualising image data by themes

As the tags are often visually motivated, my initial experiments centred around mapping the images themselves within a visualisation. (Approximately 75% of objects in the collection have photographs available online, as of March 2019). I discussed earlier the power of plotting images as data points in ‘Cultural timeline visualisations—Survey’.

Chronological arrangements of cultural collections almost always produce clustering effects, because the collections are rarely uniform over time. Closeness in time leads to clustering. Other kinds of ‘proximity’ than date can be used and several prior projects have explored visualising image data from cultural collections, clustered based on a computation of ‘similarity’ which may take many forms: in both 2D (Nasjonalmuseet 2017a; 2017b; V&A Museum, 2010) and 3D environments (Diagne, Barradeau & Doury, 2018; Diagne & Barradeau, 2017; Diagne & Doury, 2017; Bernhard, 2016). In my work here, I used tags as my chosen means of discovering similarities.

Previous projects have visualised curated, thematic connections between images in cultural data as a way to support browsing and discovery in collections. While not mapping items by time, the Städel Museum (Janßen, 2015) and Royal Academy (Huckle, 2018) have trialled interfaces with curved connecting lines drawn to link artwork images sharing, for example, (hand-curated) visual or emotional characteristics.

Visualisation of collection images by themes and time has been explored in ‘Past Visions’ (Urban Complexity Lab, 2016; Glinka, Pietsch & Dörk, 2017) where data can be filtered using curated thematic keywords. In the in-gallery context, Cleveland Museum Gallery One (Alexander, Barton & Goeser, 2013) and ‘mArchive’ (Kenderdine & Hart, 2014) offer thematic display of collection data by time.

Date information in the Cooper Hewitt collection data

At the time of my fellowship, most post-1800 objects in Cooper Hewitt’s data had a date/date span expressed numerically. Pre-1800 objects, however, often only had date information as it would appear on a gallery label (in free-text), for example ‘Created before 1870s’, ‘late 19th–early 20th century’, ‘ca. 1850’ or ‘2012–present’. (I have discussed the difficulty and subtlety of date uncertainty in my Background chapter).

In order to be able to process this data computationally, I parsed the date text to numbers using the `yearrange` library (Resig, 2014) which is written for working with museum date language. This library works by converting, for example, ‘late 18th century’ to ‘start: 1775, end: 1799’. For prototyping, this seemed sufficient, though a few cases required special consideration:

- `yearrange` does not adjust date spans for objects dated ‘circa/ca./c. X’. In these cases, I chose to extend the date span by ± 10 years around X on advice from the museum curators
- For objects dated ‘before X’, I chose to use the date X as it was not possible to generalise in this case
- `yearrange` does not parse BC dates. Since BC-dated objects are a very small proportion of the collection, I excluded these objects in my prototyping.

Visualisation design and prototyping

I started by experimenting with laying out images matching a tag horizontally by date, but was immediately faced with the issue that many objects’ dates were a (sometimes wide) span. To get round

this common issue, cultural visualisations often plot items at the beginning or median of their span. Kräutli developed the ‘temporal jittering’ algorithm (2016b:p.182) as an alternative, organising data points within a grid structure. Most of the pre-1800 objects in Cooper Hewitt’s collection have an imprecise date, so I was interested to see if I could find an approach that would avoid misleading shapes forming—for example, a cluster at 1750 for all items dated 18th Century—but was quick and simple to implement, to advance my prototyping. The approach I settled on is to plot images at a random position within their date span (image’s centre as origin). To spread the images out vertically, so more are visible, I also arranged the images vertically at a random position within bounds, see Figure 66.




Figure 66. Prototype visualisation items tagged ‘triangles’ in the Cooper Hewitt Smithsonian Design Museum collection horizontally by date.

The result more resembles a collage than a graph, encouraging the user to look rather than count. How might I hold on to this aesthetic, but ensure all the images can be inspected?

Image layout algorithm

Plotting a distribution along a single dimension, while emphasising individual points, can be achieved with a ‘beeswarm plot’ (Eklund, 2016), where points are closely-packed but not overlapping. Resembling beeswarm plots, data points represented by circles are arranged by time in Fathom’s ‘Porfiry’ (2019) or Kräutli’s ‘Britten’s Poets’ (2016a). Plotting images, ‘revisit’ by Stepaner (2010) maps tweets around a specific topic through time, distorting the time scale to highlight Twitter message threads.

I was interested to see if I could tailor a similar approach to uncertain cultural date information, keeping the collage-like appearance of the visualisation. I developed a layout algorithm taking the random positioning I had employed up to this point as a start. This layout algorithm aims to separate images so that they are not overlapping but still fairly closely packed, keeping images within their date spans. The algorithm steps are:

1. All images are positioned:
 - horizontally by their date (if the date is a span, the image is placed at a random position within the span), image’s centre as origin
 - at a random vertical position within a set range
2. For each image, the program checks if it overlaps with any other images. If it does overlap, both images are moved apart vertically and horizontally by a fixed distance in the opposite direction to the overlap. To ensure images do not stray too far from their dates, images will only be moved horizontally if some part of the image remains within its horizontal date span (see Figure 67).
3. Step 2 is repeated until there are no overlaps

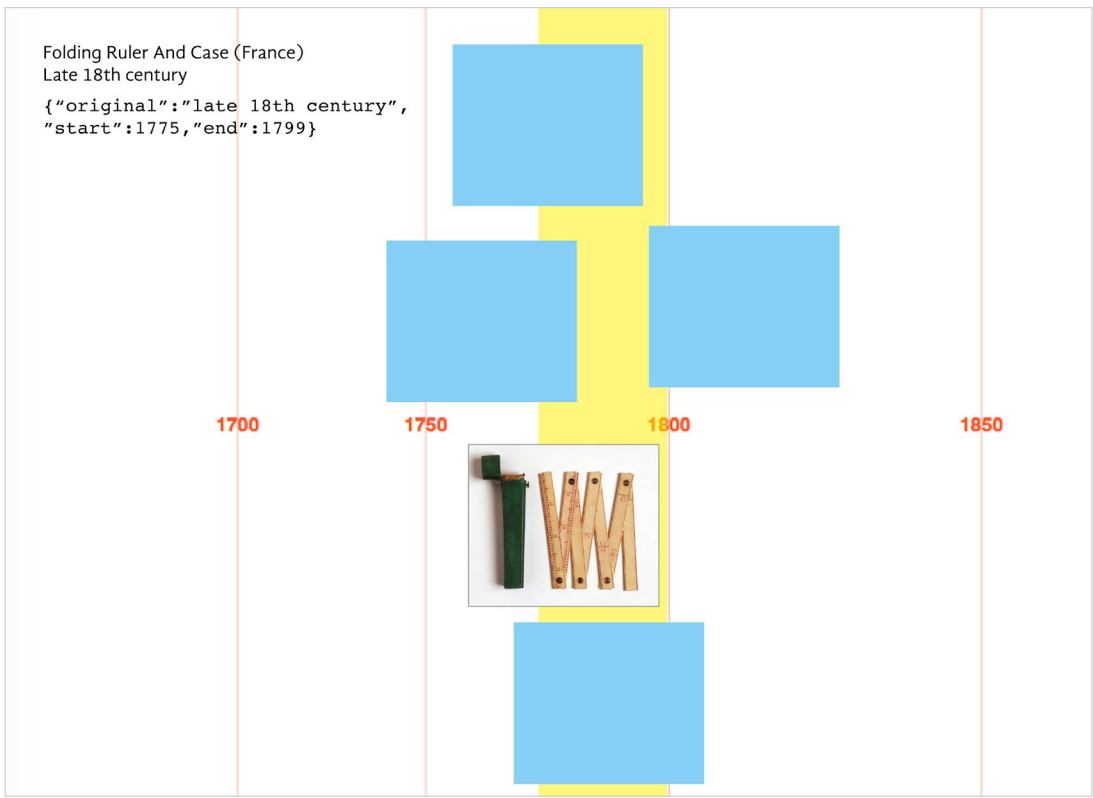


Figure 67. Visualisation annotated to show possible alternative positions (blue rectangles) for an example object, with date span shown in yellow.

The result (see Figure 68) achieves my goal: encouraging users to observe and connect the visual details in the images, which can now be inspected, over quantitative analysis. Using the images as large data markers also prevents reading off precise dates.

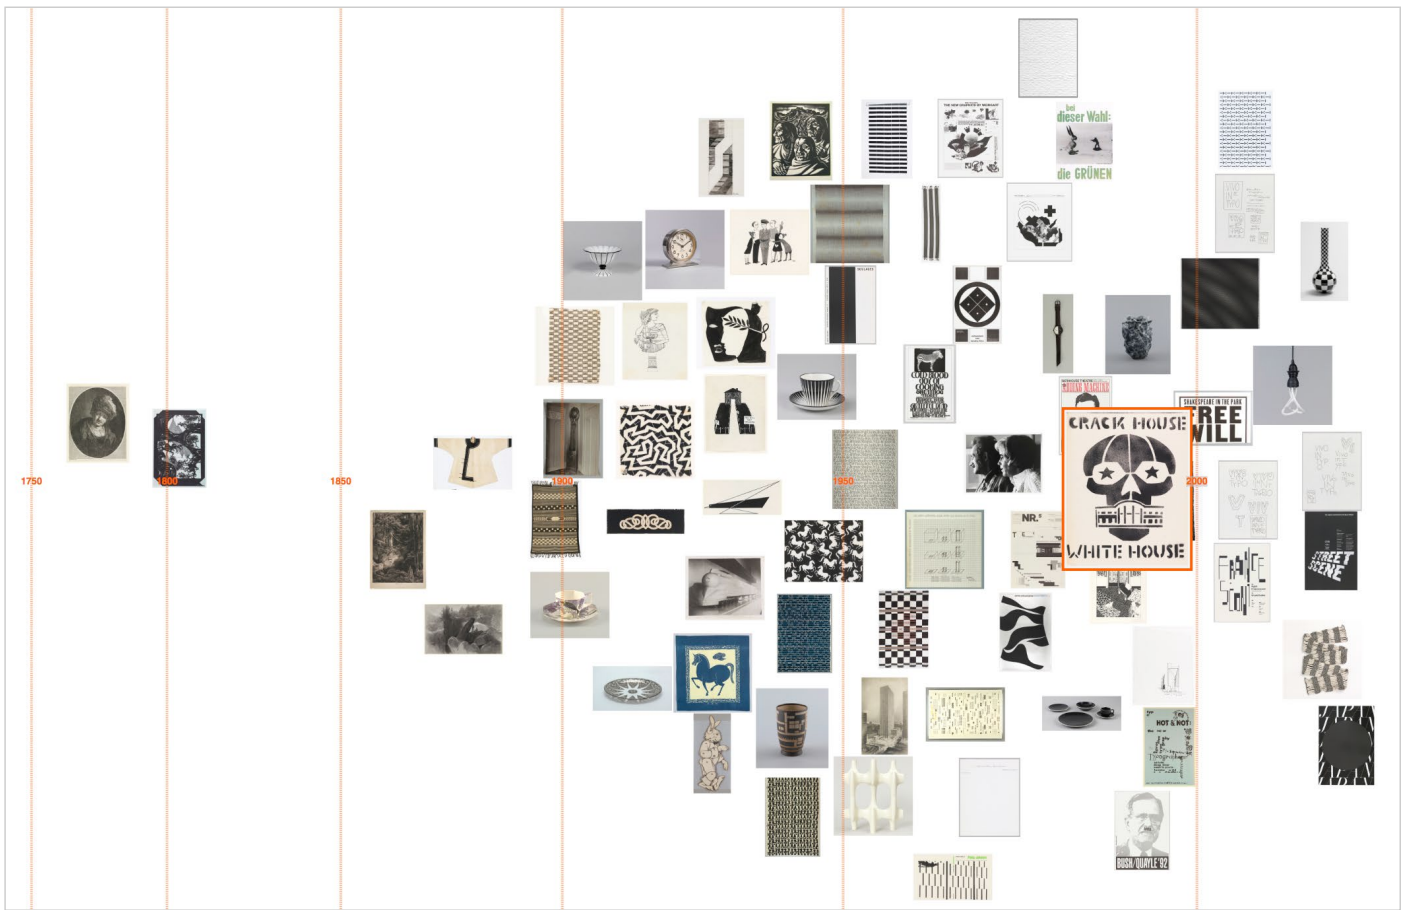


Figure 68. Visualisation of Cooper Hewitt collection data: ‘Black and white’.

Running the program on different occasions can lead to images arranged in a different order in time. Since my aim was for users to observe visual relationships rather than read off exact dates, I was not worried about exact placement horizontally. This does mean however that the timeline cannot be used to make reliable deductions about influence and causality, since a later design could appear to the left of an earlier one. When I discussed this with curators, precision and inconsistent arrangement of the items did not seem much of a concern: “our general approach is to see new connections between things, and as I said, the randomness is [not a sticking point]... as long as it’s correct to date” (Cooper Hewitt curator, 2017c).

Running this layout algorithm for large numbers of images (100+) is computationally expensive and slows the program. The prototype, though, worked fast enough to test with users, which was sufficient for my purposes. Within the scope of this project, therefore, I have not explored efficiency improvements.

Time scale

Since my intention was for users to observe trends rather than count frequencies, I adjusted the design to more evenly spread out the images. The Cooper Hewitt collection has many more objects dating post-1800 than before. So I applied a power function (with exponent 10) to the time scale (D3 API, 2019b). (The problem of data density vs. the advantages of a uniform linear timescale is a perennial one in historical visualisation (Boyd Davis, Bevan & Kudikov, 2010)).

Some tags only apply within a certain time domain, for example the earliest object tagged ‘electronic’ is from ‘ca. 1934’. Again, to even out the spread of images, I coded the prototype so that the timeline date extent would redraw every time a new tag was visualised; the timeline extent is ± 50 year padding around the lowest and highest dates for objects matching a tag, with an upper threshold of the current year on the date maximum. This again demonstrates the value of prototyping with real values where these kinds of issues, which may not have been preempted, come up.

Direction of time

I also trialled a vertical direction for time in visualisation prototyping (see Figure 69). In dialogue with the museum curators, however, I was almost exclusively told horizontal was preferred: “it’s more panoramic” (Cooper Hewitt curator, 2017a), although I suspect this was influenced by showing the prototype on a laptop with a greater horizontal screen width than height. (Boyd Davis (2012) has previously written about the question of time’s direction and orientation).




Figure 69. Visualisation of Cooper Hewitt collection data, time running vertically.

Selecting tags, selecting objects

In the prototype, tags can be chosen to visualise by 2 methods:

1. selection box with list of tags (with autocomplete on typing, driven by Chosen (2011), see Figure 70).
2. selecting an object within the timeline shown, which brings up a list of tags applied to that object for selection (see Figure 71). Visualising new timelines this way serves to display individual objects in different thematic contexts.

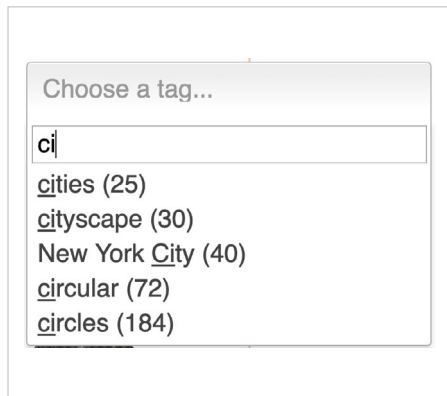


Figure 70. Visualisation of Cooper Hewitt collection data, detail: selection box.

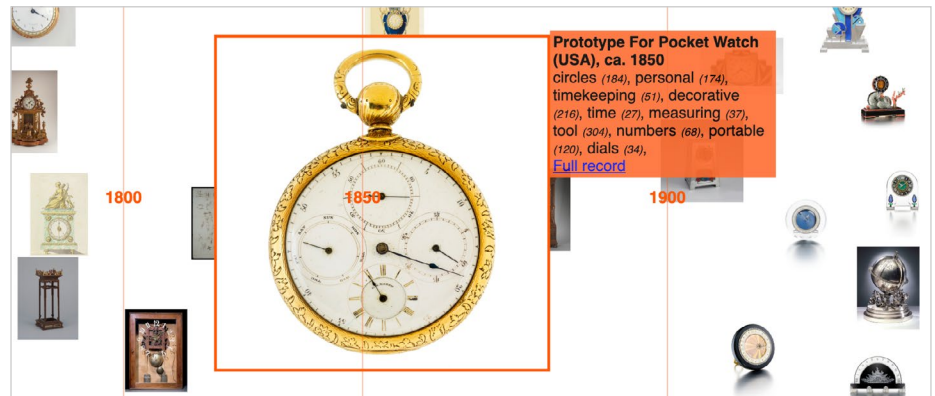


Figure 71. Visualisation of Cooper Hewitt collection data: 'Timekeeping', detail: selected object with tag information.

Images can be inspected at a larger size by selecting them. This also brings up an information box with the object's title, date as free-text, tags, and a hyperlink to its full record on the collection website.

Scale issues

Many images

While the number of objects for a tag was usually below 100, in some cases it is higher. Without reducing the image size, visualising many images is visually overwhelming (Figure 72 shows 300 images visualised), and unless you significantly expand the length of the timeline, there is a limit on the number of images that fit vertically within a time span.



Figure 72. Visualisation of Cooper Hewitt collection data: 'Flowers' with 300 images.

More than 100 images for a tag was the exception rather than the rule for this dataset. I therefore decided to place a threshold of 100 on the number of images visualised in the prototype, intending to explore the issue further with users. (I limited the number of images returned by the API, though its documentation does not make clear the ordering by which images are returned. Such lack of clarity in the documentation of software is unfortunately common, and has obvious implications for the transparency and trustworthiness of tools built using algorithms whose operations are obscure).

This issue will be more pressing in the future, though, as more objects are tagged (unless changes are made to how tags are chosen and applied—an issue I will return to in discussion at the end of this chapter). In future work, it may be worth exploring visualisation strategies for managing scale that allow users to move between focused and contextual views of the data—for instance, zooming or fisheye views (Sarkar & Brown, 1992)—(Cockburn, Karlson & Bederson, 2008).

Few images

To avoid a sparse-looking display, I coded the prototype to scale up the image size (in inverse proportion to the image number) if the total number was below 50 (see Figure 73). I also limited the tags available in the interface to those with at least 10 objects.

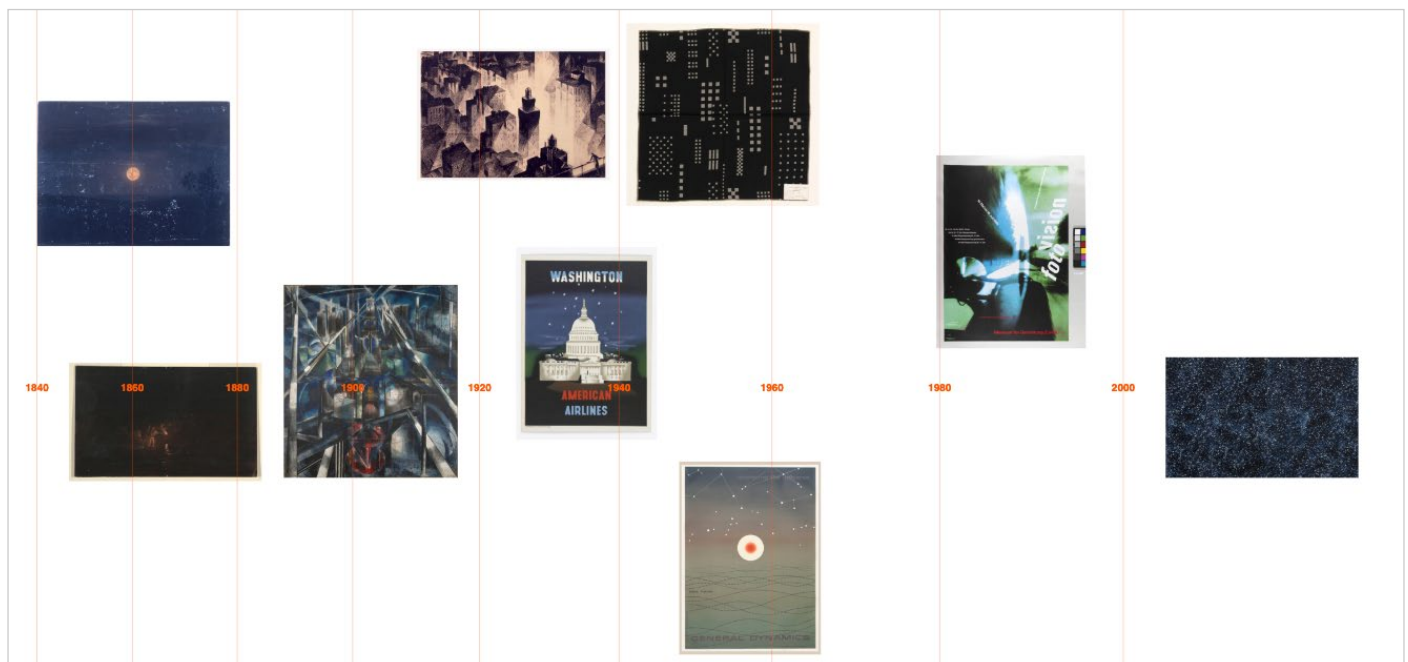


Figure 73. Visualisation of Cooper Hewitt collection data: ‘Night’ with 9 images (below 10 as not all objects in the data had an image). Showing images scaled up.

What do the tag-timelines show?

Images are irregularly clustered, ordered horizontally along a time axis. Overall, they read as a visual ensemble, but each image can also be viewed individually: visual comparisons and connections can be made.

Temporal patterns

By displaying images ordered temporally, the tag-timelines shows change, constancy, and contrast through time. Where a tag corresponds to an object type, the tag-timeline often shows evolution through time. The timelines generated for ‘Coffee and tea drinking’ (see Figure 74) and ‘Chairs’ (see Figure 75), for example, visualise designs in the collection stretching over several centuries to the present. A great diversity of forms and styles are shown, elaborate and minimal.




Figure 74. Visualisation of Cooper Hewitt collection data: 'Coffee and tea drinking'.

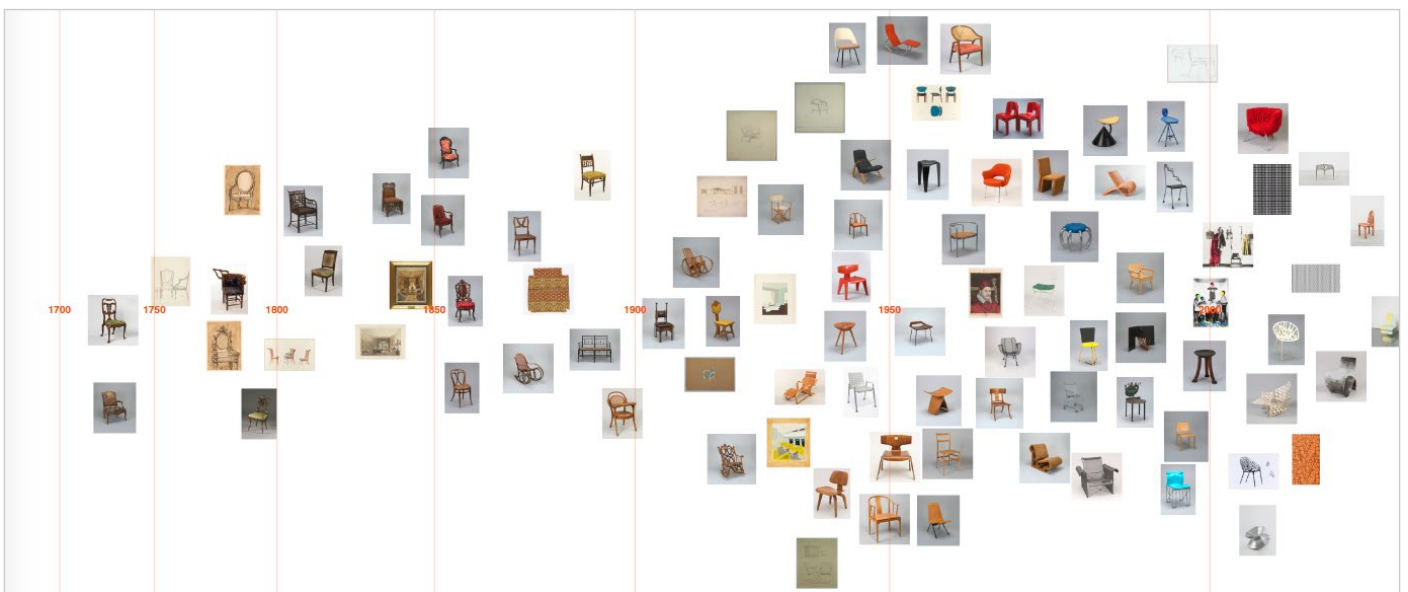


Figure 75. Visualisation of Cooper Hewitt collection data: 'Chairs'.

Tag-timelines can chart the use of a visual motif, pattern or form over time, for instance 'Checkerboard' (see Figure 76) or 'Bent' (see Figure 77). Different visual expressions of and mediums for the same visual ideas can be seen through time.




Figure 76. Visualisation of Cooper Hewitt collection data: 'Checkerboard'.




Figure 77. Visualisation of Cooper Hewitt collection data: ‘Bent’.

In some cases, the timeline displays constancy—for instance, the timeline for ‘Earrings’ shows similar forms appearing from the 19th-Century to the present (see Figure 78).



Figure 78. Visualisation of Cooper Hewitt collection data: ‘Earrings’.

Contrasts can be revealed, such as the tag-timeline for ‘Personal Environmental Control’ (see Figure 79) where 18th-Century fans are, amusingly, contrasted with modern smart thermostats.



Figure 79. Visualisation of Cooper Hewitt collection data: ‘Personal Environmental Control’.

Sometimes the tag-timeline points to a time period. For example, ‘Art nouveau’ (see Figure 80)—a style of art popular between 1890 and 1910—shows a dense column of objects in this date span.

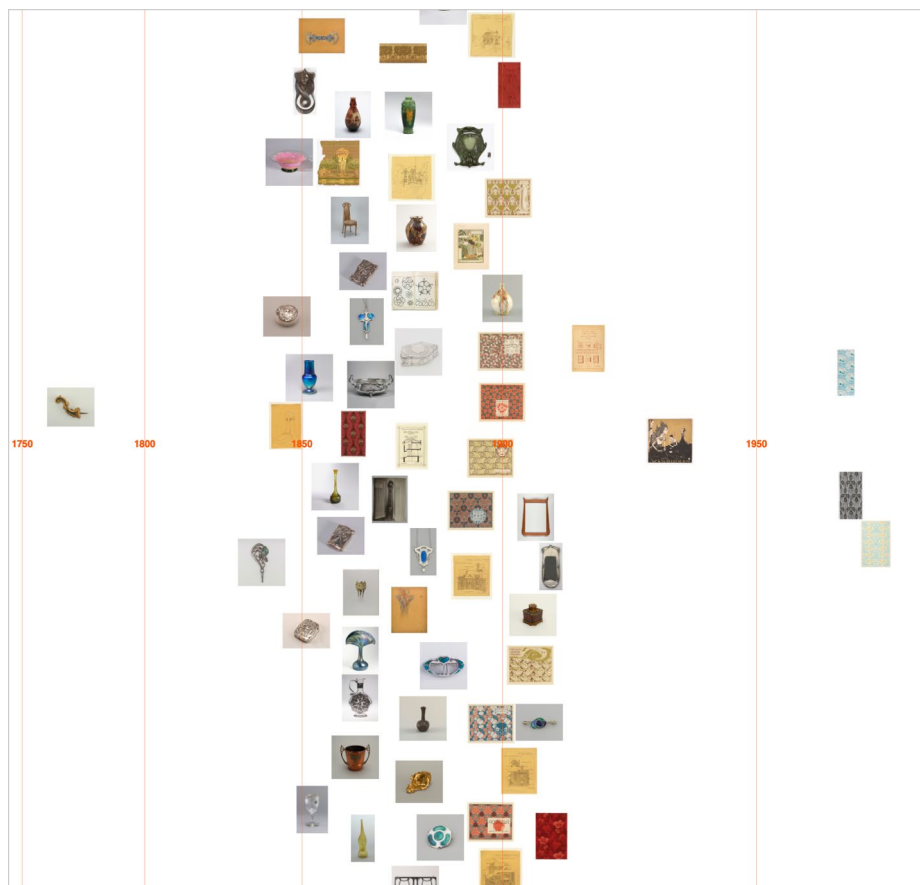


Figure 80. Visualisation of Cooper Hewitt collection data: ‘Art nouveau’.

Similarly, ‘Travel posters’ (see Figure 81) shows many examples from the 1920s-60s: a period described as the “golden age of the travel poster” (Grenci, 2014), though, as is often the case with cultural visualisations, larger historical trends and institutional collecting/cataloguing practices are entangled. Travel posters have been produced outside this ‘golden age’, but none appear in the timeline. Similarly, the museum has applied the tag ‘Art nouveau’ to designs, reminiscent of this art movement, produced in the 1960s/70s.




Figure 81. Visualisation of Cooper Hewitt collection data: ‘Travel posters’.

Historical moment

Some timelines suggest an historical moment, like this timeline for 'Skyscrapers' (see Figure 82) which shows a flurry of images from the 1920/30s onwards gesturing to a boom at that time in skyscraper construction in American cities.



Figure 82. Visualisation of Cooper Hewitt collection data: 'Skyscrapers'.

Similarly, the timeline for 'Space' (see Figure 83), contrasts 1960s Space Race souvenirs (Soviet and American) with modern telescope imaging. A lone 19th Century telescope off to the left is a reminder of the long history of human interest in space.

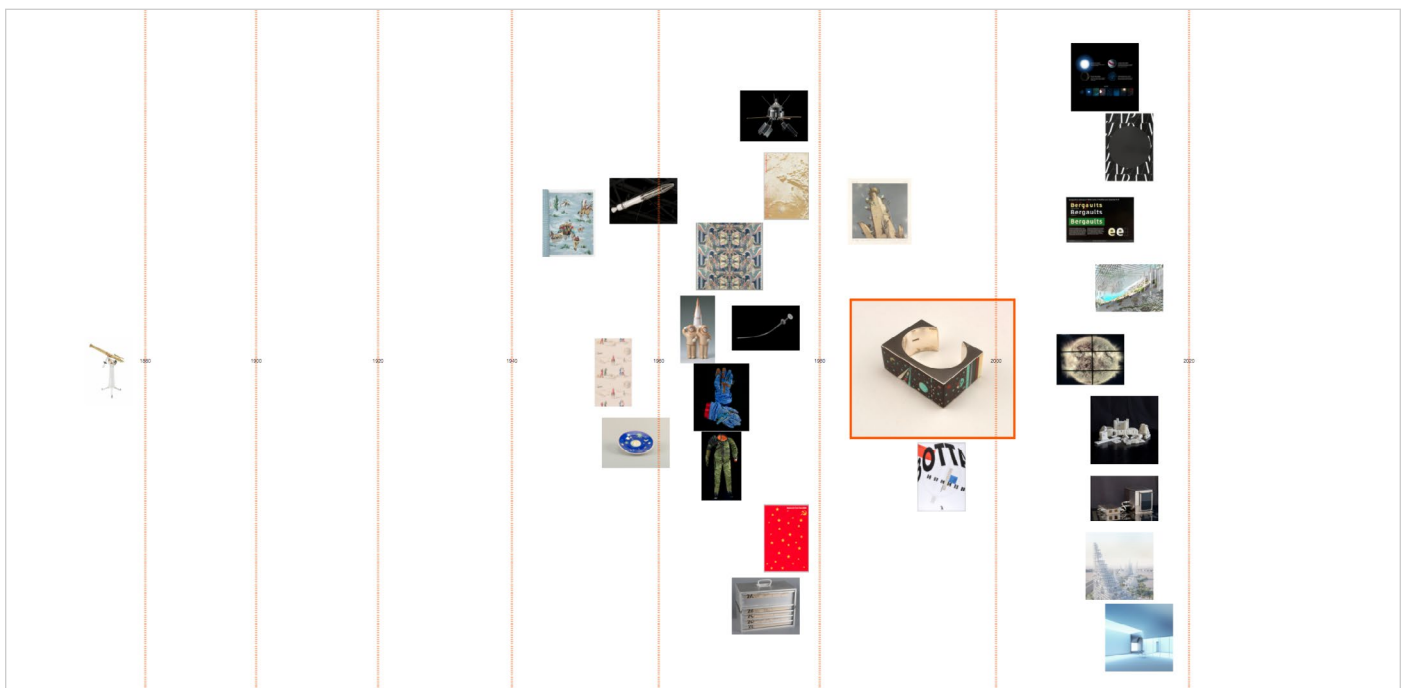


Figure 83. Visualisation of Cooper Hewitt collection data: 'Space'.

Institutional history

Some tag-timelines communicate institutional history. For example, the timeline for ‘Water’ (see Figure 84) visualises garden plans from the 16th–18th Century, through to modern interventions for accessing clean water in developing countries. This is hinting at an institutional shift at Cooper Hewitt from collecting around decorative arts—the museum’s original mission—to, now, collecting with a modern understanding of design.



Figure 84. Visualisation of Cooper Hewitt collection data: ‘Water’.

Metadata practice

The tag-timelines can shed light on the subtle ways different terminology is used in cataloguing, particularly by comparing tags which are close in meaning. Comparing tag-timelines for ‘Nature’, ‘Organic’ and ‘Biomorphic’ show ‘Biomorphic’ is more often applied to recent objects. ‘Nature’ is generally applied to depictions of nature such as landscape paintings, while ‘Organic’ is more applied to describe objects’ form.

The prototype setup only allows one tag-timeline to be viewed at a time. In future work, though, building in the ability to view multiple timelines at once may be useful for comparison, and for viewing objects against different contexts. It may also be useful to be to apply multiple tags as filters.

Abstract connections

The ‘Water’ tag-timeline also illustrates how certain tags have been applied both in the visual (eg. landscape painting and abstracted pattern) and functional (eg. drinking or boating) sense. Even when tags are applied in the visual sense, this can be in a literal or metaphorical way. The timeline for ‘Handkerchiefs’ (see Figure 85) contrasts historical handkerchief designs with a modern chair design “inspired by a crumpled handkerchief” (Cooper Hewitt, 2019d).



Figure 85. Visualisation of Cooper Hewitt collection data: 'Handkerchiefs'.

Object contexts

Since objects have multiple tags applied to them, they can be viewed against different historical contexts, for example this sundial/compass, which can be viewed as part of the timeline for 'Time-keeping' (see Figure 86) or 'Portable' (see Figure 87), or a number of other tags: each offers a different perspective on the object.




Figure 86. Visualisation of Cooper Hewitt collection data: 'Timekeeping'.




Figure 87. Visualisation of Cooper Hewitt collection data: 'Portable'.

Feedback sessions

Cooper Hewitt has considerable engagement with casual audiences and tagging its data was intended to improve collection accessibility. I, therefore, was interested to capture reactions to the design from a broader audience than my previous text visualisation tool, which focussed on scholarly users.

In the early stages of prototype development I discussed design choices with a Cooper Hewitt staff supervisor and presented my design at the museum curatorial meeting. Once the prototype design was settled and sufficiently robust/fast to demonstrate, I arranged to evaluate the tool with visitors in the museum galleries. I also demonstrated it to, and interviewed, a number of curators and education staff internally at the museum to explore their perspectives on what the design supports and its potential to serve as an internal tool. I wrote about the project on the Cooper Hewitt Labs (2017) blog, through which I received feedback online.

To seek feedback from members of the public visiting Cooper Hewitt, I built a touchscreen version of the prototype. A 32" touchscreen was installed in the galleries for 4 days. I conducted 30 semi-structured interviews lasting between 5-20 minutes, which I video recorded (see Figure 88). After demonstrating the interface, I invited participants to try it themselves. The purpose of the interviews was to:

- capture responses to the design and the ideas behind it
- get feedback on design decisions, for instance limiting the number of images displayed
- explore ideas for future development
- test the prototype's basic functionality



Figure 88. User testing/interviewing museum visitors in the galleries: video recording stills.

The museum was able to supply equipment and support for video recording equipment and I was interested to see if this would lead to added benefits for analysis. When I analysed the recordings, however, I found the additional video element added little and I returned to audio recording evaluations in the rest of this PhD. While observing visitor behaviour helped me identify pain points in the prototype's functionality, spotting these did not require reviewing recordings.

I conducted six semi-structured interviews with eight curators (one interview was jointly with three curators) of varying levels of seniority within the museum and from different museum departments (Drawings, Prints, and Graphic Design; Product Design and Decorative Arts; Textiles; Wallcoverings). I also conducted one interview with two education staff members. Interviews were audio recorded and lasted between 30-60 minutes. I demonstrated the interface to interviewees before starting the discussion.

Discussion

I have anonymised quotes taken from interviews, but distinguish between which session they derive from with letters: curator (C), education staff (E) and visitor (V).

Curators and visitors described the tool as encouraging discovery and exploration. It provides an “expansive” (C) view on the diversity of what the museum collects. It was seen as offering a playful, rich, visual experience of the collection. Visitors discussed the value the tool might bring in gallery and online settings. The curators generally imagined it as a public-facing tool for exploration. Some suggested, though, that internally it could support preliminary exploration for exhibition planning, as it offers fresh “new perspectives on historical material” (C). It was also discussed as a mirror on their metadata practice, and potentially useful for shaping future cataloguing. It was not seen as suited to focussed research, mostly I suspect because the tagging data was not seen as intended for supporting researchers.

Interviewees talked about the obvious benefits that arranging the data by time enabled: “for certain types of information it’s very useful to be able to see ... what changes and what stays the same, and how different things express themselves in different time periods, and certain ways” (C).

History, interpretation and the collection

Does this visualisation promote a particular manner of interpretation? As the curators described, the prototype emphasises the “very long continuity of design” (C) represented by the museum collection. Displaying themes by time brings the historic and contemporary together, suggesting how contemporary design may be influenced by historic examples. As one curator put it: “any time that we can encourage people to make connections with the past in design and the present, then we’re really doing our job, which is to show why historical design is relevant to the contemporary design” (C). I observed this kind of connection being made with ‘Tag-timeline’ during a visitor interview. The visitor, looking at the timeline generated for ‘lamp’, noticed two similar minimal, stalk-like designs separated by a large time span (see Figure 89): a floor lamp from the 1930s echoes the design of a simple tin candle holder from 150 years earlier. (Though in fact this visualisation conceals that the scale of one object is actually much larger than the other).

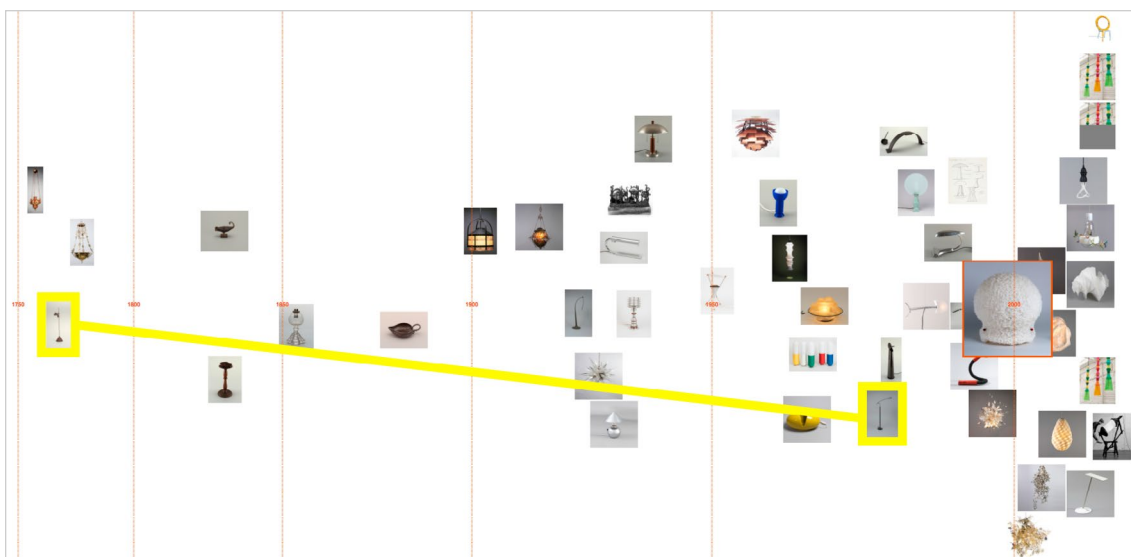


Figure 89. Visualisation of Cooper Hewitt collection data: ‘Lamp’, annotation in yellow showing similar minimal lamp designs.

Whitelaw (2015b:p.91) describes that cultural visualisation can be both a “conceptual model”—a thinking tool, built around cataloguing in the collection—and a “curatorial premise”—a statement about the collection as a whole. In this second function, ‘Tag-timeline’ aligns with the museum’s wider agenda in interpreting its collection historically: “it does execute the...underlying goals of what we say about design and the design elements that run through time...there are these threads that you can now visually see...it reinforces that” (C). One curator remarked it mirrors other ways the museum presents its collection to audiences: “what [‘Tag-timeline’ does] is kind of what we like to provide in these sort of isolated instances, which is cutting across discipline, and department, cutting across time, but you’re doing it in a broader way, so that it’s actually not just saying, ‘oh, I see one object representing arcs, from each department, from various decades’. But you look at these greater design trends” (C). Of course, this viewpoint/agenda may have contributed to the decision to create tagging data in the first place.

One curator suggested that the prototype is well suited to making sense of an idiosyncratic collection like Cooper Hewitt’s. Displaying the collection by tags serves to bring together masterpieces/design icons with more humble objects. They described the prototype as “treating it [the collection] in a very democratic way” (C).

Tags for timelines: reflecting on metadata practice

This [tool] is only as good as the tag (C).

One of the most significant topics discussed by the curators was how tagging practice shapes the experience of using the prototype, and how using the prototype might shape tagging practice. Some curators described a wish to “fill...out” (C) or “complete” (C) timeline displays: “this [prototype] would also pose the opportunity for the curators to start with the tag, and plug in all the other objects that actually should be [there]” (C).

The idea of filling in the timelines extended not just to adding more tagging metadata, but also collection acquisitions: “it’s a hole [in the timeline] and we know what it is that we would want to own from that period” (C). Timelines triggered thoughts about specific objects that were missing: “we should have considered the Fabergé! ... that would be an excellent diversion from all the glass and stuff [displayed in this timeline]” (C).

For some, there was a hankering to sculpt and define the narrative presented. As this quote describes, this curator was interested to go back and apply metadata so that the timeline narrative would match that of an historical argument:

[If we went over the objects present in this timeline for ‘Organic’ you could get] the thesis of that exhibition... the thesis of recurring organic movements that alternate with the neoclassical straight line movements would come up naturally... You’d see these sort of clusters in 1730 to 50, and the cluster in 1830s and 50s, and then a cluster of art nouveau was a very organic movement... in 1890 to 1910 (C).

A concern from education staff, however, was that with no narration or annotation on the timelines, it is difficult to identify and interpret patterns (if there are patterns, which is not always the case). Visitors, for example, may not know if the objects they see plotted in a time period are representative or pioneering for the time. This issue came up in some visitor interviews with some expressing they would like more narration. An issue to offering narration though, since tagging is ongoing, is that implied narratives in the visualisation may change as new objects appear in tag-timelines.

Do the tags speak for themselves?

The tags, representing greater design trends or historical discussions, can help guide and structure exploration in these visualisations. But are they alone enough? Do they make sense without the context in which they were applied? There was concern from curators that sometimes tag-timelines are difficult to make sense of because the reasoning behind tagging is not always explicit. Tags may have been chosen in the context of a past exhibition theme. Or the tag may relate to aspects of an object not discernable just from the visuals. For example, a dress is included in the timeline for 'Collapsible' because it is made of paper. A user would have to go into the record information to find this though; it is not obvious from the image alone.

Problems with tags and tagging

These interviews also prompted discussion about challenges—and sometimes reservations around the value of—tagging. Tags can be general/abstract (eg. 'Travel'), or highly specific (eg. 'Travel posters'). Is it better to show the larger picture? Or a tightly defined group? How general is too general to be useful? What works best on a timeline?

There was concern that tagging is inconsistent in how it is applied across the data. Different departments/individuals have different approaches (for example, in how specific a tag should be). Curators described the difficulty of deciding which themes to highlight for an object as there can be many possibilities. Since tagging is not meant to replicate existing cataloguing information, relevant tags may be omitted if objects already have the term somewhere else in their record. For some curators, there was concern about the fundamental utility of tagging; the deliberate use of non-specialist terminology that is ill-defined or non-specific meant that the connections sometimes felt "random" (C).

In visitor interviews opinion was mixed. Sometimes there was confusion: "So one [object] is from 1999 and that's from 1950 and I'm like, 'well, are they related at all? Are they from similar cultures?' I don't know" (V). But others enjoyed that the tags offer an unconventional set of connections for the museum context, that the kinds of connections can be abstract and are sometimes unexpected.

Since there is no controlled vocabulary to the tags some are very similar, for instance 'floral', 'floral bouquets', 'floral swag', 'flower', 'flowering vine', and 'flowers'. Some curators recommended that grouping or offering a hierarchy to related tags may help support exploration.

Incomplete tagging of the collection acts as an accidental curation of higher quality/of-interest objects: they are objects that have gone into exhibitions, or been highlighted for some other reason. This means the design cannot promote discovery of neglected objects/object records for curators; only well-researched objects are tagged.

Seeing everything?

I was interested to explore the importance placed on seeing everything in the collection data—particularly, following my decision to cap the number of objects visualised per timeline at 100. Some curators and visitors suggested they would like to zoom in on narrower time spans, or filter data on facets (for example, object type) to manage scale. On the whole, though, visitors did not seem much concerned about the 100 object cap or that only a portion of the collection has been tagged. Repeatedly, visitors told me they did not want to be overwhelmed—"there's something quite nice about the ... not sparsity of stuff but like graphically it's quite nice because you got everything there and it's composed"(V). Visitors said more important was that there was sufficient and diverse content representing the museum's holdings to explore.

These conditions (sufficient and diverse content) could also be applied to what makes a more satisfying individual tag-timeline—certainly, some displays are more engaging than others. For instance, Figure 90 shows the timeline for ‘Chimeras’ which could be said to be a less successful example. Aside from the fact ‘chimera’ is not a very common word, the content is a small number of visually quite consistent items, and no temporal patterns are obviously visible.



Figure 90. Visualisation of Cooper Hewitt collection data: ‘Chimeras’.

Design feedback

The design prioritises visual engagement with the material over other kinds. That is not necessarily a drawback: it depends on individual users’ preferences.

A couple of curators asked what it meant if an image overlapped with a line, suggesting they read the visualisation as organising objects into date range ‘buckets’ rather than positioned within their date spans. It may therefore be worth exploring different ways of marking dates on the time axis.

An imaginative suggestion I received in blogpost feedback concerned the possibility of pushing the collage aesthetic to the tool’s functionality too: allowing customising and curating of the timeline’s appearance. From the starting point of a tag-timeline, what would it be like if users could drag, crop, resize, hide or add in images? To construct a timeline that makes the argument you want to make?

Reflection on feedback sessions

These interviews were a reminder that the division between expert and casual audiences is not a strict one. A curator described how the tool could be useful internally at the museum for “the fellows, the graduate students, the people who are curators in other departments who are less familiar with [different parts of] our collection”. On the other side, many of the museum visitors I interviewed had connections (educational/professional etc.) to and/or knowledge of design and design history which they leveraged when interpreting the timeline displays.

Summary

This chapter introduced ‘Tag-timeline’, a collage-like visualisation of Cooper Hewitt Smithsonian Design Museum collection data by tags. This project highlighted the ways in which collection data is changing—in this case inspired by web resources—and the tensions this produces for those creating the data. It also underscored that with new types of data like this, new views on cultural data become possible as a result. Patterns—around the collection and design history more generally—aligned with some of the timeline displays, but these may change as new items are tagged. A number of the museum curators viewed these timelines as something to shape through further cataloguing. Casual users exploring with ‘Tag-timeline’ valued sufficient, diverse and representative content over access to every item in the museum’s holdings.

'Dive into Color'

Cooper Hewitt Smithsonian Design Museum



Figure 91. Olivia Vane. 2018. 'Dive into Color'. Photo: Caroline Koh Smith.

Throughout my fellowship at Cooper Hewitt developing the Tag-timeline prototype, I worked closely with staff at the museum. Discussing the prototype with the curators prompted them to suggest that the template I had developed could work well for visualising the collection by colour. As a curator explained: “[visualising by] colour, I think, is useful for the purposes of the study of the taste for different colours, but it is also a more interesting exercise for the public just to be able to do and get pleasure out of” (Cooper Hewitt curator, 2017c). Colour is enjoyable—it is eye-catching and vibrant—and it offers a visual way to explore a digitised collection without needing specialist knowledge. With a design collection like Cooper Hewitt’s, tracing colour through history also serves for looking at fashions and innovation in colour technology.

Colour was topical as the museum curators were, at the time, working towards an exhibition around colour theory and design. The stimulus of the Tag-timeline prototype and the timely opportunity for novel interrogation and presentation of the collection together had produced a new brief. In response to this new area of interest, I worked on developing the prototype in this new direction.

My fellowship had finished at this point so I was no longer based within Cooper Hewitt. I began to explore how Tag-timeline might be adapted to colour, keeping in contact with museum staff over email and video call.

Previous colour visualisations of cultural data

Colour has previously been used as a search facet in collections interfaces and as the basis for visualising collections. Hinchcliffe's (2016) 'Tate Explorer' offers colour as a search facet paired with a timeline. Barrett-Small's 'ColourLens' (2014) searches over Rijksmuseum and Walters Art Museum data by colour with a graphic for each item visualising its colour proportions. And Google Arts & Culture's 'Art Palette' (Doury & Ferrier, 2018) is a search engine that finds artworks based on a chosen colour palette.

Collections visualised by colour include the Library of Congress, where Wrubel created 'Library of Congress Colors' (2018) for viewing the colour palettes of objects over a collection, and Thorp (2018) visualised the colour names present in the titles of works. Also using Cooper Hewitt data, Abad created a visualisation of the colours present by decade in Cooper Hewitt's objects (Cooper Hewitt Labs, 2014). Lev Manovich and collaborators have visualised artworks, for example Mondrian and Rothko paintings (software studies initiative, 2011), by colour characteristics including hue and saturation. Reyes (2015) visualised Paul Klee's paintings by hue. And Foo's visualisation of the New York Public Library digitised collection (2016) has an option to organise items by colour.

All of these examples rely on computational techniques to identify colours rather than manual cataloguing. This is a classic case of technologies changing what can be researched, as these projects would all be near-impossible without digitised images and the ability to compute the colours within them.

Colour theory and design

I was interested to trace colours through the time dimension. Since the planned museum exhibition centred around colour theory and design, the curators also raised if exploring the collection by colour harmony (see Figure 92) could be integrated into a timeline visualisation.

Westland et al. (2007) describe colour harmonies in the context of art and design as colour combinations that give the impression of "visual rightness and balance"(p.1). They list pervasive contemporary schemes:

1. Monochromatic colour harmony (where colours are chosen with the same or nearly the same hue)
2. Complementary colour harmony (this is always represented as referring to opposite colours on a hue circle)
3. Analogous harmony (where colours are chosen with similar hues)...

Other examples ... include triadic colour harmony (three colours whose hues are each separate [sic] by about 120 degrees in the hue circle) and tetradic colour harmony (basically a double complementary scheme) (Westland et al., 2007:p.11).

Artists and designers create different visual effects, for instance balanced or contrasting, with different harmony types (Koenig, 2003:chapter 9). Westland et al. (2007), however, warn that "what is considered harmonious is to a large extent subject to fashion, personal preference and other cultural influences" (p.1) and "which combinations give rise to the pleasing effect is a question that has been of great interest for hundreds of years and shows no sign of abating" (p.1). This inherent subjectivity is an issue I would run into in prototyping.




Figure 92. Colour harmony examples.

Colour in Cooper Hewitt’s collection data

Cooper Hewitt had already enabled searching their collection by colour online (Cooper Hewitt, 2018). Colour data was extracted using the RoyGBiv module (Parvaneh, 2013), described on the Cooper Hewitt Labs (2013b) blog. Roughly, RoyGBiv works by checking the colour value of each pixel in an image, clustering colour values that are similar enough to be considered the same and returning up to 5 dominant colours in an image.

The colours extracted from Cooper Hewitt’s collection with RoyGBiv are good on the whole, but errors sometimes occur. The background colour is sometimes picked up. The effect of light and shadow on a 3D object can introduce multiple, illusory colours, for example Figure 93.



Figure 93. Light and shadow on 3D objects can introduce multiple, illusory colours when pixel values are used to extract colours for objects. Object IDs (left to right): 35460799; 18806261. © Cooper Hewitt, Smithsonian Design Museum.

As always there are quirks working with digitised collections, like tiles (see Figure 94) which had coloured stickers on them when they were photographed, or buttons (see Figure 95) photographed next to dots of yellow paint.



Figure 94. Tiles photographed with stickers. Tile wall facing (Netherlands), ca.1725. Object IDs (left to right): 18726001; 18725273; 18725393. © Cooper Hewitt, Smithsonian Design Museum.




Figure 95. Buttons photographed with yellow paint. Object IDs (left to right): 18309573; 18309683; 18309545. © Cooper Hewitt, Smithsonian Design Museum.

Lace was sometimes photographed against a dark background for contrast, see Figure 96.




Figure 96. Lace photographed against a dark background. Object IDs (left to right): 18569397; 18569395; 18318023. © Cooper Hewitt, Smithsonian Design Museum.

Different perspectives on reliability of metadata

Interestingly, my conversations around colour data when I had been working in the museum revealed there were very different views on the reliability/quality of this extracted colour data. In general, the curators characterised the colour data as “very unreliable” (Cooper Hewitt curator, 2017c), particularly because background colours were sometimes picked up. Staff working in Digital at the museum, by comparison, described that the technique had been very successful. The colour calculations are correct in a simple, perhaps rather stupid, sense. The problem is that the result does not correspond to what a human would be likely to say about the objects depicted, such as “this chair is all the same shade of red” or “this lace is simply white”. Is it worth visualising data that is considered unreliable (by some) to begin with? Will the resulting visualisation be inherently untrustworthy?

Since the possible number of unique colours extracted across the collection is huge, searching by colour on the Cooper Hewitt website, at the time, was simplified by snapping extracted colours to the closest value in a standardised palette (the default was the CSS4 web colour palette, but the CSS3 and Crayola palettes were also options). On the Cooper Hewitt website, you could search the collection by 116 CSS4 colours (see Figure 97). Both the original and snapped-to colour palettes are available in the Cooper Hewitt data (Cooper Hewitt, 2017) – all stored as hex codes (six hexadecimal digits representing the levels of red, green and blue).

Data issues

Although Cooper Hewitt has an API, there was no method for returning all the objects matching a colour. Instead, I used a collection dataset Cooper Hewitt had put on GitHub (2017) earlier that year. While APIs may be seen as a straightforward way to give external developers access to collection data, their design may not anticipate how developers wish to use the data.

Even with access to Cooper Hewitt's collection data on Github though, the dataset only contained about 130,000 unique objects (of approximately 200,000 total in the collection). Accompanying documentation did not explain why and following up with staff did not recover an explanation other than a suspicion the Github dataset was out of sync with internal collection data updates. Since the Github dataset still covered a large number of objects, I decided to proceed with prototyping with the data available.

With this new brief, I was interested to find out: what the issues are in visualising a subjective quality like colour harmonies, what a timeline visualisation of Cooper Hewitt data by colour/colour harmonies would show and reveal, and whether there would be differences in visualising computationally-extracted data (like colour)—known to contain errors—versus manually-added data (like tags)?

Visualisation prototyping

As a first step, I adapted my code to visualise collection items matching a CSS4 colour along a timeline (see Figure 98 to Figure 100).

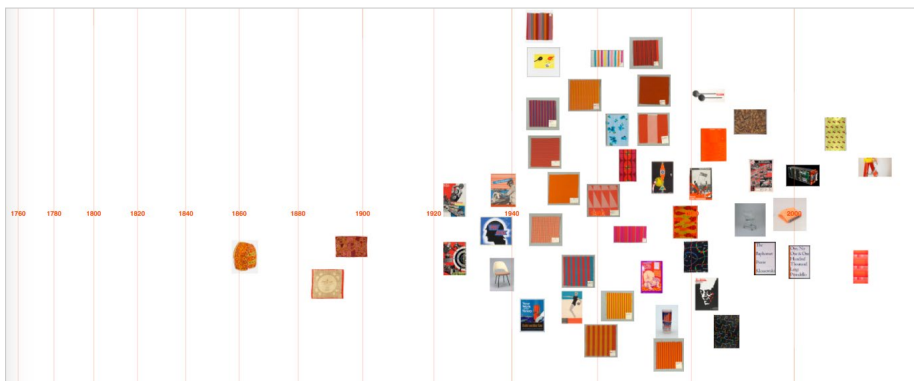


Figure 98. Visualisation of Cooper Hewitt collection data: 'Orangered'.



Figure 99. Visualisation of Cooper Hewitt collection data: 'Steelblue'.




Figure 97. Colour search on the Cooper Hewitt website (Cooper Hewitt, 2018). Screenshot taken Aug 2018. © Cooper Hewitt, Smithsonian Design Museum.




Figure 100. Visualisation of Cooper Hewitt collection data: 'Olivedrab'.

My next step was to explore how I might visualise objects featuring a colour harmony, first trying complementary colours (opposites on a hue circle). I initially tried to do this by matching a chosen CSS4 colour with the nearest CSS4 colour of opposite hue. The HSL and HSV (hue, saturation, lightness/value) colour systems define hue as an angle round a circle (0-360°), so I inverted hues by converting hex codes to HSL. The visualised results were unsatisfying though, as often the search failed to find any matches.

This does not mean to say there was a lack of objects with complementary colours, but that my search was too precise. Firstly because the colour data is imprecise anyway (colours have been snapped to a palette and the colour extraction method is not 100% reliable). More crucially though, complementary colours are opposites on a hue circle, but within what bounds? Quantifying colour harmonies has proved very difficult and there are no metrics in consensus (Westland et al., 2007:pp.12-13).

I tried extending the reach of my search to matching several colours close to the inverted hue, but it felt very frustrating not to have a visual reference to the range of colours in a region and what colours were being searched over. I wanted to *see* my search.

So I started experimenting with using a colour wheel input as a way to pick colour combinations and simultaneously see possible hue relationships. I first tried mapping the colours from the standardised palettes by HSL round a circle (see Figure 101).

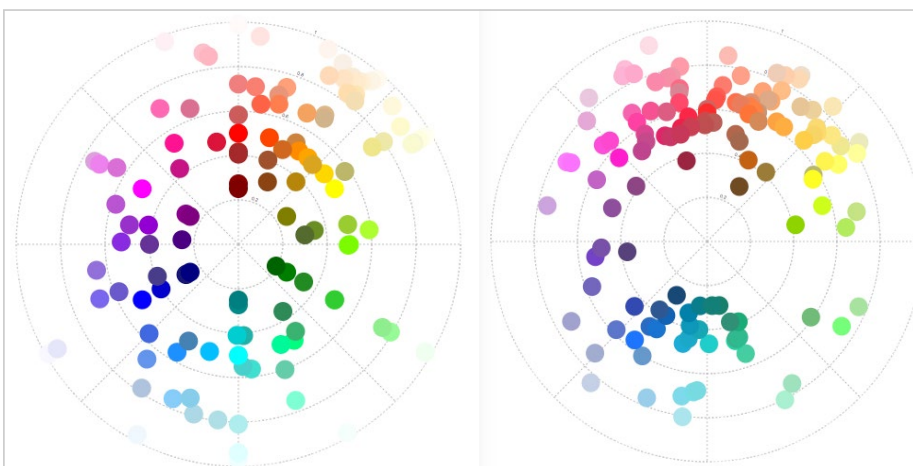


Figure 101. Snapped-to colours in the Cooper Hewitt collection. CSS4 (left) and Crayola (right) palettes mapped by hue (in HSL). Angle = hue, radius = lightness.

To make it easier to see the possible colours and so the design would more closely resemble a traditional colour wheel, I wrote some code to map the CSS4 palette colours to a wheel (see Figure 102).




Figure 102. CSS4 palette colours (left) mapped round a wheel (right). Hue (HSL) = angle, ordered by lightness.

Colour wheels

I realised at this point, though, that the resulting design does not match a typical artist's/pigment colour wheel (which has red, yellow, blue – RYB – primary colours) (Harkness, 2006:p.222). If complementary colours are opposite colours on a hue circle—which hue circle?

HSL is a simple transformation of RGB colour space, and therefore the wheel has red, green, blue primaries (Harkness, 2006:p.221)), see Figure 103. If this colour wheel is used to search over design artefacts, surely it would be more appropriate to use a design closer to the norm for artists and designers?

There is “no one perfect or ideal colour wheel” (Harkness, 2006:p.229), but converting my HSL colour wheel to something closer to a pigment version (using code from Knight's (2016) implementation of Adobe's 'Kuler' colour wheel) seemed a reasonable compromise here, see Figure 104.

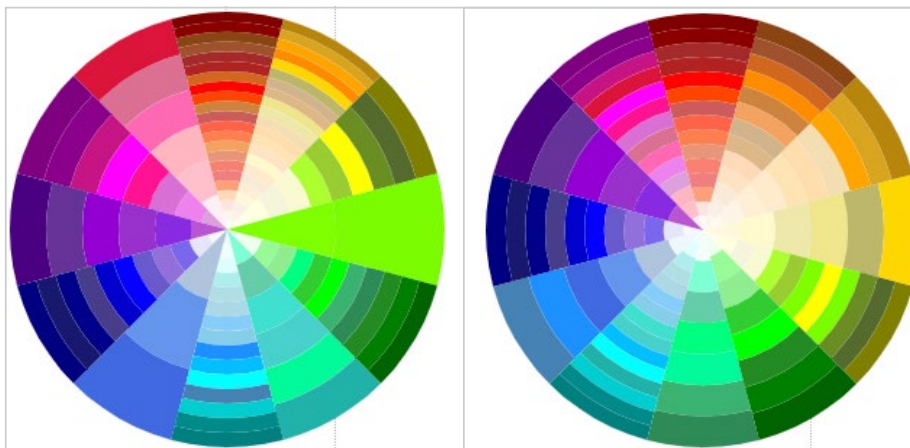


Figure 103. (left) Colours mapped by hue (HSL) resulting in an RGB colour wheel. Angle = hue, radius = lightness.

Figure 104. (right) Colours mapped by hue (HSL), adjusted to more closely resemble a pigment (RYB) colour wheel . Angle = hue, radius = lightness.

Using this colour wheel as a guide and an input, I could see and choose which colours to query. By searching over multiple colours, the visualised results were better. In Figure 105 and Figure 106, white and black borders around tiles in the colour wheel indicate the searched-over colour combinations.



Figure 105. Visualisation of Cooper Hewitt collection data: Purple & olive.




Figure 106. Visualisation of Cooper Hewitt collection data: Orangered & cyan/blue.

Querying against colour data in HSV

While the results looked better with this prototype, the user interface is a mess and complicated to use. And the search query was not excluding objects that featured other colours in addition to the searched colour combination.

Sticking with the CSS4 palette was greatly complicating the task, so I abandoned using it. I converted all the original extracted colours (not snapped-to) from hex codes to HSV and created my own Elasticsearch (2019a) index with the colour data stored as a nested datatype (Elasticsearch, 2019b): a way of storing data with multiple constituent values (here hue, saturation and value) that ensures queries treat each trio of values as a single entity. This way I can: search over a hue range, with a threshold on saturation and value; exclude objects also featuring other hues; and it is easier to generally define more complex colour harmony searches (e.g. analogous, triadic, quadratic and split complementary) (see Figure 107).




Figure 107. Different colour harmonies tried in prototyping.

Colour wheel graphic for objects

As a by-product, I realised I could repurpose my code to map individual object palettes round a colour wheel too. Thus, you get a compact graphic describing the colour relationships present in a single design (see Figure 108 and Figure 109). This is a nice example of the serendipity of designing (discussed in my Methods chapter: ‘Cultural visualisation prototyping’), where you identify new possibilities as a result of seeing what you have already made.



Figure 108. (upper) Object with colour palette, (lower) palette mapped to a wheel, easy to see complementary harmony.




Figure 109. Colour wheel graphic examples.

Colour palette-to-wheel algorithm

1. Convert all palette colours from hex codes to HSL
2. Exclude low saturation, or very dark colours.

Based on experimenting, I set a minimum threshold of 0.15 (out of total 1) on saturation (S) and lightness (L).

3. Assign an angle for a colour.

In HSL, hue (H) is measured (0-360°). Adjust H angle to bring it closer to artist/pigment colour wheel angle using Knight's (2016) code.

4. Colours are assigned to 1 of 12 sectors around a circle, based on which sector the angle falls in.
5. Each sector is divided into ring (annulus) sectors by the number of colours assigned. The ring sector areas are ordered by lightness (lightest in the circle centre) and coloured with the assigned palette colour.

In the visualisation interface, I adapted a simple artist's (RYB) colour wheel to use as an input, (see Figure 110). There were few hits for the more complex harmonies (triadic, tetradic, split complementary) and the results felt less convincing. I had widened the hue range to search over in order to increase the small number of hits, so the results were less visually cohesive. In conversation with the museum curators, we decided to drop these more complex harmonies from the visualisation.




Figure 110. Prototype visualisation exhibited at Royal College of Art's January 2018 'Work in Progress'. Photo: Olivia Vane.

Manual data and query adjustments

As mentioned earlier, it is tricky to know the numerical bounds of a colour harmony: what angle of hues should queries cover? This is further complicated because the circle angles/proportions covered by the different hues are unbalanced: for the pigment colour wheel, blue covers 125°, while green covers only 45° (Harkness, 2006:p.228).

To improve search results, I manually adjusted the hue angles over which the different harmony schemes search (and depending on what colours are being searched over). I made adjustments by personally judging the quality of results. For monochromatic and complementary, I made the search query exclusive—results exclude matching objects with additional non-matching colours. But for analogous (colours are chosen with similar hues), I made the search non-exclusive as otherwise the results were dominated by monochromatic examples. Because colour harmony is a subjective quality, I felt more comfortable making adjustments behind-the-scenes to the visualisation function to, by my eyes, improve the visualisation.

At this stage, since this setup only allows a fixed set of possible searches, with repeatable results, it was worth it for me to do some manual editing of the colour data to remove obvious errors. In addition, there were some objects I manually excluded, for example textile sample book pages (see Figure 111), where multiple designs are included in 1 photograph as it arguably does not make sense to include these in the visualisation.



Figure 111. 'Sample Book (France), 1938; Made by Textile Argus'. Object ID 18431057. © Cooper Hewitt, Smithsonian Design Museum.

Sharing my prototyping with the museum curators, their impression was my manual adjustments were “very effective” and the resulting colour/colour-combination timelines were “very satisfactory” (Cooper Hewitt curator, 2018). The decision was taken to exhibit a polished version of the visualisation in the upcoming exhibition—settling the visualisation audience as museum visitors.

In order to avoid reducing the size of the images (so it is still possible to see what the objects are), following museum visitor feedback from the Tag-timeline project, I capped the number of visualised objects to the 100 most saturated in colour. For some searches, there were many results (for one monochromatic search on a yellow hue, the results were approximately 5,000). For other queries though the total results were below 100, so all matching items are shown in those cases.

Final interface design

The final interface design—titled ‘Dive into Color’—was exhibited at Cooper Hewitt (2019f) in the exhibition ‘Saturated: The Allure and Science of Color’ May 2018–March 2019 (see Figure 112 and Figure 113). On suggestion from the curators, for the final design I adopted a colour wheel input (see Figure 114) inspired by a Hilaire Hiler design (Cooper Hewitt, 2019e) in Cooper Hewitt’s collection (see Figure 115). This wheel is a ‘perceptual’ colour wheel (Harkness, 2006:p.222) with 4 colour primaries (red, green, yellow, and blue). It features 30 hues. The interface has 4 harmony options: monochromatic, complementary, analogous and spectrum (a rainbow colour option).



Figure 112. Olivia Vane. 2018. ‘Dive into Color’ installed at Cooper Hewitt. Photo: Caroline Koh Smith.



Figure 113. Olivia Vane. 2018. 'Dive into Color' installed at Cooper Hewitt. Photo: Caroline Koh Smith.



Figure 114. 'Color Wheel, 1936–37; Hilaire Hiler'. (Cooper Hewitt, 2019e). © Cooper Hewitt, Smithsonian Design Museum.



Figure 115. Detail from 'Dive into Color' showing Colour Wheel picker, inspired by Hiler's design.

My query structure, developed for harmonies, was not suited to the spectrum option. Plus the colour extracting method only identifies up to 5 dominant colours in an image. I had some limited success finding spectrum objects using my hue querying technique. In the end I manually coded in spectrum matches, including objects identified by text searching the collection data. Figure 116 displays the spectrum results.

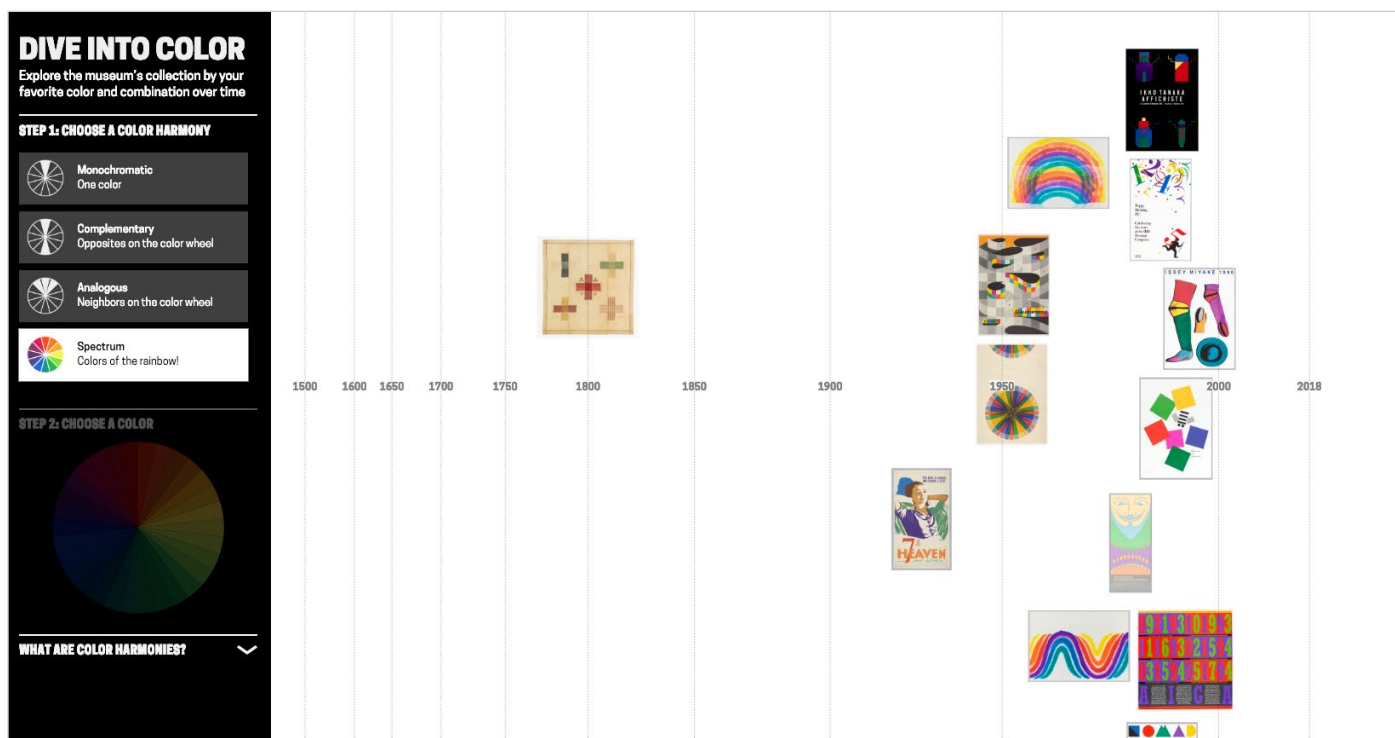


Figure 116. Visualisation of Cooper Hewitt collection data: Spectrum.

What does 'Dive into Color' reveal?

Visualising the Cooper Hewitt data with this approach gives some sense of when colours appear in time. There are no results for purple pre-19th Century (see Figure 117 and Figure 118), perhaps because of the difficulty and expense of producing purple before synthetic dyes/pigments were developed in the 19th Century (Barnett, Miller & Pearce, 2006:p.448; *ibid.*:p.451)?



Figure 117. Visualisation of Cooper Hewitt collection data: Monochromatic, purple.

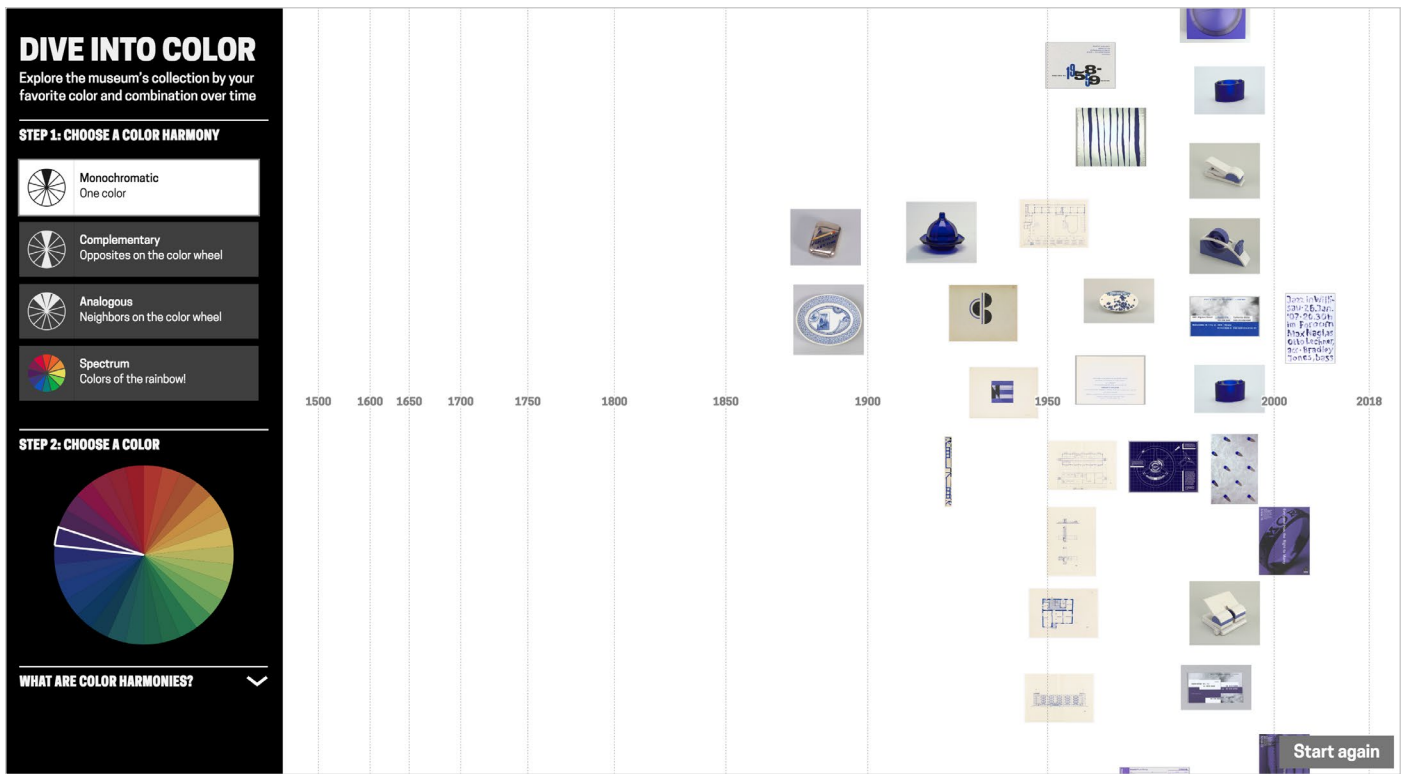


Figure 118. Visualisation of Cooper Hewitt collection data: Monochromatic, purple.

As often is the case interpreting collection visualisations, it is difficult to disentangle historical trends from the biases and character of what has been collected and how it has been catalogued (and bear in mind I am only visualising the 100 objects most saturated in colour for a search). For example, these green and purple Japanese prints of irises (Figure 119) are clearly part of a set rather than indicating some colour trend around 1910. Using the images themselves as data points is helpful for diagnosing this, and emphasises how misleading a traditional quantitative data visualisation might be.

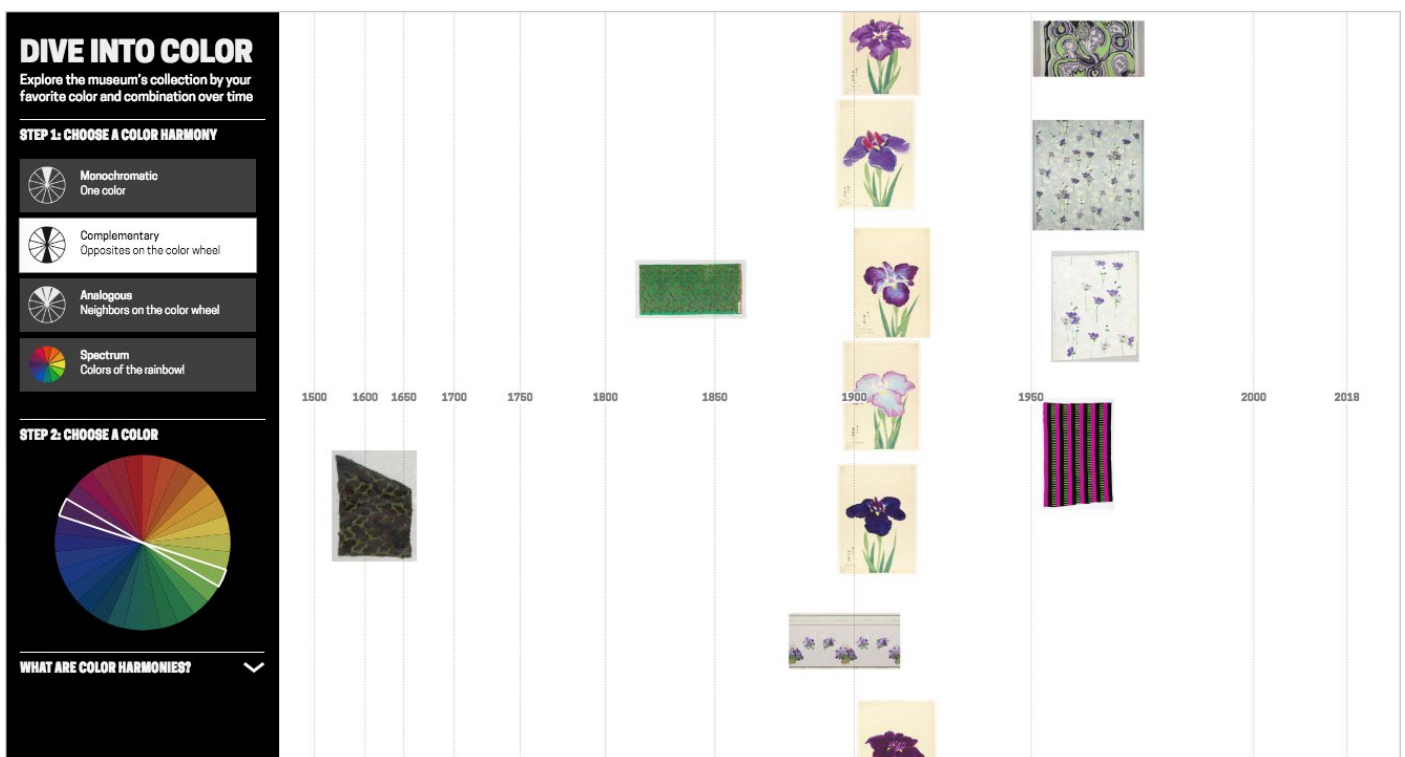


Figure 119. Visualisation of Cooper Hewitt collection data: Complementary, purple & green.

Similarly, the visualisation for red (see Figure 120) shows a column at 1900: the result of choosing to catalogue each playing card from a deck as a separate record.




Figure 120. Visualisation of Cooper Hewitt collection data: Monochromatic, red.

The purple-green visualisation though demonstrates how the tool can connect artefacts across time, in this case by similar colour scheme/design (see Figure 121 and Figure 122):



Figure 121. 'Frieze (USA), 1890–1910; Manufactured by Hobbs, Benton & Heath'. Object ID: 18500019. © Cooper Hewitt, Smithsonian Design Museum.



Figure 122. 'Sidewall, Anemone, 1960–66; Designed by Phoebe Hyde'. Object ID: 18459631. © Cooper Hewitt, Smithsonian Design Museum.

The tool surfaces colours, used in a particular material, that are strongly attached to design types. For example blue and white ceramics manufactured in the Netherlands in the late 17th Century and early 18th Century (see Figure 123 to 126):

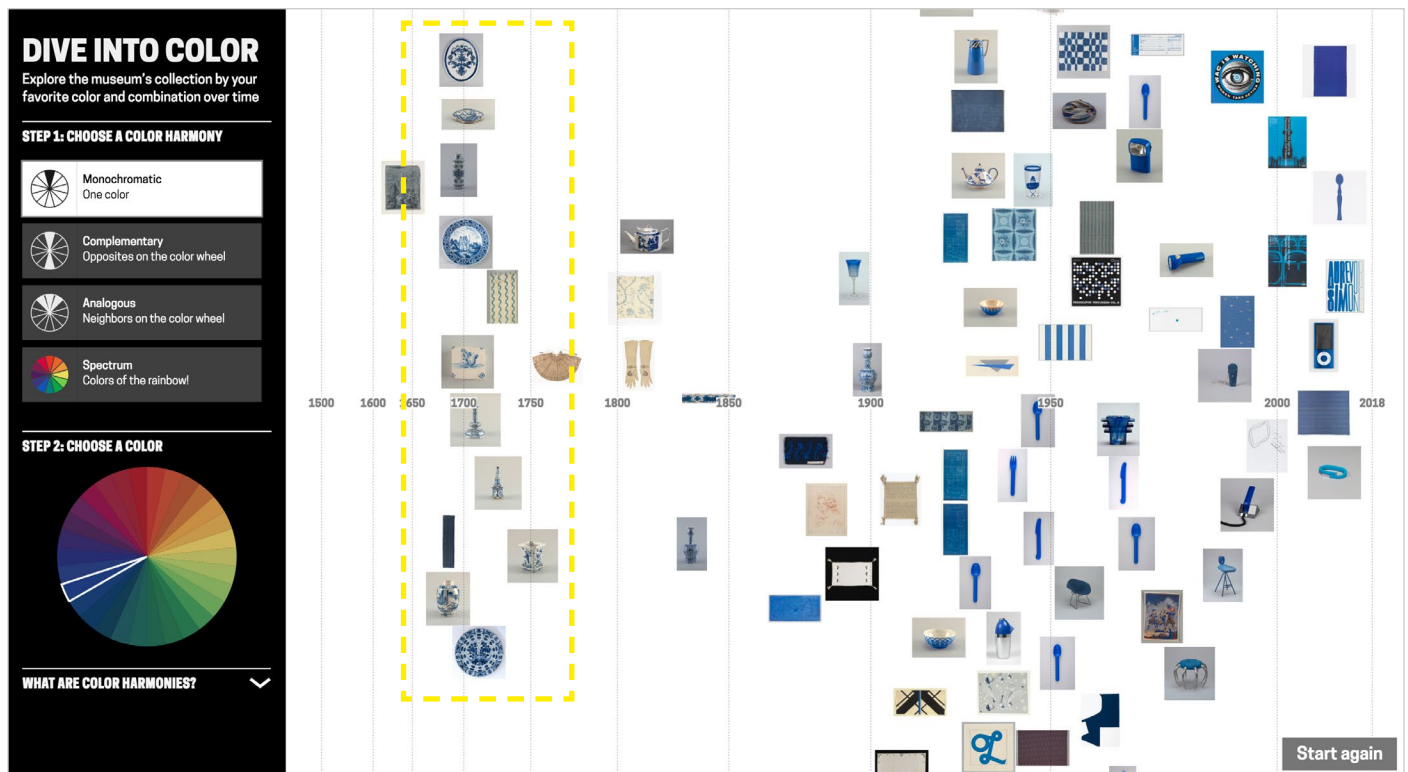


Figure 123. Visualisation of Cooper Hewitt collection data: Monochromatic, blue. Annotated to show late 17th Century/early 18th Century Dutch blue & white ceramics.



Figure 124. 'Plate (Netherlands), 1675–1725; tin-glazed earthenware'. Object ID: 18621479. © Cooper Hewitt, Smithsonian Design Museum.



Figure 125. 'Plaque (Netherlands), 1675–1725; tin-glazed earthenware'. Object ID: 18621427. © Cooper Hewitt, Smithsonian Design Museum.



Figure 126. 'Obelisk (Netherlands), ca. 1700–25; tin-glazed earthenware'. Object ID: 18621483 © Cooper Hewitt, Smithsonian Design Museum..

French red and white textiles in the late 18th/ early 19th Century (see Figure 127 to 130):

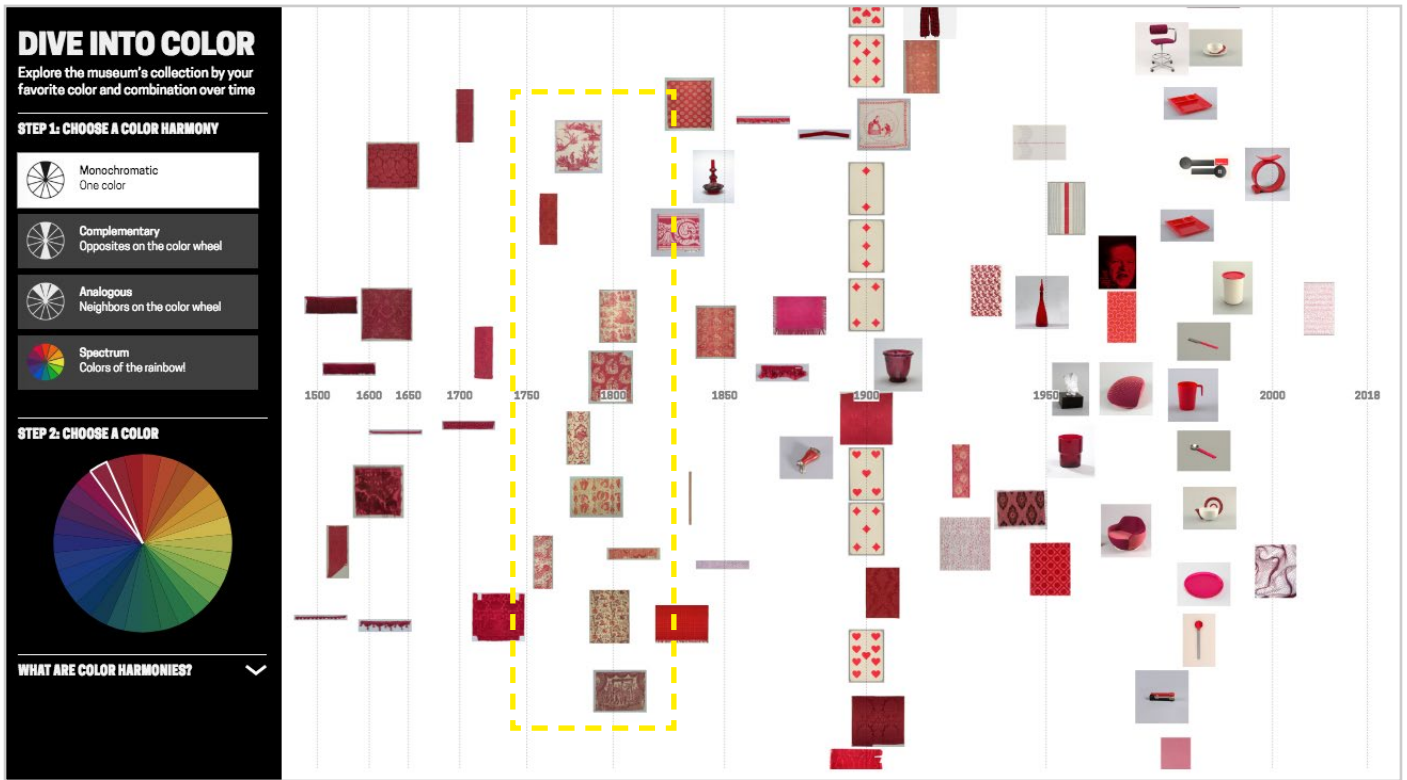


Figure 127. Visualisation of Cooper Hewitt collection data: Monochromatic, red. Annotated to show French red and white textiles in the late 18th/ early 19th Century.



Figure 128. ‘Textile (France), late 18th Century; cotton’. Object ID: 18438769. © Cooper Hewitt, Smithsonian Design Museum.



Figure 129. ‘Textile (France), ca. 1850; cotton’. Object ID: 18671007. © Cooper Hewitt, Smithsonian Design Museum.



Figure 130. ‘Textile (France), 18th Century; cotton’. Object ID: 18482057. © Cooper Hewitt, Smithsonian Design Museum.

Or these English vivid blue and white, Wedgwood-type, late 18th Century ceramic buttons/ medallions (Figure 131 to 134):

Again one should be cautious making quantitative assessments from the visualisation. These buttons actually all come mounted together on one board (see Figure 135); they have individual and group digital records. It is unclear who assembled them and why.

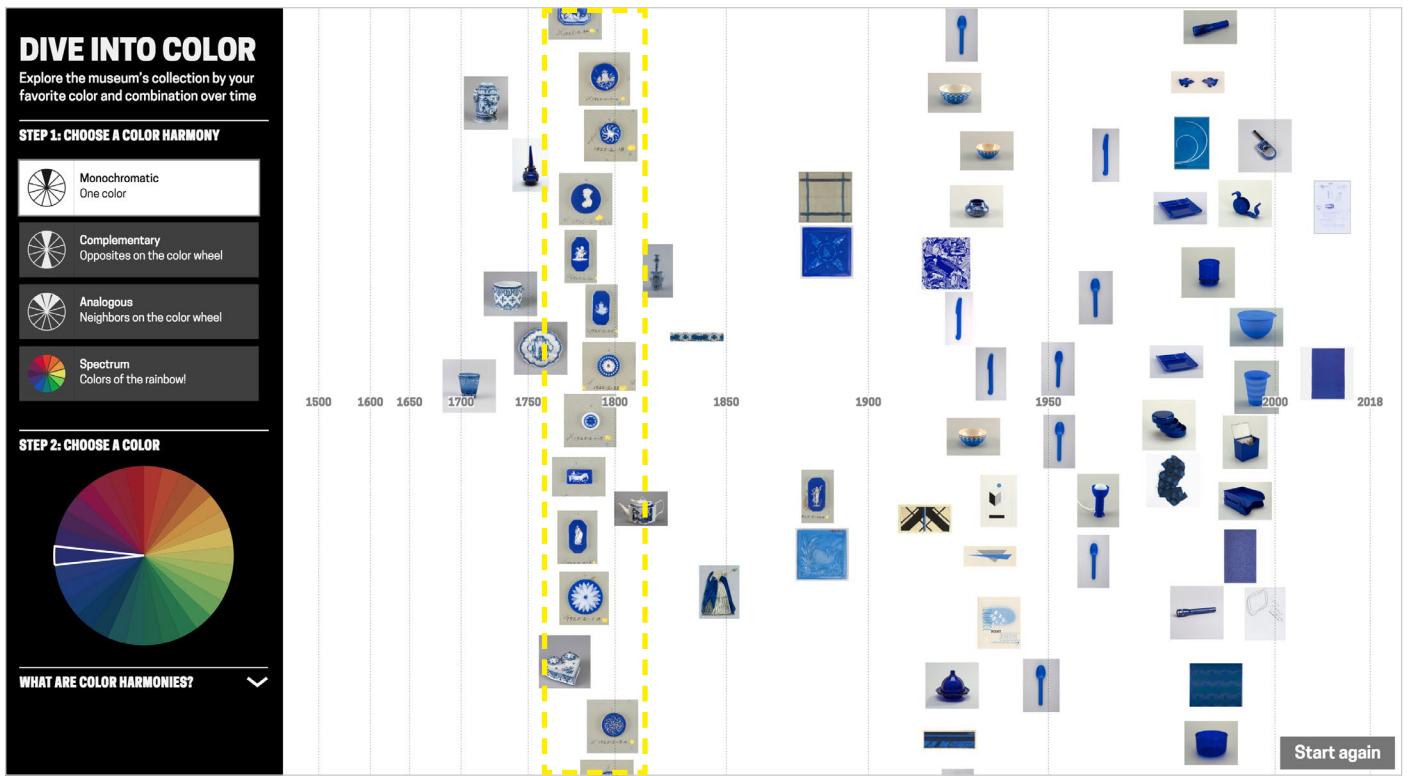


Figure 131. Visualisation of Cooper Hewitt collection data: Monochromatic, blue. Annotated to show English blue & white, Wedgwood-type, late 18th Century ceramic buttons/medallions.



Figure 132. 'Medallion (England), late 18th Century; stoneware'. Object ID: 68730605. © Cooper Hewitt, Smithsonian Design Museum.



Figure 133. 'Medallion (England), late 18th Century; stoneware'. Object ID: 18711563. © Cooper Hewitt, Smithsonian Design Museum.



Figure 134. 'Button (England), late 18th Century; stoneware'. Object ID: 18310009. © Cooper Hewitt, Smithsonian Design Museum.




Figure 135. Paper title on board reads: 'Buttons and Mounts, Wedgwood type; England and Belgium, 18th and 19th Centuries'. Object ID: 69112317. © Cooper Hewitt, Smithsonian Design Museum.

Visualising blue-yellow shows more saturated colour from mid-19th Century onwards (Figure 136). Is this signaling changing fashion, or the availability of new synthetic dyes/pigments? Can we connect the more saturated harmony in designs from the mid-1800s with Chevreul’s influential text from the time, ‘The Law of Simultaneous Color Contrast’, describing how colour harmony can be used to create a more vibrant effect? (Westland et al., 2007). Though a number of the earlier objects are textiles and the colours will have faded over time, possibly a combination of these factors is at play here.



Figure 136. Visualisation of Cooper Hewitt collection data: Complementary, blue & yellow.

Feedback sessions

I captured reactions to the final design from museum curators through dialogue over video call (with two curators). As the project had progressed, I became aware of other prospective audiences—in addition to museum visitors—where such a tool might be relevant, particularly researchers looking at history of design and colour. I therefore arranged a number of additional semi-structured interviews, to explore impressions of the tool from these audiences:

- six interviews with eight Masters students in History of Design from the Royal College of Art/ Victoria and Albert Museum programme (two interviews were with two students at a time).
- three interviews with colour history specialists (two historians with expertise in architectural colour and colour systems, and the collection manager of an historical colour library).

In addition, I was interested to explore if the design would be relevant to designers using collections for inspiration and reference, which I return to at the end of this chapter.

In the following discussion, I have anonymised interview quotes, but distinguish between which session they derive from with letters: curator (C), history of design student (HoD S) and historical colour specialist (HCS).

Discussion

The design was seen as playful, visually appealing, and supporting serendipitous discovery in the collection, giving a snapshot of the diversity of the museum's holdings. Different audiences talked about the visualisation as a new 'image' of the collection: "it's a kind of picture that appears" (HoD S); "[it is a] composition of what's in a gallery" (C). And there was appreciation of the minimal aesthetic employed: "it's just very [visually] clean... everything looks brilliant from a visual perspective" (HoD S).

Narrative and history

The curators discussed the possibility of adding narration/annotation to the timelines to emphasise the thesis of the exhibition; historical colour theories proposed at various times could be plotted within the timeline to suggest how the designs that followed may have been influenced.

This visualisation may lay out the collection data in a very democratic way, but key historical moments are not highlighted. Further, by relying on colour data which has errors, key historical moments cannot be guaranteed to appear anyway. As one of the curators noticed, "Perkin's mauveine scarf, that is the invention of purple dye in 1856, doesn't show up on here. But we're calling it a major moment in color history [as it is the invention of the first synthetic dye, which then takes the fashion world by storm (Cooper Hewitt, 2019h)]...If there are key points, we want to make sure they aren't missed in this interaction". On investigation, this scarf (see Figure 137; Cooper Hewitt, 2019h)—incidentally tagged 'synthetic dyes'—is absent from the purple timelines because the colour extraction method misses the intense purple colour in this object—it is not in the colour data. (I previously pointed to cultural visualisations' tendency to draw attention to errors and idiosyncrasies in the data in my 'Background' chapter).

Similarly because the timeline focusses on colour harmonies, it does not capture historical fashions if they do not fit the harmonies structure: "what about when we go into the later Victorian Age? What about the polychromy issue, where it [interiors and decorative arts] becomes really multicolored?" (HCS); this is not visible in 'Dive into Color'.

Asking questions with/of the data

Capping the visualised objects at 100 for a colour search was seen as fitting for exhibition purposes and for serendipitous exploration: "if it's just to understand a snapshot of what is available, what yellow things could be there, then that's perfect"(HoD S). (As a reminder, due to issues accessing data at the beginning of this project, the visualisation did not search over all the collection data anyway).

While the design has potential for use in supporting researchers and collection managers, there would need to be adjustments to the design allowing more control in search; filtering on different data facets, time-spans and combining with text search were all recommended. Researchers—looking for trends and patterns, rare and elusive objects, and making connections currently invisible in the collection—or collection managers—making sure there is a record for everything and that it is discoverable—would require access to the complete collection data. Though, as illustrated earlier, completeness does not negate that analysis of the visualised collection data needs to be conducted with caution.



Figure 137. Scarf Sample (France), mid-19th Century. (Cooper Hewitt, 2019h). © Cooper Hewitt, Smithsonian Design Museum.

While clusters of objects (and also absences) may point to trends in fashion and taste, patterns of influence, the availability of new dyes or pigments, or makers/design types renowned for using particular colours, the visualisation on its own does not lend explanation. “[It] poses questions rather than necessarily answers them” (HCS).

Colour functionality

I could never [have] thought this kind of project can exist, 'cause number and colours are just too different (HoD S)

The colour harmony functionality fits the exhibition narrative, but to ask wider questions about colour in the collection, various interviewees described how it would be useful to be able to custom pick colour combinations and be able to search for black, white and grey.

The history of design students discussed examples from their own work where such a tool could be useful. These projects were not looking at pigment history, but hue choice in design. Example projects included: tracing the use of blue through time in visual communication of the anti-vaccination movement to convey trust; or pink clothing in the history of women's protest movements. In both these examples a hue, rather than a more specific colour, was of interest.

The current setup does not always allow the user to isolate the colour they want. When first discussing the project with curators at Cooper Hewitt, one curator suggested they would be interested to test a hypothesis that orange, as a colour used in interiors, rose in popularity in Western Europe and America from the 1920s onwards (Cooper Hewitt curator, 2017c). Searching for orange objects in this interface, however, brings up lots of darker hued, brown objects including items made of wood (and you cannot filter only items relating to interiors). It cannot isolate the items that would help investigate this theory.

For the specialist colour audience, this tool design has some issues. Not least because I do not know how accurately the photos represent the true colours of the objects (what lighting conditions the photographs were taken in, if they have been retouched etc.). While the overall extracted colours seem generally good, they may not be precise enough for some. The control in searching by colour—only by hue—may also be too limited for some needs. Though it was pointed out that developing a more sophisticated tool would require a high level of customisation (there is no consensus on colour systems) for, likely, a very small audience size.

Seeing is believing?

Colour data is computationally extracted, in contrast with manually added metadata. Do these different cases require different considerations in visualisation design? In this project, I have used arguments like it 'looked better', or the results were 'more satisfying' to explain my design decision-making. I was working with colour data I knew had errors to visualise subjective qualities. As a result, I was more comfortable adjusting parameters in my search queries and editing obviously wrong colour data to return what looked *better* to me. Because colour is an immediate visual quality, you can easily see when images appearing in the visualisation do not match. (Though, of course, looking at the visualised results will not tell you if there are absent items).

In interviews, I asked whether this data massaging to produce more satisfying results bothered interviewees. I found it difficult, though, to explain what I had done concisely and avoiding technical language. While some said for research purposes it would be important to know the search criteria, I was often told the person did not mind as you can see if the results are 'right' (“I think it's quite self-explanatory” (HoD S)), but they would worry if there were obvious errors appearing in the visualisation.

While prototyping the design in conversation with curators at Cooper Hewitt, we discussed the possibility of different versions of this tool: offering more control in search and not massaging parameters for in-depth researchers. But there is also value in visually satisfying results and my editing was considered effective. As a curator expressed it:

I have a large number of professional designers and design students who come here... scholarly research is [can be] visual research. Just seeing beautiful examples of how people have used particular color schemes is research. The visual, the satisfying...seeing the most compelling works has a value as well. For the professional designer to say, 'Wow. This is really incredible, the use of this color scheme. I wanna share this with my students' (Cooper Hewitt curator, 2018).

Textile designers feedback session

I was also interested in how the tool might serve as a design/creative resource, for reference and inspiration. I arranged a feedback session with textile design students (Masters students in Textiles from the Royal College of Art). I chose to approach textile students because many of the images that appear in 'Dive into Color' are of historical textiles.

This session was not very successful. It proved comparatively difficult to recruit design students. Three attended the feedback session, two of whom were not native English speakers and communication was not completely fluid. The feedback I received from the students mirrored descriptions of the tool as playful and supporting serendipitous discovery, but I did not collect new insights about the tool in this different use case. On reflection, I think my prepared questions and outlook were too rooted in viewing the collection material from a historical point of view; my questions could have been more open or better tailored to the different goals of browsing for design inspiration.

Summary

Building on 'Tag-timeline' in the previous chapter, I designed a timeline visualisation—'Dive into Color'—offering a novel implementation of colour harmony search across cultural data, and illustrating how artist colour systems can be used in visualisation as an alternative to digital colour palettes. The visualisation reveals patterns hinting at Colour History and the collection, but, playful and visually appealing, it is also very accessible for casual audiences. Designing a visualisation to trace fuzzy, subjective qualities in data known to contain errors, I was more comfortable using my own judgement of what 'looked better' or 'right' to guide my decision-making; the result was judged by the museum curators as effective for exhibition purposes. While the visualisation is complex in terms of the data delivery mechanism, the immediate visual quality of colour helps indicate when results in the visualisation may be an error (though it will not highlight missing results).

Exploratory visualisation

Royal Photographic Society collection, Victoria and Albert Museum, London

Following the 'Dive into Color' project, I spent 3 months on a research placement within the Victoria and Albert (V&A) Museum, London, sitting in with the museum's Digital Media team. The brief was to work on an independent timeline visualisation project. I found the placement unrestrictive and I was free to explore with prototyping.

Particular digitised collections within the museum were highlighted as of potential interest to work with, including the Royal Photographic Society (RPS) collection (V&A Museum, 2019a) which I chose to work with. The RPS collection had recently been transferred to the V&A from the Science Museum Group and was in the process of being digitised: photographed and digital cataloguing. The collection showcases the history of the art of photography and includes over 270,000 photographs. The photographs are of varying forms (small and large prints, glass slides etc.) and were made using a range of techniques, including early experiments in the development of photographic technology. It is a diverse collection (for example, see Figures 138 to 143) including photographs made by important early pioneers as well as plentiful, widely-produced types (Royal Photographic Society, 2017).



Figure 138. Samuel Bourne. 1866. 'Evening in the Mountains'.
Museum number: RPS.347:57-2017.
© Victoria and Albert Museum, London.



Figure 139. George Washington Wilson. Before 1891. 'Rotten Row'.
Museum number: RPS.1925-2018.
© Victoria and Albert Museum, London.



Figure 140. Camille-Léon-Louis Silvy. ca. 1860s. 'Woman in a dress'.
Museum number: RPS.1728-2018.
© Victoria and Albert Museum, London.



Figure 141. Lewis Carroll. c. 1876. Xie (Alexandra) Kitchin. Museum number: RPS.2246-2017. © Victoria and Albert Museum, London.




Figure 142. Paul Martin. 1970-2000. 'Trafalgar Square at Dusk'. Museum number: RPS.5176-2018. © Victoria and Albert Museum, London.



Figure 143. Glass lantern slide. Museum number: RPS.1876-2018. © Victoria and Albert Museum, London.

At the time of my placement, the collection was only at the beginning of digitisation. Images and cataloguing information were being added daily. This digitisation was happening alongside preparations for the opening of a new photography gallery within the museum (Victoria and Albert Museum, 2018), which would showcase RPS items alongside items from the museum's wider photography holdings.

During the placement, I shared and discussed prototypes with curatorial staff working on RPS digitisation and Digital Media staff. I conducted three audio-recorded interviews: two informally, mid-prototyping, with a cataloguer and with a Digital Media staff member, and a further more formal semi-structured interview with a curator after the end of the placement (of roughly 30 mins). I presented work-in-progress visualisations in wider meetings within the museum around collection data practice and the new photography gallery, which were also opportunities for feedback.

Prototyping

My initial steps were to familiarise myself with and explore the affordances of this dataset. I made rough plots laying out the collection data to get a better sense of what the collection consists of, how it distributes in time, and what attributes might be interesting to visualise. (I discuss this approach in my Methods Chapter: see 'Cultural visualisation prototyping'). I was working with a static copy of the RPS collection data (an Elasticsearch index snapshot as of June 2019: 5,472 digital records—about 2% of the physical collection—of which 2,155 had an image and a date). This data was a copy of the records available on the museum's website, as opposed to the internal CMS (collection management system). These are not identical.

The RPS collection showcases artistic photography and I was interested in highlighting the visual content, so again opted for plotting the images themselves as data points. I started by roughly plotting all the data by date. In the V&A collection, date information for a record is available as a date span in ISO 8601 format (International Organization for Standardization, 2019); the museum's internal CMS translates manually-input free-text dates to date spans using a formula. I chose to plot the 'photographed' date here (the date a print was made can be different).

Similarly to my starting point at Cooper Hewitt, I started by plotting each item at a random position along its date span. To spread the data apart vertically, I plotted items at a random vertical position within bounds.

While the result (see Figure 144) is an expected mess, there is some sense of the photographs' distribution in time and columns of similar images reveal photograph sets (see Figure 145). I was told that photographs in albums are catalogued in a group record in the internal CMS, but the items are published as individual records in the web data version.




Figure 144. RPS data plotted horizontally by date, and at a random vertical position.

Sets of photographs in the RPS collection can be identified in this crude plot: left column is the photographs album ‘The Royal Visit to Gwalior’ 1905; right column is the lantern slides for ‘Mount Everest expedition of 1921’ and ‘Mount Everest expedition of 1922’.

How to pull the data apart? My first thoughts were to explore segmenting the data by photographic technique (examples in the RPS dataset included ‘albumen process’, ‘daguerreotype’ and ‘heliography’). Different techniques were introduced, and were popular at different times (V&A Museum, 2019b) and I was interested to see if mapping the data this way would show popularity shifting over time from one technique to another.

I mapped the data in horizontal bands for each technique (see Figure 146). I scaled the height of the bands with a power scale (D3 API, 2019b) to accommodate imbalance in the number of items for different techniques.

This plot gives a sense of the different aesthetics and image content typical for techniques, as well as hinting at the periods when techniques were used (see Figure 147).

The main takeaway from this visualisation, though, is that it betrays the decisions made about what to prioritise for digitisation. In a collection where only 2% had so far been digitised, these biases are inevitable. The visualisation’s shape is dominated by albumen prints and daguerreotypes (the top 2 bands), which had been digitised first. Dense isolated clusters elsewhere often represent photographs by key figures in the early history of photography, such as Julia Margaret Cameron (see Figure 148).




Figure 145. RPS data plotted horizontally by date, and at a random vertical position.




Figure 146. RPS data plotted horizontally by date, and vertically by photographic technique (techniques ordered from most to least items).

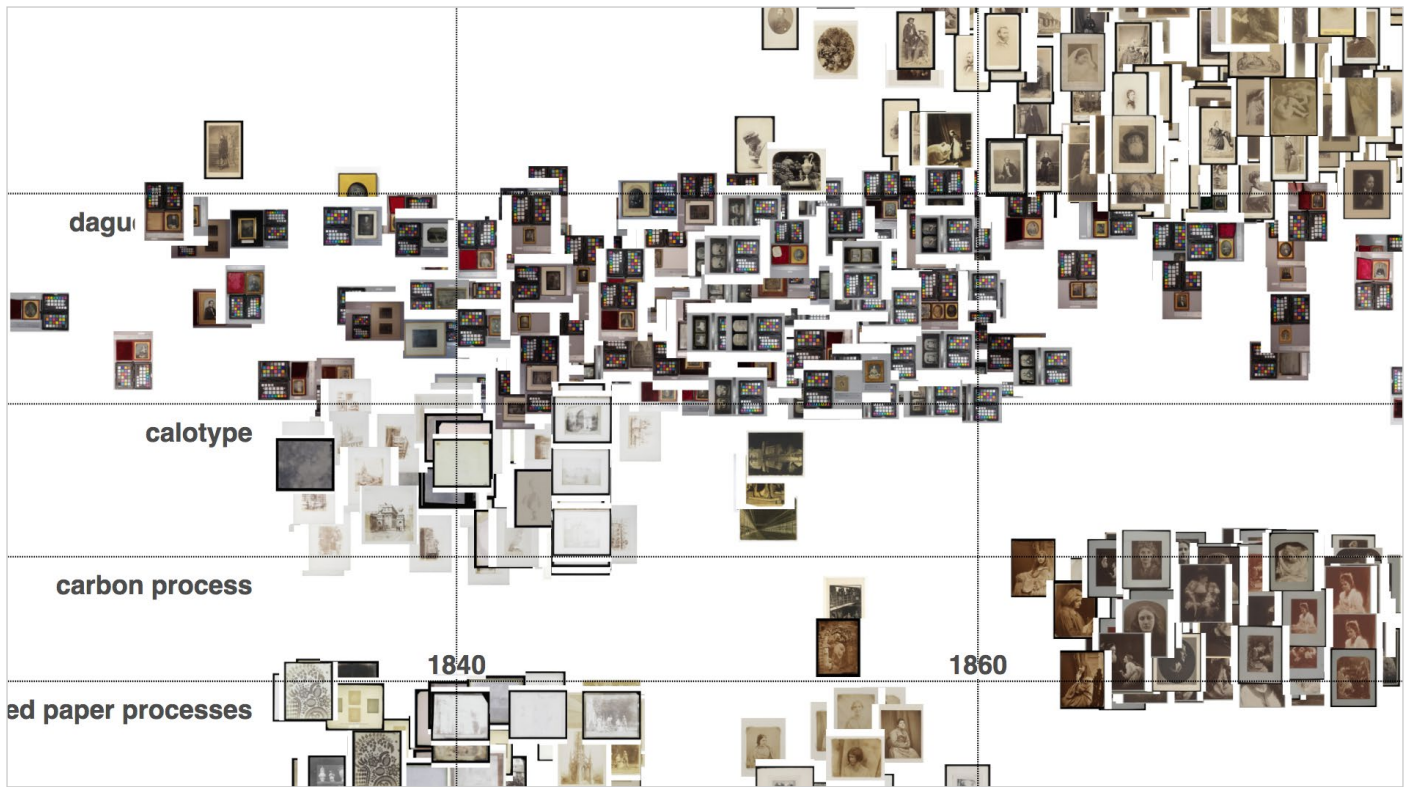


Figure 147. Detail from RPS data plotted by time and technique: different techniques have different tone and image content.



Figure 148. Detail from RPS data plotted by time and technique: Julia Margaret Cameron carbon process photographs.

Dialogue with the museum revealed a convergence of practical reasons and wider institutional factors—largely the opening of the new photography gallery and associated editorial content—behind these biases. “What you see here is probably the result of some curatorial decisions either about the narrative of the actual photography centre spaces or around particular stories we’re telling within it like [the creative/making] process...a very V&A...story” (V&A Digital Media staff, 2019). The bias towards digitising earlier items was explained because RPS items were planned to play a large role in telling the early history of photography in the photography gallery exhibition.

Aside from this, the basic plot is useful for spotting errors in the data. In Figure 149 an isolated albumen print (labelled 3) and calotype negative (2) are both positioned too late because of errors translating their textual dates to date-spans. However, not all anomalies are errors: the isolated carbon print (1) is correctly positioned for its record. Further digitisation may fill in the gap.

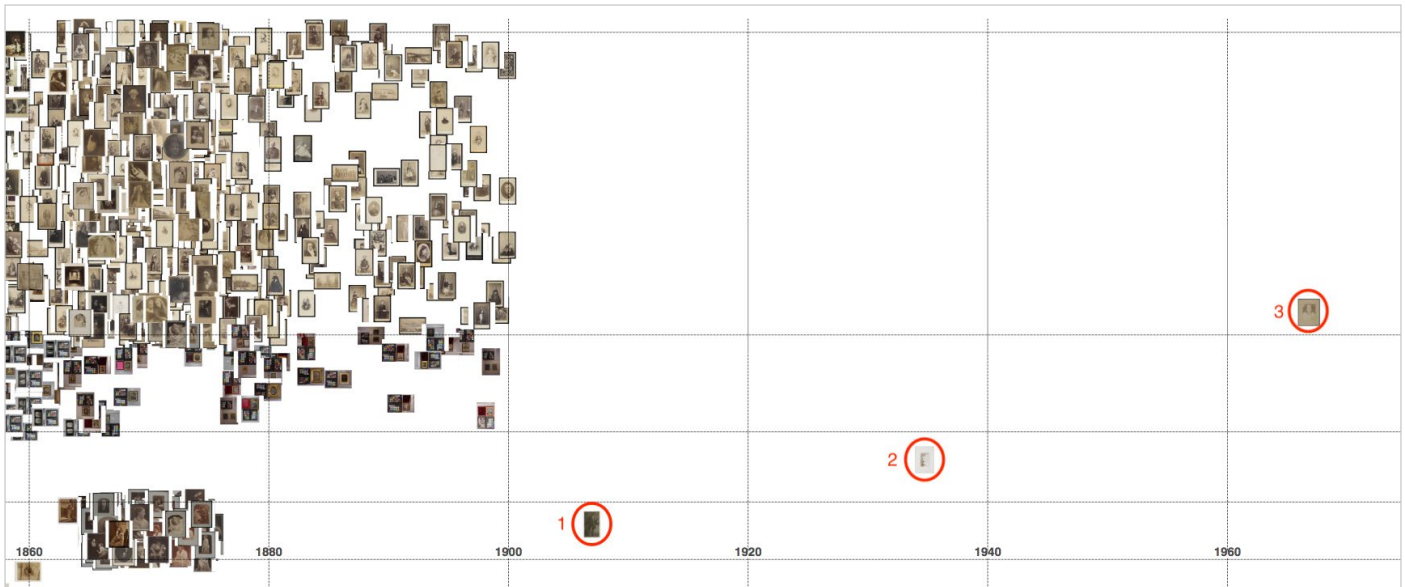


Figure 149. Annotated detail from RPS data plotted by time and technique showing anomalous data points.

Similarly, the distribution of the daguerreotypes (another early photography technique) betrays an error with dates. Their distribution is most dense 1840–60, with a sparser scattering either side (see Figure 150). Checking the records reveals many of the items are dated “mid-19th century”. The script translating this to a date span had tripped on the hyphen, converting the span to the full century: 1900–1999. An RPS cataloguer speculated these digital records may have been inherited from the Science Museum Group as this free-text date form is different from V&A convention, hence why the CMS script was ill-prepared to translate it.



Figure 150. Detail from RPS data plotted by date and photographic technique: daguerreotypes.

This technique for spotting errors has a shortcoming though. If an item’s date span has been erroneously widened, plotting it at a random position within its span will only sometimes position it outside its correct date.

I next tried segmenting the data by ‘category’ (see Figure 151). ‘Category’ is used in the V&A data both to identify themes (for example ‘Landscapes’, ‘Children and Childhood’ or ‘Military’) and for institutional grouping ‘photography’ or ‘Royal Photographic Society’.

Discussing this visualisation with a V&A curator revealed the institution’s cataloguing approach is changing. Thematic categories are actually new at the V&A, they have only existed for the last 6 years (V&A curator, 2019). Partly, as I noted earlier, this responds to the changing questions that audiences approach the curators with:

20 years ago they’d [researchers] come in and say, ‘I want to see works by this artist’, ‘I want to see works by Constable’... Now, they’re more likely to come in and say, ‘I’m researching early feminism’ or ‘I’m researching black British history’. It’s much more thematic and much more social history based perhaps. So they’re looking for specific imagery that we

didn't necessarily think to record at the time. We were more focused on who painted it or who took the photo (V&A curator, 2019).

It is difficult to adapt existing collection data that was not catalogued to accommodate thematic searches: "historically we haven't made it easy [to explore by theme]...all our sources are geared towards finding the artists" (V&A curator, 2019). Further, the way curatorial work is divided up across the institution perhaps does not lend itself to this kind of data creation: "if you think about how these categories relate to our curators and curatorial divisions [hierarchies within museum departments]...You've got curators who look after areas of stuff. They don't look after themes... You don't have a curator here of Transport [for example]" (V&A digital media staff, 2018).

These new categories also speak to the changing ways that the digitised collection now serves the wider digital experience and editorial content on the museum's website: "these meaty themes...[get] people into the content and the collections" (V&A digital media staff, 2018). For the cataloguers though, compared to recording the maker, dimensions, or material of an item, applying thematic categories can feel less rigorous: "it's something that you just feel like yeah, that fits...It's not...possibly as regimented as it could be" (V&A curator, 2018).

In this visualisation, I only looked at the thematic 'categories', manually excluding the others.




Figure 151. RPS data plotted horizontally by date and vertically by 'category', excluding the categories 'photography' and 'RPS'.

Multiple 'categories' can be applied to any record. At this stage of making rough plots, I chose simply to pull the data apart as much as possible employing a rule that an item will be plotted in the 'category' with the lowest total items—in effect, pulling all the data down as much as possible. (Actually, multiple photographic 'techniques' can also be applied, but with this dataset, only 2 records so far had more than 1 'technique'). If taking these visualisations further, a filtering approach (eg. Urban Complexity Lab, 2015) might be a better way to deal with visualising non-mutually exclusive facets.

Mostly this visualisation revealed the many portraits, the most widely applied category, in the data (see Figure 152).



Figure 152. Detail from RPS data plotted by date and category: 'portraits'.

Some of the other 'categories' were beginning to make compelling groups around themes (see Figure 153), but had not been so widely applied.

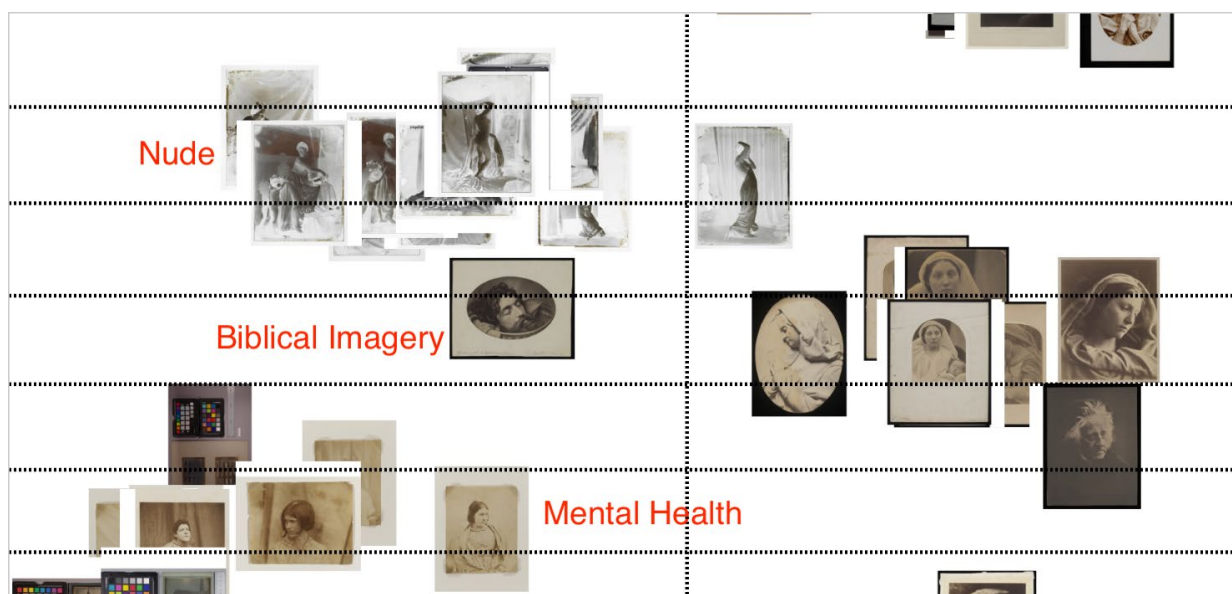


Figure 153. Annotated detail from RPS data plotted by date and 'category'.

My experience with the RPS data argues for the value of taking multiple perspectives in exploratory visualisation as a way to pick up on errors and inconsistencies. For example, further visualisations of the RPS data plotted by time and photographer showed that photographs by one early photographer, Linnaeus Tripe, were not visible. Following this up revealed the Tripe records' data was structured slightly differently to other records; this had not been noticed up to this point. Organising the data by photographer also revealed there were two authorities in use for pioneering photographer Julia Margaret Cameron (the other being 'Cameron, Julia Margaret'), which split her photographs into separate groups in the visualisation.

Discussing these plots with a curator, however, raised concerns around forming conclusions about the collection from them. Firstly, the curator felt the visualisation conceals the data's origins, which are especially pertinent given the digitisation project was in-progress:

I wouldn't necessarily trust information displayed like this in the same way as I would trust looking at our catalogue...in our catalogue I can see the data behind it...I can also see who put it in and when. And...what else is in [the] record. Whether it was a quick record that we did just to have a record when we digitised the object or whether it was something that somebody really spent time and expertise doing (V&A curator, 2019).

Further, by relying solely on the numerical date information attached to records, the visualisation can be lacking in terms of temporal ordering: "I might not know the exact date of a photograph but I might at least know that it happened before or after another photograph was taken" (V&A curator, 2019).

These rough, unpolished plots can help with spotting data errors and offer an overview of the unfolding digitisation project. If the shape of the dataset is unrepresentative of the collection, however, can a visualisation offer a useful structure for sensemaking without spotlighting the overall shape? Further, some of the more plentiful photography types in the RPS data overwhelm other content, for example there are many *carte de visites* (a type of small photograph popular in the late 19th century, see Figure 154). How might I design a visualisation that gives a taste of items like these without drowning out other context? As the images are a fundamental part of the collection, I was also interested to explore other visualisation approaches combining focus and context (Cockburn, Karlson & Bederson, 2008) enabling the images to be inspected while surveying a wide temporal range.



Figure 154. Detail from RPS data visualised by date and technique: lots of *carte de visites*.

Image similarity across the RPS collection

While the 'category' facet had not been widely applied, an alternative way to trace connections across the collection is by visual characteristics; this is not possible by relying solely on cataloguing information. One way to explore an image collection visually is to extract feature vectors from the images. Feature vectors are (multi-dimensional) numerical representations of images describing their visual characteristics. By comparing feature vectors, it is possible to quantify the visual similarity and dissimilarity between images. There has been increasing interest in the application of feature vectors (following recent progress in the field of Machine Learning) to support search in cultural image collections (Pim, 2018; Yale University Library Digital Humanities Lab, 2017). There are also artistic cultural visualisation examples (eg. Nasjonalmuseet, 2017a; 2017b and Diagne, Barradeau & Doury, 2018) built off these techniques.

I extracted feature vectors from the RPS images using the Keras (2019) library written in Python. I ran the RPS images through a pre-trained VGG16 model (Simonyan & Zisserman, 2015), a convolutional neural network—a type of machine learning model particularly suited to image analysis—and extracted a feature vector for each item. (In this case the feature corresponds to the penultimate layer of the network). The model was already trained on the Imagenet (2019) database (over 14 million modern photographs) and I did not do any further retraining. There are other pre-trained models available but, as the results for the RPS images seemed fairly successful, I did not test results for other models.

The feature vectors VGG16 outputs have 4,096 dimensions, as if measuring 4,096 characteristics of each image. To speed up computation over them, I reduced down the dimensions to 300 using a statistical process called Principal Component Analysis (scikit-learn, 2019a), that aims to bring out patterns in data by prioritising the most significant differences.

To illustrate the results, the examples that follow show a chosen (larger, top) image with the 5 images (below) that the model considers most visually similar. Visual similarity is measured by comparing the distance—using the Cosine distance measure (SciPy, 2019)—between feature vectors.

From its training the machine-learning model has learnt to recognise characteristics of images that are useful for numerous image-processing tasks. The feature vectors, therefore, capture a variety of visual characteristics. So when I say ‘visual similarity’, we see from the results that this describes an entanglement of different characteristics including content, texture, tone and composition.

Some of the results are more obvious (see Figure 155 to 160): duplicates, sets of photographs from an album, and those by an individual photographer with a distinctive style.




Figure 155. Reverse image search results from the RPS collection. Julia Margaret Cameron portraits of women.




Figure 156. Reverse image search results from the RPS collection. Julia Margaret Cameron portraits.




Figure 157. Reverse image search results from the RPS collection. Magic lantern slides from Mount Everest expedition of 1921 and Mount Everest expedition of 1922.




Figure 158. Reverse image search results from the RPS collection. Photographs from The Royal Visit to Gwalior album, 1905




Figure 159. Reverse image search results from the RPS collection. Salted paper prints by William Henry Fox Talbot, all pages from the same album 1844.




Figure 160. Reverse image search results from the RPS collection. Stereoscope daguerreotypes of displays at the Crystal Palace, Sydenham c.1855.

Whereas in the examples above the similarity test produces quite predictable results, in other cases the model cuts across the collection—particularly when searching across popular photograph types.

The chosen image in Figure 161 is a full-length portrait carte de visite. The search results are more standing portrait cartes de visite, but taken at different photography studios at different times in different places. Searching across the collection, the model picks up on conventions for staging portraits over this period.

An image's structure can be the basis for similarity: Figure 164 shows photographs all mounted in a circle/oval.




Figure 161. Reverse image search results from the RPS collection. Cartes de visite: full-length, standing figures.




Figure 162. Reverse image search results from the RPS collection. Cartes de visite: (mostly) sitting men.




Figure 163. Reverse image search results from the RPS collection. Cartes de visite: head and shoulders portraits.




Figure 164. Reverse image search results from the RPS collection. Photographs of people mounted in a circle/oval, including [middle] *The Head of John the Baptist on a Charger* by Oscar Gustav Rejlander 1857.

As can the image content: Figure 165 shows landscapes with trees.




Figure 165. Reverse image search results from the RPS collection. Landscapes with trees: more cartes de visite, [middle] Agnes Warburg’s coloured Avenue of Trees, c.1920, and [right of middle] a lantern slide from Mount Everest expedition of 1921.



Figure 166. Reverse image search results from the RPS collection. Photographs of people, some in pairs, inside or by buildings.



Figure 167. (left) Reverse image search results from the RPS collection. Landscapes, including photographs [top, left, right of left] by James Ravilious 1970, and [right] Olives and almonds by Agnes Warburg 1930.

Of course, when there are no similar images in the collection, the model will fail to retrieve any (see Figure 168—though conceivably, the image of lace [bottom right] has a similar visual texture to the water beads photograph).

Neural networks are notorious for being difficult to interpret the way they work, though their interpretability is a growing area of research (Olah et al., 2018). In Figure 169, the chosen image is a glass slide and 3 other slides from the same set are considered similar. But there are also 2 not-obviously similar images, other than they all have borders. It is unclear why these have been retrieved. This emphasises the difficulties of working with processes whose inner workings are not accessible.




Figure 168. Reverse image search results from the RPS collection. Mixed images, [top] Goutte de rosée sur une feuille de capucine by Guyula Hálasz Brassai, c. 1935 and [right] Lace by William Henry Fox Talbot, early 1840s.




Figure 169. Reverse image search results from the RPS collection.

To overview how the model organises the collection, I reduced the feature vector dimensions to 2 (using the ‘t-SNE’ algorithm (Maaten & Hinton, 2008; scikit-learn, 2019b)), losing some of the detail but allowing me to plot the results (see Figure 170).



Figure 170. ‘t-SNE’ plot of the RPS data.

Direction has no meaning in this plot, but images considered similar will be close together. The layout looks slightly different every time the algorithm is run because there is an element of randomness in how the algorithm works.

Before dissecting this plot, what are the two clusters off on the left? Those are the daguerreotypes (a type of early photographic process) which, in the RPS data, happen to be photographed with a colour chart. Machine learning models like this can pick up on this kind of accidental visual distinction. Because of the colour charts, the model considers the daguerreotypes alien from the the other images. And it further splits them into 2 distinct groups—colour chart placed to the side (Figure 171), or above/below (Figure 172): a meaningless difference (unless one is interested in colour calibration techniques). In further work, it would be worth cropping out the colour charts to avoid this.

Interestingly, when I asked why only daguerreotypes had been photographed with colour charts, the cataloguer was surprised to see them: “[the V&A museum] stopped doing the color chart I don’t know how many years ago...So they [daguerreotype photographs that feature a colour chart] must’ve been done previous to that time... [or] I wonder if they’re images that we’ve taken from Science Museum?” (V&A curator, 2018). This example highlights how the institution’s identity shapes even how photographs are taken in the digitisation process; as an art and design museum

there is a “tension between object photography and beautiful aesthetic photography” (V&A curator, 2018) in documenting collection items.



Figure 171. Unknown photographer, St. Paul's Cathedral, daguerreotype, mid-19th Century. Museum number: RPS.518-2017. © Victoria and Albert Museum, London.




Figure 172. Unknown photographer, Portrait of a Man and Boy, daguerreotype, mid-19th Century. Museum number: RPS.409-2017. © Victoria and Albert Museum, London.

Other than this accident of clustering the items with colour charts, there seems to be some logic in how the model clusters the images (see annotations in Figure 173). This technique offers understandable, visual paths through the collection that were not previously possible with cataloguing data alone.

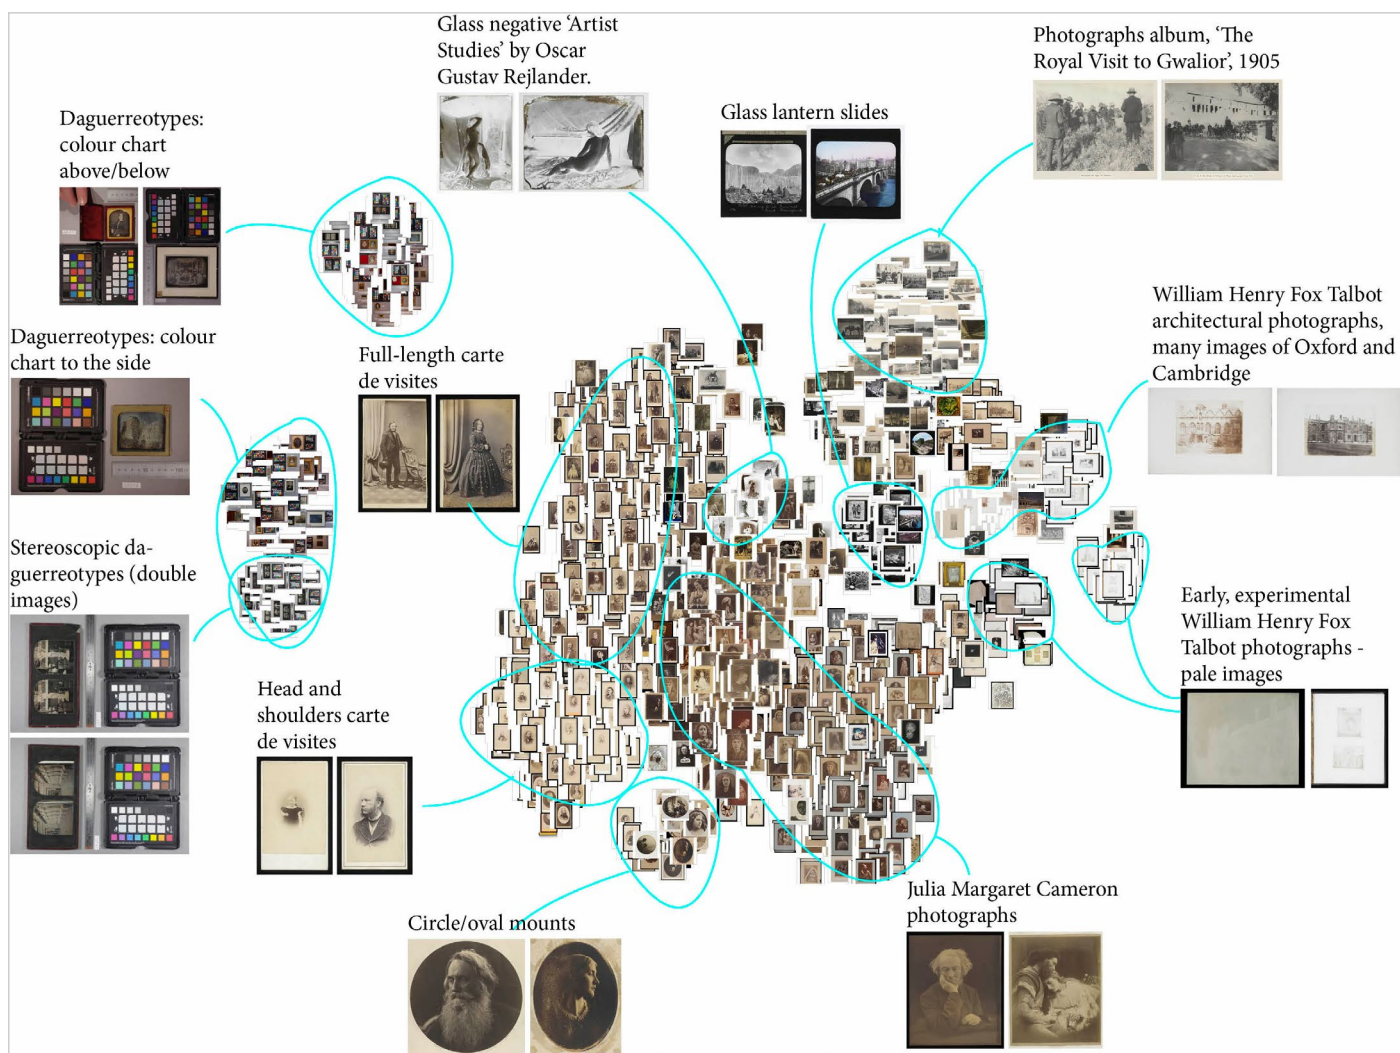


Figure 173. Annotated 't-sne' plot of the RPS data.

‘Time-tiers’ visualisation

How might these visual similarities across the collection be explored in timeline visualisation? In addition, I was eager to build off the insights made in my earlier exploratory visualisation. One approach to visualising connections across a collection without spotlighting the overall collection shape is sampling (previously raised in the ‘Cultural timeline visualisations—Survey’ chapter). Further, my exploratory visualisation had demonstrated that very plentiful items in the collection, for example the cartes de visites, were drowning other content. Since photographic types are often constrained to a time period, sampling across the full timespan would be a way to showcase diverse items from across the collection.

Illustrating this idea, Figure 174 shows a mocked up visualisation design: around a central, selected item there are three tiers: each is a timeline covering a, from top to bottom, widening time period extent around the selected item’s date. A sample of visually similar images are plotted along these timelines; images are displayed at a large enough size to make out details. The purpose of these multiple timelines is to combine focus (showing items from a similar time) and context (showing items from across the whole collection period). It is an alternative approach to, say, a fisheye view (Sarkar & Brown, 1992), but the timelines all have a uniform, linear scale.




Figure 174. Mockup of tiered timeline visualisation.

The aim here is to offer temporal context to a selected item at various scales, and without spotlighting the overall dataset shape which is skewed by digitisation priorities.

To test this idea, I first built an interactive prototype with three tiers (see Figure 175). In this design, a selected item is displayed in the centre at a larger size. The selected item is assigned a random date within its date span; the top two timelines have extents 10 years and 50 years respectively around this date. The bottom and widest timeline has the full RPS collection temporal extent.

Visually similar items across the collection are plotted across these timelines. Items are plotted on the timelines, prioritising the most visually similar (according to the extracted feature vectors), and prioritising placing items on the top timeline and then the middle timeline if their date span overlaps with these. Up to a maximum of 2 items can be placed on the top tier, 4 for the middle, and 8 on the bottom tier. I put a minimum similarity measure threshold on visualised items, so sometimes few items are visualised. Items are positioned on the timelines at a random point within their date spans, constrained by the timeline extents. If items overlap they may not be visible,

though hovering over an image brings it to the top of any overlap and reveals a tooltip with the item's title.




Figure 175. Prototype visualisation of RPS collection data.

A new central item can be selected, visualising new contextual timelines, by refreshing the interface which selects a random item or selecting one of the smaller items. Clicking on the central selected item opens its record on the V&A collection website in a new tab.

I found the three tier design difficult to interpret; content does not neatly divide in an obviously meaningful way between the top two lines. I, therefore, adapted the design to just two tiers: the top covering 10 years, and the other covering the full collection time span. I also added serpentine arcs as a minimal way to indicate where the top timeline maps to the bottom. Examples show the visualisation contextualises selected items with visually similar items in the collection through time: by content, trees, Figure 176, or full-length portraits, Figure 177; by pose, Figure 178; or by tone and image structure, Figure 179.

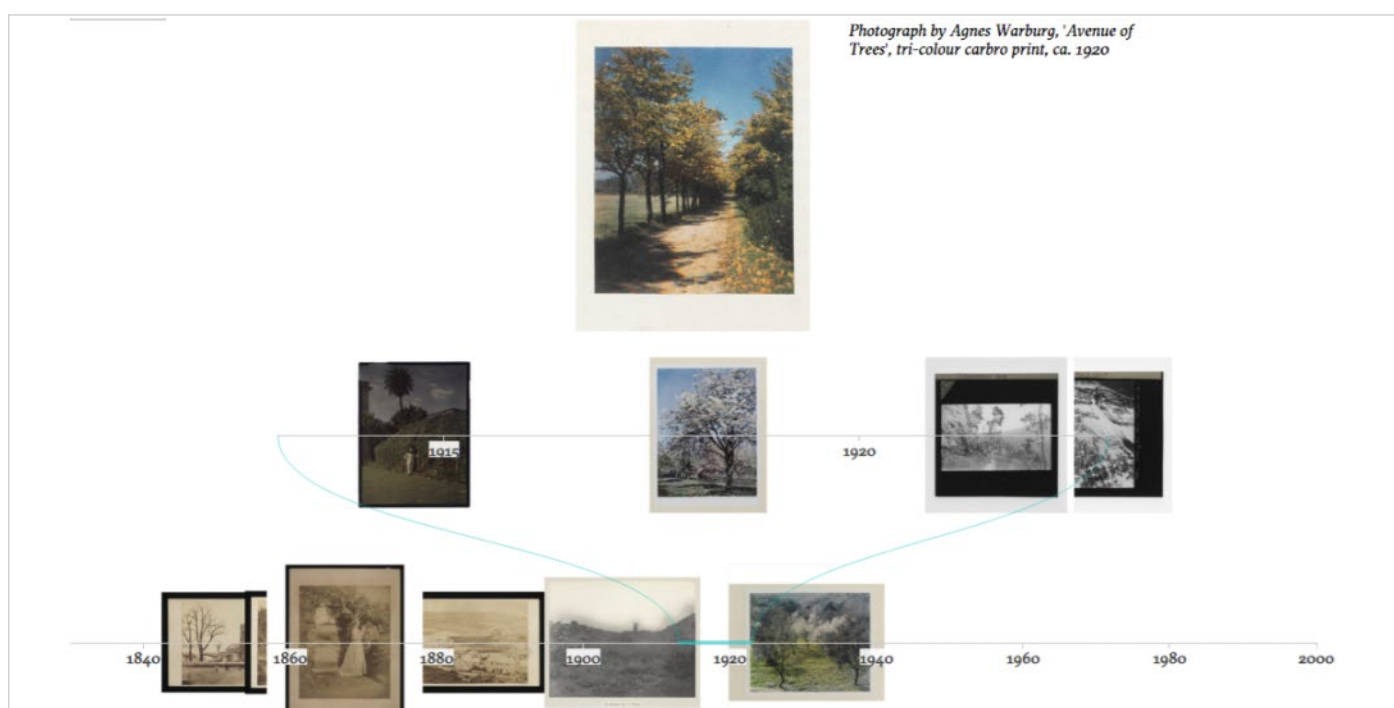


Figure 176. Prototype visualisation of RPS collection data.



Figure 177. Prototype visualisation of RPS collection data.



Figure 178. Prototype visualisation of RPS collection data.



Figure 179. Prototype visualisation of RPS collection data.

The visualisation sometimes reveals when there are multiple very similar items (see Figure 180), for example if multiple prints have been made from the same negative. Placing items at a random position in their timespan, however, means items may, confusingly, be positioned at different points even if they were produced at the same, but uncertain, time.



Figure 180. Prototype visualisation of RPS collection data.

Prioritising items only by visual similarity means sometimes the other visualised items are simply others from a set (see Figure 181 and Figure 182). In these cases, the visualisation fails to show wider context. This could be improved by sampling over a set of time spans making up the full timeline extent, rather than just the most visually similar across the whole collection.

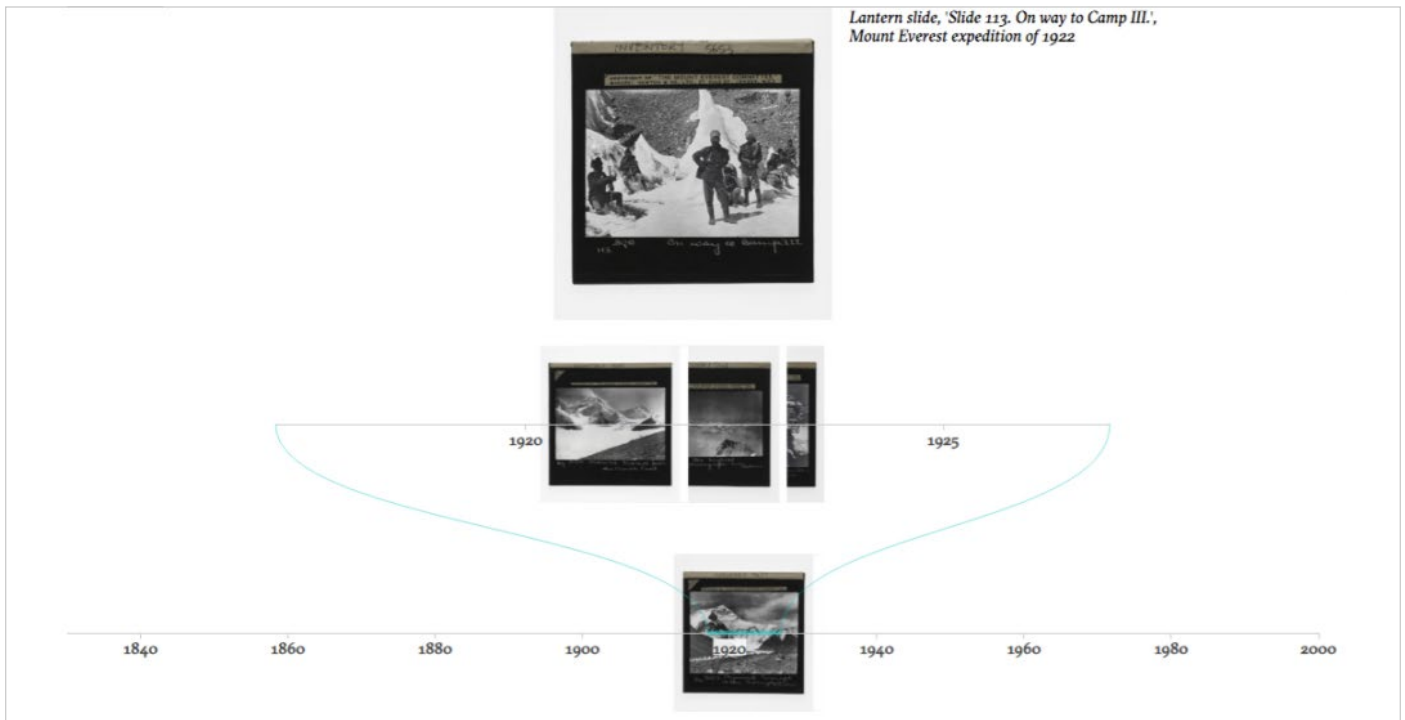


Figure 181. Prototype visualisation of RPS collection data.

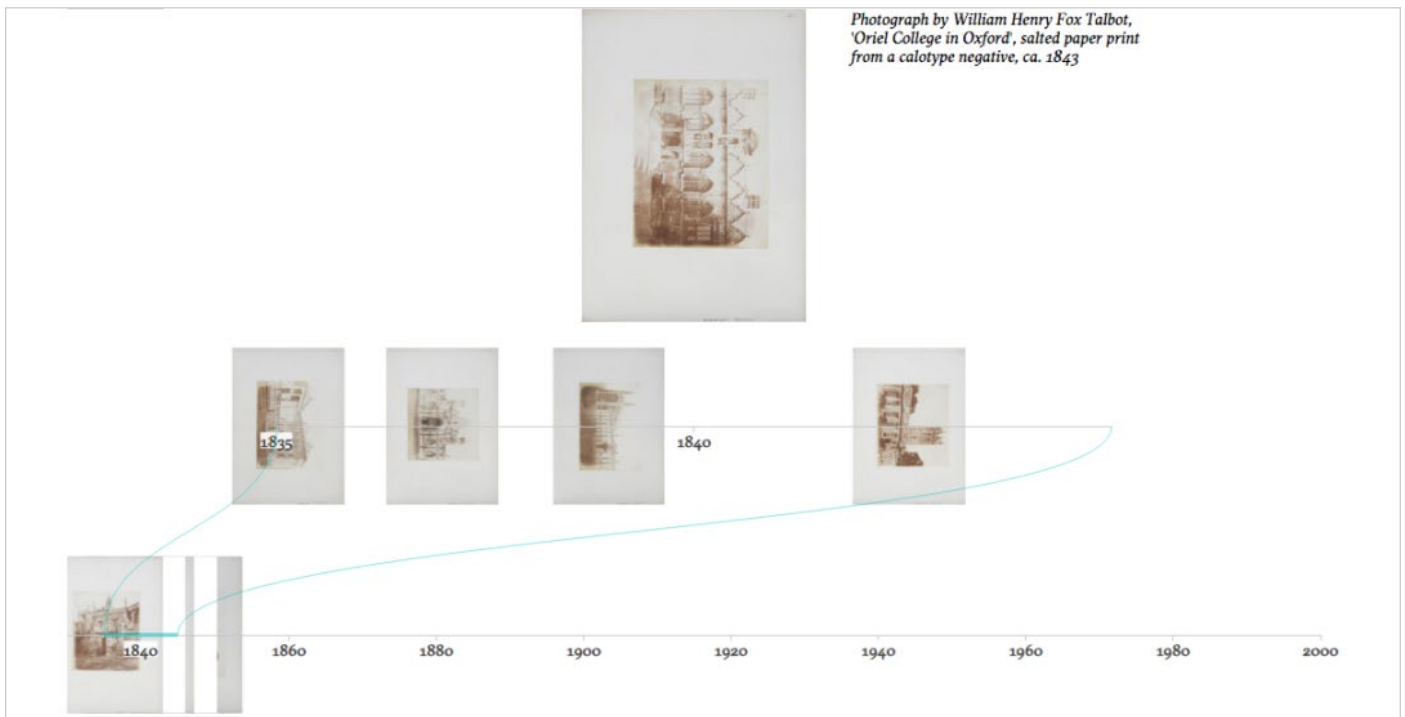


Figure 182. Prototype visualisation of RPS collection data.

The visualisation also reveals where the connections made using the image similarity data are naive. For example, in Figure 183, the visualisation is centred around an image of blotchy paper, which checking the V&A record (V&A Collections, 2019) turns out to be the back of a photograph. Using the similarity data, the visualisation connects this with photographs of grass. While the blotchy paper is not an image of grass, it could be said to be texturally similar? Is this an interesting, lateral connection or confusing and unhelpful?

This is a fairly crude way to contextualise items across the photography collection; the feature vectors tangle up a variety of different visual characteristics across the images. This project work, though, gestures to what may be possible in the future if developments in techniques make it

possible, for example, to specify what kinds of visual characteristics are of interest; the resulting visualisation would likely be more meaningful and more useful. A V&A curator, for example, speculated how useful it would be to be able to specify searching for all studio photographs with the same painted background used, as an indicator that they were taken at the same studio (this is not always known), or to trace photographs that feature the same person.



Figure 183. Prototype visualisation of RPS collection data.

This project was truncated by the end of my placement at the V&A, though I was able to continue working on these ideas and designs with a different collection, which I discuss in the next chapter.

Summary

This chapter began with exploratory visualisation of the Royal Photographic Society (RPS) collection data at the V&A Museum. Visualisation, in dialogue with collection experts at the V&A Museum, pointed to the shape of digitisation progress, cataloguing errors, the consequences of sharing data between different cultural institutions, and changes to what item characteristics are now being recorded in the V&A catalogue. As a way to cut across the collection data by visual content—which is not possible with the cataloguing data alone—I explored the potential of computationally extracting visual similarity measures across these images. While this approach is somewhat successful, it pays attention to non-meaningful visual distinctions such as the presence of a colour chart, and generally raises interpretability issues. Finally, I introduced a novel timeline design, inspired by insights made exploring the RPS data, for browsing collection data combining focus and context with multiple tiers.

'Faces of Sweden'

Nordic Museum, Stockholm

Following my placement at the V&A Museum, I was approached by the Swedish National Heritage Board to work on a cultural visualisation project with Swedish Open Cultural Heritage data (SOCH) (K-samsök, 2019a). SOCH aggregates digitised cultural heritage collections across Sweden. The brief was to design a visualisation prototype geared towards casual browsing and I was free to choose what data to work with.

Portraits in the Nordic Museum collection

After exploring collection data available on SOCH, I chose to develop a visualisation around portraiture in the Nordic Museum, Stockholm (for examples, see Figures 184 to 187). I chose this data because most records had date information, and there were enough records (785 with an image/production date) for my purposes. (Initially, in fact, I was interested to work with data representing Swedish folk art—collection records had been distinguished for a Nordic Museum exhibition on this topic (DigitaltMuseum, 2019c)—but these items often had no date information other than accession). Relevant to building on the 'Time-tiers' prototype in the last chapter, most of the records had high-quality images attached and many of the images are compelling.



Figure 184. (left) Nordic Museum. Man's portrait. 1620-1660 (Estimate). Identifier: NM.0060722. (CC BY-NC-ND).

Figure 185. (right) Nordic Museum portrait. Woman's portrait. Production 1830s. Identifier: NM.0122896. (CC BY-NC-ND).



Figure 186. (left) Nordic Museum. Miniature portrait. Production 1600-1659. Identifier: NM.0300957. (CC BY-NC-ND).

Figure 187. (right) Nordic Museum portrait. Silhouette portrait. Production 1770-1820 (estimate). Identifier: NM.0046024. (CC BY-NC-ND).

Unusually for the project work in this PhD, rather than working with a standalone collection as defined by an institution, I retrieved this data via a text search query: ‘Porträtt and tavla’ against the SOCH API (K-samsök, 2019b). ‘Porträtt’ meaning portrait, and ‘tavla’ meaning framed picture. Not all portraiture in the Nordic Museum collection had been categorised as ‘Porträtt’; this is work in progress within the institution (Nordic Museum curator, 2018). Using a free-text query, therefore captures when the terms appear in free-text descriptions in the record, returning a greater number of relevant works. After data cleaning, there were 785 portraits (with an image and production date) returned for this search (see Figure 188). They are mostly oil paintings, but also prints, silhouettes, photographs etc. and date between 1570-2003.

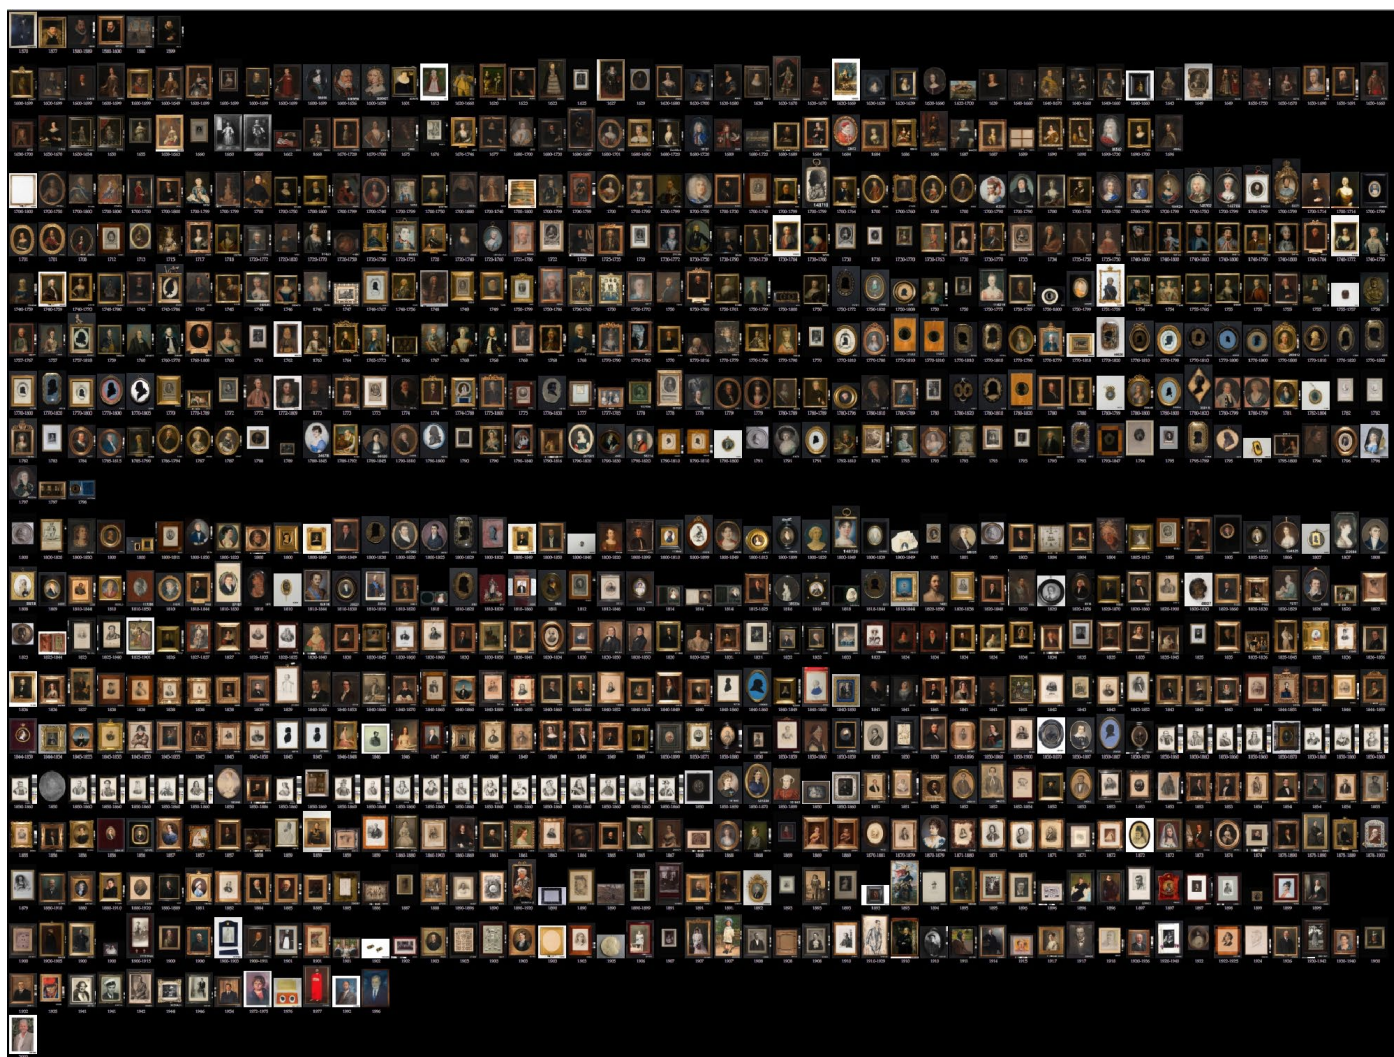


Figure 188. 785 portraits in the Nordic Museum collection, ordered by date and chunked vertically into centuries—ordered earliest to latest by (beginning of date span for) production date.

During this project I was in contact with individuals at the Swedish National Heritage Board who could advise me on SOCH data structuring and the API functionality. Midway through this visualisation project, I conducted an audio-recorded semi-structured interview (roughly 90 mins) in person with a curator at the Nordic Museum, Stockholm to discuss portraiture in the museum collection, and my developing prototype.

Portraiture in the Nordic Museum collection has been collected for its cultural heritage value—“we want the pictures to have a documentary value” (Nordic Museum curator, 2018); the paintings serve as an historical record of what people in the Nordic countries looked like, what kinds of people were depicted, and the artwork that decorated people’s homes. A visualisation tracing these items through time should, in theory, support tracing these characteristics through time. Not all

the items returned for this search, however, show a portrait in the attached image. For example, an additional frame for a portrait painting had its own record (DigitaltMuseum, 2019a). And the record image for a portrait on a ring only showed the box containing the ring (DigitaltMuseum, 2019b). Again we face the problem of cultural data as an historical vs. administrative tool. In this project, I decided to manually exclude these records, prioritising exploring the portrait images over time.

The obvious thing to say here is that SOCH data is in Swedish, which I do not speak. The data (retrieved as JSON-LD format) is in name value pairs: the 'name' in English and the 'value' in Swedish (for example "title: Tillverkare"). In many respects I found working with this data insightful, as it introduced cataloguing approaches I had not previously encountered including: recording multiple dates for one record, labelled with meta-types (such as 'create') for the event a date represents; recording a depiction date for images; and dating records by historical periods. But even with automated translation support on the web, working with this data was not as fluid an experience as with that in my native language. This contributed to my choice to work with image collection data that is relatively accessible to interpret.

Visualisation design

I adapted the 'Time-tiers' prototype, building in more interactivity to serve casual browsing and to enable close comparison between different time periods. In this visualisation—titled 'Faces of Sweden'—there is a main, central timeline covering the full dataset temporal extent. Two secondary timelines of narrower temporal extent, above and below, map onto the main one (shown by curved, connecting lines). These each display a sample of up to seven images from the selected time spans, showcasing possible items from across the time period. A user can drag and resize the coloured bars overlaying the main timeline to select different time spans.

By selecting small timespans, a user can focus on and compare different times: see Figure 189.

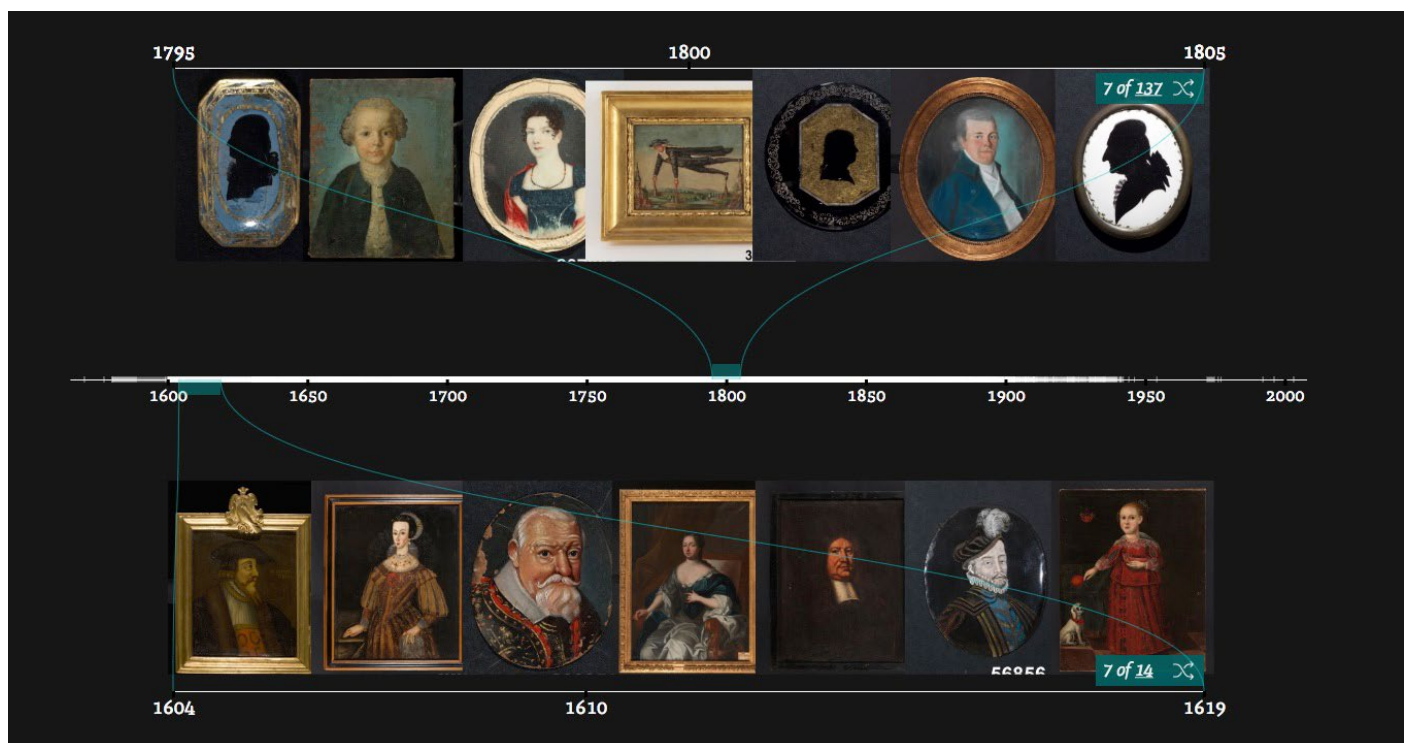


Figure 189. Secondary timelines showcasing items from around 1610, and around 1800.

Or by selecting a wide time span, they can cut across time to observe changes (see Figure 190).

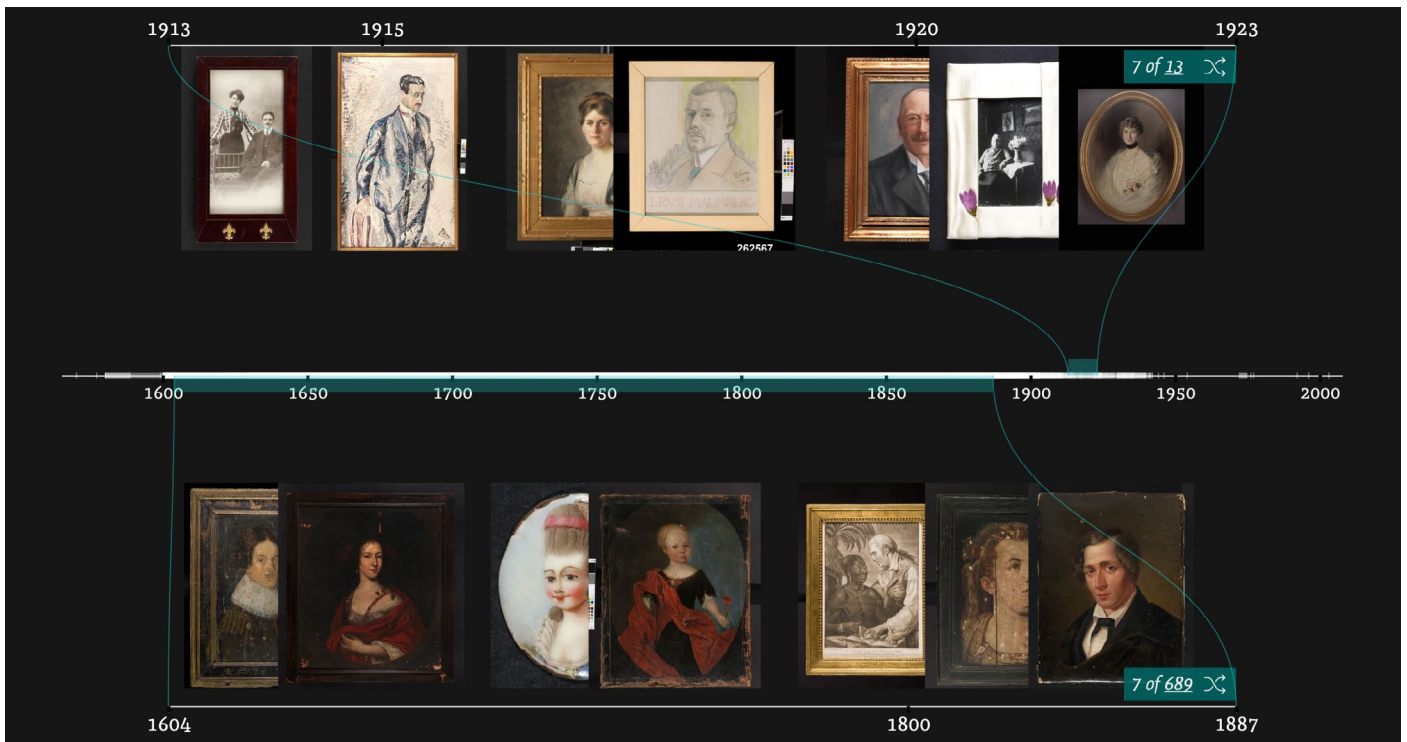


Figure 190. Bottom secondary timeline showing portraits from 1604-1887.

To even out the distribution of images across time, the sampling on the secondary timelines works by splitting the selected time span into seven smaller spans and sampling a unique image for each of these. (If I just sample seven images over the full time span, the distribution skews to when there are most portraits in the data.)

If an image has a precise date, it is positioned at that date (image centre is origin). If not, it is mapped at a position halfway along the smaller timespan it was retrieved for. This way, there is a more even spread of images (less overlap, more images visible). Hovering over images reveals a tooltip with the record title and date span.

I was interested to try different numbers of secondary timelines (see Figure 191). Mostly, I found the constraint here is the display screen size and I settled on two for desktop display.

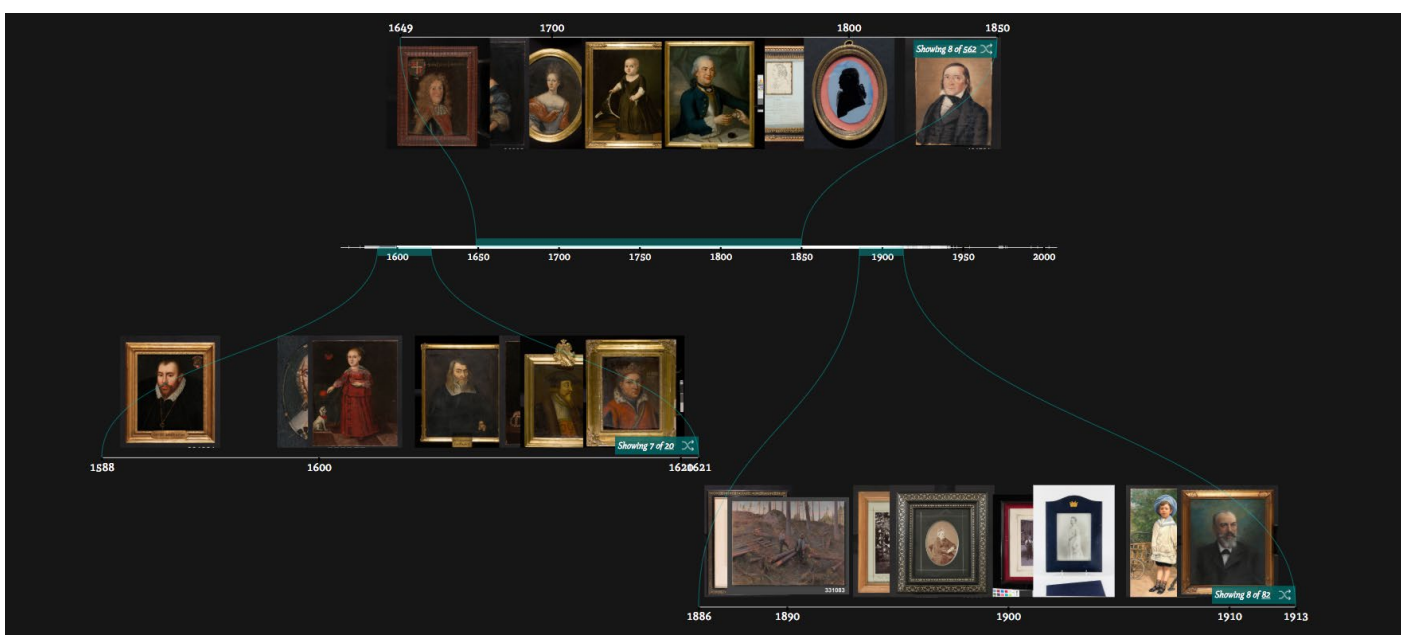


Figure 191. Additional secondary timelines.

The main (central) timeline shows the distribution of the data over time, but highlights where there are and are not records, rather than the quantitative distribution. Many items have a wide date span (some items are dated to the century). So, to indicate where portraits can be found, I draw white rectangles over the timeline for each item's date span, with greater transparency for larger time spans. The overall visual effect communicates to the user where along the timeline items can be found (see Figure 192).



Figure 192. Main timeline, showing when in time there are/are not items by the strength and thickness of the line.

Since the images visualised are a sample, an annotation on the secondary timelines indicates the proportion of items displayed (see Figure 193). Clicking the 'shuffle' (crossed arrows) icon visualises a different random sample of 7 images. Clicking the total (in Figure 193: '180') visualises all the images overlapping that time span in a tile layout below, organised from earliest to latest (see Figure 194).




Figure 193. Annotation on secondary timeline. Clicking the total, '180', generates a tile view below of all 180 portraits for this time span. Clicking the shuffle icon will display a different 7 images.

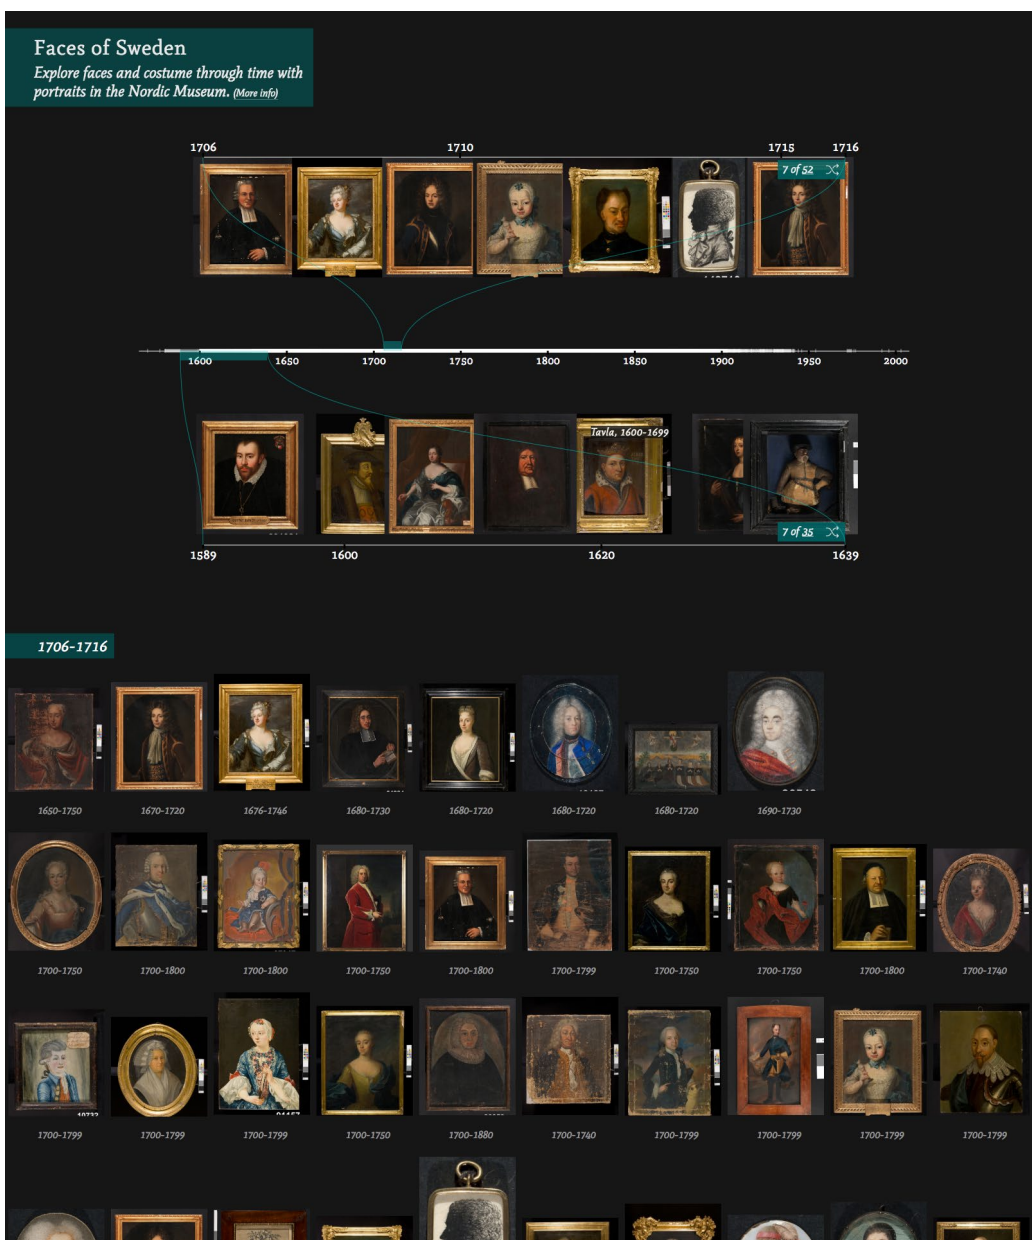


Figure 194. Tile view below timeline.

Clicking on any image, opens a lightbox viewer (powered by PhotoSwipe, 2019; see Figure 195). For this prototype, only the record's title appears here hyperlinked to the full record on the Nordic Museum collection's web presence (DigitaltMuseum, 2019d). In a future more developed version, it would be worth presenting more information from the museum record in this lightbox view for a more seamless experience.



Figure 195.
Lightbox viewer.

Discussion

The timeline design, prioritising visual engagement with the portraits, supports comparison of different time periods and there is a playful serendipity in using sampling. The immediate appearance of the design provoked a strong positive reaction from the museum curator on demonstration: “Oh, I love it! It’s wonderful!”. The curator described the design as supporting easy, fluid exploration of the dataset by time: “it’s so easy to move...around the timelines”.

As the museum’s collection is geared towards cultural history, it is perhaps no surprise these are the kinds of trends—for example, costume—hinted at in using the timeline, rather than, say, a progression of art movements. Exploring ‘Faces of Sweden’ suggests popularity of powdered wigs in the 18th Century, the changing cut and structure of women’s dresses, trends of more and less colourful clothing, changing male collar/tie styles etc. (Though, examining costume through these portraits is mostly limited to the waist/shoulders upwards).

The visualisation also highlights changing portraiture techniques represented by the museum collection through time, hinting at fashions and technology for producing images. For example, it draws attention to silhouette portraits that are common in the late 18th/early 20th Century in the Nordic Museum collection, though I noticed a pattern in that many of the silhouette portrait records (see Figure 196) are assigned ‘estimate’ dates beginning in 1770. The museum curator explained when no exact date is stated during acquisition of silhouettes like these, they are assigned the date 1770-1810 (estimate) as this is the period when most are known to have been made. The records keep this date until a curator has the opportunity to make a more precise estimate based on the portrait’s hairstyle and fashion (Nordic Museum curator, 2019). So while many of these items do not have explicit known dates, the timeline still projects knowledge about this fashion.




Figure 196. Example silhouettes in ‘Faces of Sweden’, all dated: estimate 1770-_____.

One of the use cases flagged in dialogue with the museum curator was education:

“[The timeline visualisation is] perfect if you’re working with schools systems. You can ask questions [about costume, hairstyles, pose, symbolism, how different genders are depicted etc.], and they [students] can try to look it up and get the answers through the timeline... It’s more useful that there are different techniques [being displayed], because you ask questions and get some information about what was possible to do in different periods”.

As the curator described, “the paintings bear a lot more information than the obvious”, but without explicit narration or more of the underlying cataloguing information surfaced some patterns are not apparent. For example, as the museum curator pointed out, while many of the early portraits depict nobles, some are of priests—this is not obvious to a non-expert from the images alone. In fact, a number of the most recent painted portraits are of previous directors of the museum. From the 1830s/40s, the bourgeoisie become a more important social group in Sweden and they are then represented in portraits in the collection. From this group, several of the portraits come as pairs representing married couples (see Figure 197). That some of the portraits are one of a pair, however, is concealed in the visualisation as items are sampled at item level. This problem of identifying items that ‘belong together’ is an example of how the constraints of the dataset limit visualisation design. In the case of Figure 197, neither record indicates the picture is one of a pair. And where paintings in a set have been grouped as one record (see Figure 198), the timeline is coded to only visualise the first image (as often a supplementary image is of the picture’s back). Manual editing would be required to show and group these.



Figure 197. Pair of portraits in the Nordic Museum collection. (left) Identifier: NM.0236954. (CC BY-NC-ND) (right) Identifier: NM.0236955. (CC BY-NC-ND).




Figure 198. Trio of paintings including 2 portraits catalogued in 1 record. Identifier: NM.0030050A-C. (CC BY-NC-ND).

While ‘Faces of Sweden’ can be used to explore changing historical costume, the clothes worn by the person depicted in a portrait are not always from the time the picture was produced. For example, this data includes a number of lithographs from 1850-60 depicting earlier historical figures (see Figure 199).



Figure 199. Lithographs in the Nordic Museum collection, 1850-60, depicting historical figures. (Example record identifier: NM.0269513). (All CC BY-NC-ND).

And even if the person depicted is of the time, the costume may not be. For example, Figure 200 shows a 19th Century actor dressed in Renaissance-looking costume portraying Hamlet.

I also encountered composite objects, with multiple production dates attached. For example, the frame shown in Figure 201 significantly predates the photograph. In my code I opted to preference the earliest production date in the record. Though, as that example demonstrates, this is not bulletproof for selecting the portrait ‘image’ date. The question here really rests on: what is the intention in visualising this data?




Figure 200. (left) Nordic Museum. Portrait miniature, 1850-1897, showing actor Edvard Mauritz Swartz in costume as Hamlet. Identifier: NM.0094514. (CC BY-NC-ND).

Figure 201. (right) Photograph, 1983, in frame, 1890-1920. Identifier: NM.0323027A-B. (CC BY-NC-ND).

Often portraits in the Nordic Museum are dated to the century (or half/quarter century), so the most pronounced changes in costume can be seen by selecting a time span bridging different centuries. As Figure 202 shows, selecting a time span of 17 years around 1800 can place images (see Figures 203 and 204) side-by-side that might, in the popular imagination, characterise 18th and 19th Centuries respectively. Their date information mirrors this: 1700-1799 and 1800-1899. Beyond the choice of century/half-century date markers along the time axis, the visualisation design does not emphasise historical time as periods, but this way of sectioning up our view on the past still surfaces in the visualised data.

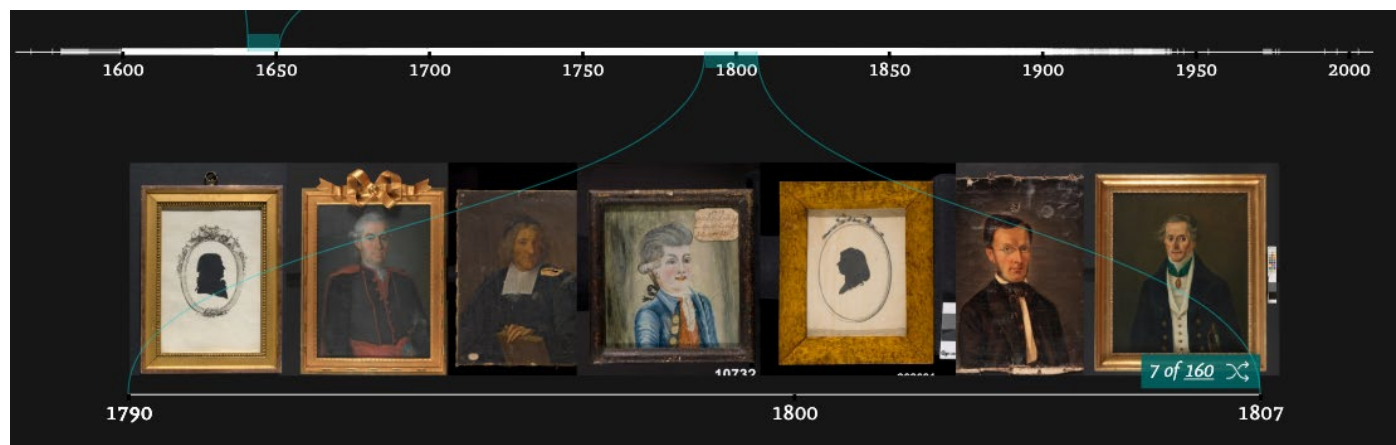


Figure 202. Detail from visualisation of Nordic Museum data.




Figure 203. (left) Production date 1700-1799. Identifier: NM.0010732. (CC BY-NC-ND).

Figure 204. (right) Production date 1800-1899. Identifier: NM.0082269. (CC BY-NC-ND).

The Nordic Museum curator observed the search query ‘Porträtt and tavla’ misses a number of, particularly popular, historical items in the collection which could be considered portraits: for example ‘kistebrev’, a type of hand-coloured broadsheet (Bringéus, 2008). Ought these to be included too? Would it make sense to mix in data from other museums? There are few 20th-century-onwards portraits in this dataset. Adding in data, for example from the Nationalmuseum (the national gallery of Sweden), could be a way to include more. But the basis by which artworks have been collected differs between institutions: items in the Nationalmuseum are primarily collected for their artistic value. Would the resulting visualisation end up *doing* something else? And the possibility of mixing in data from international institutions brings up the issue of wider context. The curator was wary about adding in data from beyond the Nordic countries without separating the data geographically. They suggested that, historically, artwork in the Nordic countries was influenced by artistic practice in, for example the UK or the USA, but a temporal delay is evident in styles travelling. If the data is not grouped geographically, the curator speculated “the timeline will look a bit strange”.

Though the curator was pleased there is an option to view the full dataset, they were satisfied with only mapping a sample of items on the timeline itself for casual browsing. They commented there are parallels with endeavouring not to overwhelm visitors in exhibition-making: “I’ve always been afraid of too much information, because I think that if it’s too much information [for a casual audience], [then it’s] no information”. They speculated the sampling approach is perhaps a better fit for this collection: “I think it’s better to have fewer. This is art. This is not only portraits”. Further, it is not necessary to see every single individual item to observe visual patterns in the images: “if you have the men here with their wigs, then you need only three men, not 30, if it’s what you want to say”.

Summary

This chapter introduced ‘Faces of Sweden’, a timeline visualisation for casual browsing of Nordic Museum portraits data building on prototyping from the previous chapter. The design allows users to interactively focus on time spans of their choice to see sampled items, and supports comparison. This project brought up particular issues with cultural image data: depiction date vs. production data, and the tension between general object photography vs. what visual characteristic is of interest. The kinds of patterns evident in this visualisation align with the museum’s collecting focus on cultural history (rather than, for example, fine art). Although without explicit narration, wider social trends such as the shifting social positions of those portrayed in collected portraits over time or that some items are part of a set, are not made visible. Discussing the prototype with a curator from the museum highlighted that care ought to be taken mixing in similar data from different data sources which may have different collecting contexts, as the patterns visible may be confounded.

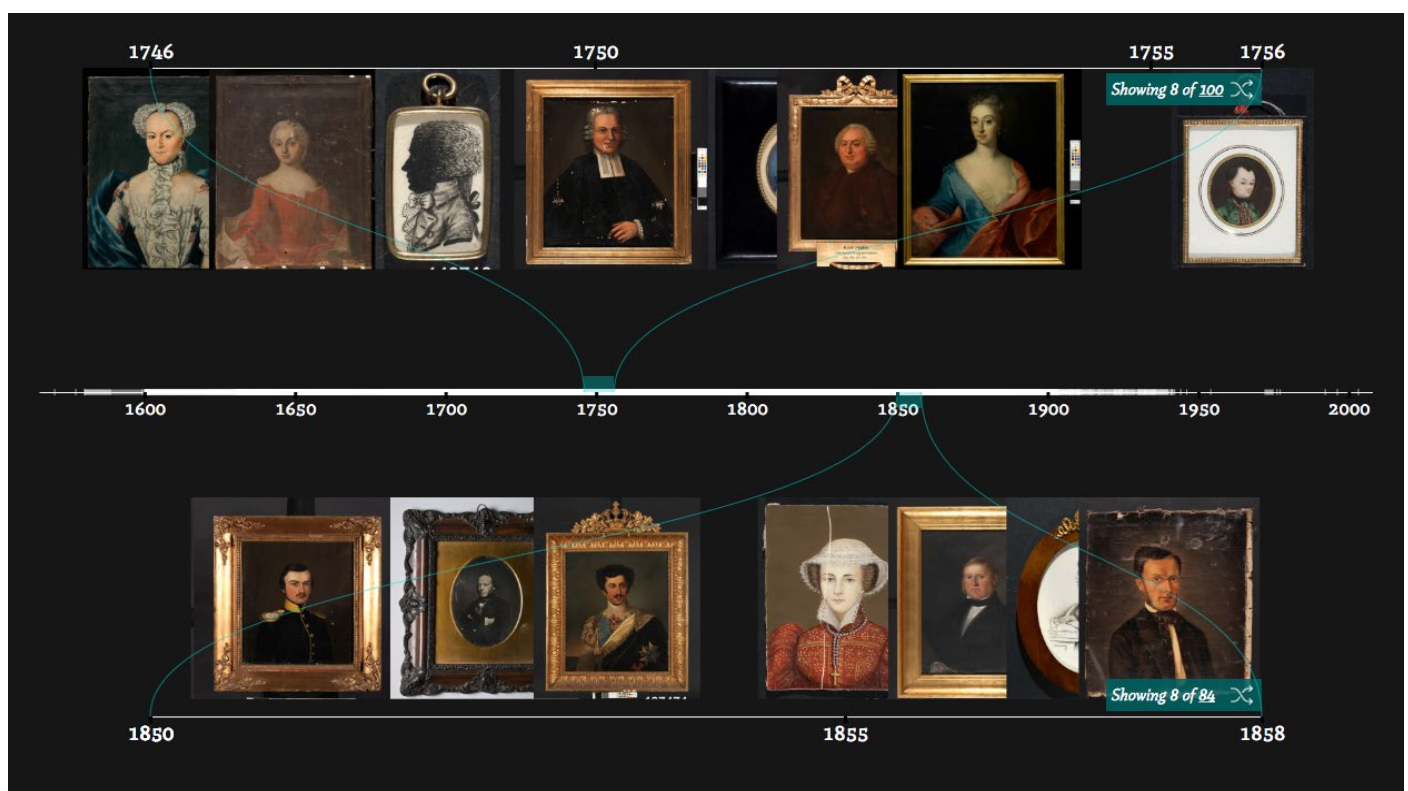


Figure 205. ‘Faces of Sweden’ visualisation.

Discussion

Is ‘just the data’ enough?

At the beginning of this thesis, I posed the question: **is ‘just the data’ enough?** Through cataloguing, cultural institutions create value and meaning around the items in their collections. By plotting these items by time, using the categories with which they have been described, historical narratives may be made visible. But to what extent is ‘just the data’ sufficiently well organised to visually tell the stories that people find valuable, versus the desire for authorial intervention?

What the data affords

Crucially, a timeline visualisation can only display what the data affords. Visualising a dataset a particular way is dependent on what connections have been made in the data and what is computationally possible. You can only extract colour data, for example, if there are high enough quality photographs of colourful items. (Though an institution can decide to create new cataloguing or further digitisation). Not only are datasets limited, but the provided APIs may further limit what can be done with them. I found this in the ‘Dive into Color’ project where there was no API method for what I required and I had to seek the data a different way.

As I set out in ‘Background’, different kinds of cultural institutions have traditionally had different ideas about what their role is in terms of interpreting and delivering collections to audiences, which are reflected in cataloguing conventions. Their data, therefore, looks different, which has consequences for how it may be visualised. Museum datasets typically include a greater number of facets for making connections in visualisation, contextualising items, than libraries or archives. The ‘Steptext’ prototype, for example, was developed with the Medical Officer of Health reports which have minimal cataloguing. The visualisation offers a transparent and interpretable mechanism for exploring commentary across the texts, but patterns may be subtle and a researcher must use their existing knowledge and refer to other resources to interpret them.

What the data affords is further underpinned by the cultural institution’s identity, as this determines what is in the collection and how it has been catalogued. Visualising portraits data from the Nordic Museum, Stockholm—a cultural history museum—by time unsurprisingly hints at trends in changing costume, rather than the progression of art movements. Similarly, the tags in the Cooper Hewitt collection are created to highlight threads relevant to design.

Just as a museum may loan objects for an exhibition to fill gaps in an argument due to the collection’s extent, a cultural dataset may not necessarily be enough on its own to tell certain stories that align with patterns and trends within it. The visual narratives that result may be fragmentary and incomplete.

Data reliability also plays a role in what the data affords. If the data is not completely reliable (for example, extracting colour data for the Cooper Hewitt collection in ‘Dive into Color’), key items cannot be guaranteed to appear in a timeline visualisation. In ‘Dive into Color’, Perkin’s mauveine scarf—an item the museum curators flagged as representing a significant moment in colour history—is absent from the timeline visualisation because of errors in the colour data. Exploring the Royal Photographic Society collection by visual similarity is hindered by the technique’s mishandling of items photographed with colour charts.

What cultural data affords, though, is changing because cultural data is changing. Many of these datasets are considered works in progress by the institutions holding them. New kinds of characteristics are being recorded, web resource norms are shaping expectations for how audiences experience collections online, accessibility plays a role, and new computational techniques open up new possibilities.

What the design affords

The choices made in how data is delivered in a visualisation will also determine what patterns may appear. In ‘Dive into Color’ because I chose to filter the data by colour harmonies only these patterns in colour data will be visible. The visualisation cannot show historical colour fashions that do not fit the harmonies structure, for example the late Victorian fashion for polychromy (multicoloured).

Do the timeline visualisations created in this PhD promote a particular kind of view on these collections? As Cooper Hewitt curators described in the ‘Tag-timeline’ project, this visualisation emphasises long threads through time—the historical continuum—encouraging connections between contemporary and historic items. These timeline visualisations could be said to be democratic in their display (ignoring that certain items may have been preferenced in cataloguing). Grouping items by shared characteristics in the data does not pay attention to provenance or rarity. Humble and prized items can appear side by side. On the other hand, this means items that might represent more significant moments in an historical narrative are not prioritised or even pointed out. Perkin’s mauveine scarf, for example, if it had been included in ‘Dive into Color’, would not have been any more prominent in the timeline than anything else.

Does the data speak for itself?

Are the patterns and narratives displayed obvious in these visualisations? A straightforward way to answer this question is clearly no, or I would not have required the help of the many collection experts I worked with in this PhD. The interpretation of, and engagement with, a cultural visualisation is highly dependent on the user and the use: what knowledge and expectations they bring to the interaction. One of the strengths of ‘Dive into Color’ is that it bridges popular appeal—it is playful and aesthetically pleasing without requiring knowledge about the collection or colour history, but with the further possibility of drawing attention to historical episodes, patterns and narratives in what is displayed.

The tags at Cooper Hewitt, visualised in ‘Tag-timeline’, are intended to provide a curatorial view of themes and connections across the collection: these links have been made in advance and can therefore be visualised. Some museum visitors I interviewed describing feeling this structuring made sense of the data for them:

“it wouldn’t have occurred to me to browse those aspects of it, and also suggests if I click on those there might be other things in the collection that were relevant to [the tag] ‘nautical’. Which I might not think would be true at this particular museum. It gives me new ways to think about the object that you’re looking at that I might not have thought of myself...And part of the reason you go to the museum is to have somebody knowledgeable helping you organize your thoughts about something. Like the expert is actually someone in control, reminding you of structure” (Museum visitor).

The terminology used in cataloguing plays a role here too. Concerned that some of the language is scholarly and not accessible enough, education staff at Cooper Hewitt recommended that, if adapting ‘Tag-timeline’ to a public-facing interface, it would be worth curating which tags to use in the interface and providing explanation for what some mean. It is also useful to consider the

context of data creation. Tagging data at Cooper Hewitt sometimes responds to an exhibition's thesis; this context is lost/confused as more data is later added to the same tag.

The kinds of stories people may want to tell by visualising a collection can be confounded when the visualisation shapes and patterns relate to an entangling of different factors: the collection extent, the institution's collecting history, the choices made in cataloguing items, historical patterns etc. Exploring the V&A collection data with timeline visualisation, grouping the visualised data by photographic technique begins to suggest historical periods when particular techniques were popular but, to a greater extent, the visualisation shape tells about the decisions made in prioritising what items to digitise first (shaped by various practical factors). The tag-timeline for 'Water' in the Cooper Hewitt collection (see Figure 206) traces items connected to the concept of water through time, but it also reflects historic changes in what the institution has collected, as its collecting mission shifted from the decorative arts to a modern understanding of design. The timeline displays garden plans from the 16th–18th Century, through to modern interventions for accessing clean water in developing countries.



Figure 206. Visualisation of Cooper Hewitt collection data: 'Water'.

There may be patterns in the data that are not separated out as data facets and, thus, without intervention are not obvious. In 'Faces of Sweden', plotting collection images by date prioritises visual patterns across the data, such as in costume. Other kinds of patterns, such as wider social trends (for example, that portraits from the 1830s/40s represent the rising bourgeoisie in Sweden at this time), are not obvious without existing knowledge or annotation. Similarly, some of these bourgeoisie portraits are one of a pair—for a married couple—though this is not indicated in the data. Without manual intervention, this is also concealed in the visualisation.

Shaping the narrative

Sometimes the curators I was working with internally at an institution wished to shape the narratives presented in timelines: either to change the items displayed in particular visualisations and/or to offer annotation (moving a timeline design from an 'implicit' to 'explicit' narrative). In the 'Tag-timeline' project, for example, conducting interviews with the museum curators revealed a desire to shape the timeline displays: to 'fill in' timelines by tagging further collection items or even use the visualisations to consider collection acquisitions (consider narratives not yet possible

to tell with the existing collection items alone): “I think we could produce a much more complete timeline if those of us tagging thought about it in reverse” (Cooper Hewitt curator, 2017c).

For ‘Dive into Color’, the curators were interested to bend the visualisation displays to the exhibition thesis through narration/annotation; historical colour theories proposed at various times could be plotted within the timeline to suggest how the design items that followed may have been influenced. Though, narration/annotation will likely be difficult for a visualisation running off live collection data which may change over time.

Customisation might play a role in negotiating this desire to tweak the visualised display. In addition, customisation could be helpful dealing with the kinds of inherent quirks found in working with cultural data. For example, in the ‘Faces of Sweden’ project where the portraits data included a record for a spare picture frame. From the starting point of a procedurally-generated timeline, what would it be like if users could drag, crop, resize, hide or add in visualised data points? To shape the visualisation to make the argument you want to make? Gibbs (2016) talks over this idea, raising questions about the implications this would have for a visualisation’s honesty:

How much do complex, data-driven visualizations need to [be] developed solely through computational means?

When creating representations of data largely done through software (and especially at large scales), must representations remain free of direct manipulation after an initial algorithmic rendering? Is it acceptable to alter a computed representation in order to highlight a particular feature? To what extent might that be considered subversive or misleading? To what extent is that simply better communication? Is the visualization more about the unadulterated output of the tool (even if unfortunately treated as a black box) or about communicating an interesting historical phenomenon? (Gibbs, 2016).

This point touches on whether visualisation is considered a final output or, as is the approach in this PhD, iterative visualisation *as research*. By exploring and experimenting with data through prototyping, the designer/developer/researcher can look at what they have just made, sometimes with their collaborators. Sometimes they may follow the lead prospected by what they have just made/discovered. Sometimes they may refine the coding/layout and other choices to get closer to the original intention.

Trust

Trust was an unexpected issue to emerge in this PhD. There is a long running debate around the use of visual embellishments in visualisation design, concerning what approach best conveys information to a user (Bateman et al., 2010:p.2574-5). The absence of visual rhetoric is, for some, connected with a visualisation’s perceived trustworthiness. For Tufte (1983:pp.106-121), for example, honest visualisation design requires being true to the data; its virtues are simplicity and clarity. However, minimal designs can also be misleading, and visual embellishments have benefits of their own (eg. Bateman et al., 2010; Hullman, Adar & Shah, 2011). As Boyd Davis (2017:p.4) puts it “there is no simple equation: decorated charts bad, undecorated good”.

Mayr et al. (2019) recently reviewed literature around trust in data visualisation, focussing on trust building mechanisms. However, this area has so far been little explored in a humanities visualisation context.¹ Mayr et al. define trust in data visualisation as “the user’s *implicit or explicit*

1. Though, perhaps, this is looking to change. The theme of the IEEE VIS conference’s 4th Workshop on Visualization for the Digital Humanities 2019 is “Building Trust: Process & Interpretation” (VIS4DH, 2019).

tendency to rely on a visualization and to build on the information displayed” (p.25). They argue that trust in this context depends on the trustworthiness of the data and of the visualisation design, and also the user’s trust perception. They raise a number of important points relevant here. Firstly, trust in data visualisation is highly context dependent: the data, visualisation, user, their goals and planned action (Marsh & Dibben, 2003). Trust is not relevant to all data visualisation interactions: it is important when there is some risk involved for the user (ibid.:p.471). And there are different mechanisms—depending on the context—for trust building: “users might elaborate the information deeply and gain a good understanding of [its uncertainties and quality]...But they might also use superficial cues [like usability and a positive user experience] as indicators for trust” (Mayr et al., 2019:p.25). The projects in this PhD provide some preliminary examples of extending trust in data visualisation to the cultural data context.

Trustworthiness of the data

The trustworthiness of a dataset depends on a number of factors; Kelton, Fleischmann & Wallace (2008) list accuracy, up-to-dateness, completeness, bias, validity and stability over time. In the case of the Medical Officer of Health reports, for example, the data could be described as highly accurate: the OCR text transcriptions are very high quality. On the other hand, the biases of these sources and their completeness for answering particular questions is also relevant. It is necessary for a scholar to understand the larger context of the text collection: who authored these sources and why, what may be missing, and are they alone sufficient for answering certain questions?

These projects also demonstrated that the accuracy of data produced through computational means can be complicated by mismatches between what may be accurate in a simplistic sense and what a human would likely say. For example extracting colour data from a photograph where the collection item had a coloured sticker on it (see Figure 207): a human would recognise the sticker’s colour is not relevant. Similarly, connecting items using feature vectors extracted from the Royal Photographic Society collection images could sometimes be viewed as naive. Blotchy paper is connected to photographs of grass (see Figure 208). Blotchy paper is not a photograph of grass, but could be said to be texturally similar.



Figure 207. Tiles photographed with stickers. Extracting colour data from these photographs includes the sticker’s colour. The colour data is ‘correct’ in a simple sense, though a human would treat it as irrelevant. Object IDs (left to right): 18726001; 18725273; 18725393. © Cooper Hewitt, Smithsonian Design Museum.

Different groups can have very different perceptions—even at odds—about the trustworthiness of the same data. In ‘Dive into Color’, the computationally-extracted colour data was known not to be 100% accurate; the museum curators described it as “very unreliable” whereas digital staff considered the technique had been successful. This perhaps reflects different values. Commenting on barriers to the development of Digital Art History as an emerging field, Zorich (2013) describes art historians as “perfectionists, a trait that serves them well in the discipline, but is at odds with the nature of digital research, where nimbleness—being able to work quickly to release research in preliminary and iterative stages—is vitally important” (p.15). This clash is recognisable in my interactions with many curators during this PhD (for example those uncomfortable creating tagging data which often seems highly subjective).



Figure 208. Prototype visualisation of RPS collection data. Using computationally-extracted visual similarity data connects an image of blotchy paper (top)—a photograph’s back—with photographs of grass (bottom middle); this could be considered naive, though the images are texturally similar.

Trustworthiness of the visualisation

In ‘Steptext’, historian interviewees underlined the importance of trust in the visualisation, connecting the simple aesthetics and mechanism of the visualisation with trustworthiness. The virtue of simplicity in digital humanities tools has been noted before. Gibbs and Owens (2012) argue that insufficient usability is an issue with many digital humanities tools; a simple and clear interface promotes ease-of-use and transparency.

But a simple mechanism is not always possible. Designing ‘Dive into Color’, I was working with colour data I knew had errors to visualise subjective qualities. Rather than being able to precisely explain how the visualisation worked (which I found a challenge in evaluations), a number of interviewees remarked that the persuasiveness of the visualisation lay in the results looking ‘right’. This is perhaps because the filtering characteristic—colour—is an immediate visual quality, so it is obvious if there are results present that should not be there (though not if data is missing).

The opaque workings of data, algorithms and software was an issue I encountered throughout these projects and may have implications for a visualisation’s trustworthiness. Historians I interviewed in the ‘Steptext’ project were concerned about the interpretability of relevancy ranking. Using a pre-trained machine learning model and applying dimensionality reduction techniques with the RPS photographs requires taking the validity of the model and algorithms on trust. There can be a lack of clarity in software documentation. The Cooper Hewitt API documentation did not make clear how results were preferenced if a limit was placed on the number returned. In these cases, I was weighing up the consequences of such opacity with my expectations for the visualisation’s use and audience/s.

Where collection data is a work in progress, it may be that some records are higher quality than others, and taking stock of a record as a whole helps indicate the information quality. Visualisation, in contrast, is a mapping of some parts of a cultural item’s record and may not encourage close

engagement with the full records. Raising concerns about the exploratory RPS visualisations, a curator said:

I wouldn't necessarily trust information displayed like this [in a timeline visualisation] in the same way as I would trust looking at our catalogue...in our catalogue I can see the data behind it...I can also see who put it in and when. And...what else is in [the] record. Whether it was a quick record that we did just to have a record when we digitised the object or whether it was something that somebody really spent time and expertise doing (V&A curator, 2019).

In this case, however, the curator made this remark prior to being shown that the full record for each visualised data point could be accessed by hyperlinks through to the collection website. Mayr et al. (2019) point out the process of trust building for a user changes and evolves with their interactions with a data visualisation over time (p.26). And since trust translates most explicitly into a user's actions (whether they do or not build on the visualised data), "an assessment of trust should at least measure a user's behavioral intentions (Would you rely on the facts in this...[data visualisation]? Would you use the information for a decision?), if not his actual behavior" (p.28). More focussed investigations of trust in cultural visualisation would likely benefit, therefore, from studying longer interactions and exploring users' behaviour and intentions.

Quantitative, qualitative visualisation

My project work in this PhD has underlined the need for caution in applying quantitative timeline visualisation to cultural datasets. Particularly, the issue of what exactly constitutes a single item has the power to greatly distort the shapes produced. Providing a route back to the original data source is productive for avoiding misinterpretations.

This PhD documents a variety of cultural visualisation prototypes exploring the qualitative visualisation space. It is sometimes assumed that the design space for 'distant reading' is a large one, but that for 'close reading' is not. This research project highlights that exactly what kind of close engagement is desirable depends on the user, the use, and the dataset in question. Evaluating 'Stext' with historians demonstrated, for this audience, what the texts say was much more important than what the page images look like. But it is also relevant, in this case, that the transcribed text data was very high quality and, therefore, checking for OCR errors against the page images is less essential.

Completeness

Having access to the full collection data, 'seeing everything' and knowing what is being searched over is very important in a research context. For popular audiences using timeline visualisation for more casual exploration, however, completeness was seen as less essential. Evaluations with visitors at Cooper Hewitt suggested they did not want to be overwhelmed and that sufficient and diverse content representing the museum's holdings was more important. As a museum visitor said: after all, "you only see part of the collection when you...[physically come to the museum] too". There is a further question of whether visualising everything in a digitised collection equivalently is *necessarily* the most representative of an institution's holdings anyway: should each digital record be treated equally when they may represent completely different physical scales or importance of items?

Are some collections more 'comprehensive' than others? Perhaps in relation to how tightly-defined the collecting agenda or collection boundaries are? It might be said, for example, that the Medical

Officer of Health reports collection is more ‘comprehensive’ (every surviving copy of this report type) than, say, the Cooper Hewitt collection—a collection of historic and contemporary design—the boundaries of which are unavoidably more nebulous. And, in that sense, is ‘completeness’ in visualisation more important for some collections than others?

How many interface designs do we need?

This PhD has introduced a number of new cultural visualisation designs. But how many new interface designs do we need? Could there be one super-visualisation that would suffice for all instances, times, users etc.? After all, an individual visualisation design can support multiple modes—analysis, exploration and presentation. The heterogeneous content and character of cultural datasets and the varying ways users wish to engage with them suggests barriers to such an ambition. What design is desirable depends on the user, the use, and the dataset in question. (Though Haskiya (2019) points out developing custom designs is expensive, which can be a barrier to adoption for cultural institutions). Also, just as cataloguers are finding what aspects of the collection audiences are interested in is changing, likely what is desirable in cultural visualisation design will also change with time.

Reflections

Designing cultural timelines

The timeline prototypes presented in this PhD could all be described as conventional timelines (linear, horizontal, time running left to right) as opposed to, for example, what Silva & Catarci (2000) call 'innovative' timelines which make use of distortion, visual metaphors, 3D etc. I do not see this as a problem. This thesis demonstrates there is room for innovation in conventional timeline designs for cultural data. Further there are benefits for visualisation users in that a familiar form that can be understood more quickly and easily (Carpendale, 2008:p.20). Boyd Davis (2012) has previously explored the implications and perception of time axis orientation and direction in timeline design.

In 'Background', I described different modes (analysis, exploration and presentation) of cultural visualisation (though a design can straddle more than one). The prototypes developed in this PhD, while tending around the exploration mode, also overlap into other territories: 'Steptext' into analysis, and the other prototypes more towards presentation. I also quoted Schnapp et al. (2009) who call for digital toolkits that serve the humanities' core methodological strengths: "attention to complexity, medium specificity, historical context, analytical depth, critique and interpretation" (p.2). The visualisations portfolio in this PhD respond to this call, paying attention to the complexity, character, and context of cultural datasets. In addition to "medium specificity" these projects highlight the importance of domain specificity, by which I mean it is not just important that the digitised items in the Cooper Hewitt collection represent objects, it matters that they are a design collection. 'Dive into Color' explores the colours and colour harmonies present across the collection—especially relevant to design objects. Cataloging within an institution pays attention to what for the institution is important, defined by the domain it covers. Choices in how digitisation is carried out are affected too (for example, the V&A Museum wanting, as an art and design museum, high-quality aesthetically beautiful photography of items (V&A curator, 2018)).

This PhD takes a 'Research through Design' approach. As expanded on in 'Methods', prototyping designs can progress research by helping to define and redefine the problem being investigated. This is something I experienced in this PhD work where, through prototyping with cultural datasets and discussing those prototypes with collection experts, I became increasingly aware that cultural datasets are a shifting landscape and there is a clash of values from involved parties in negotiating these shifts. This became a growing focus in my investigation, informing my understanding of my research questions.

Intuitive, aesthetic decisions

Throughout the practical work in this PhD, I sometimes made design decisions on intuitive, aesthetic grounds (making things look/feel 'right'). In the 'Tag-timeline' project I designed the prototype to even out data distribution across the timeline. This design decision means it is more difficult to compare the data's temporal distribution across tags. Instead, the displays are less sparse and the temporal gaps between images are smaller, perhaps encouraging users to view the displayed items as a temporal continuum. As a museum visitor said: "there's something quite nice about the ... not sparsity of stuff but like graphically it's quite nice because you got everything there and it's composed" (Visitor).

In the 'Dive into Color' project, I often made decisions on these grounds and found the Cooper Hewitt curators in agreement over my adjustments. I and the curators decided to exclude the more complex harmony options (quadratic, tetradic, split complementary) because the results *felt* less convincing and the results were less visually cohesive. Since colour harmony bounds are not well defined numerically, I manually adjusted my query searches, by personally judging the results quality. In the end, the museum curators considered my manual adjustments "very effective" and resulting timelines "very satisfactory". And as a museum curator said: "the visual, the satisfying...seeing the most compelling works has a value as well" (Cooper Hewitt curator, 2018). As the direction of the project shifted to producing an interface for public exhibition, designing for a satisfying experience was further judged as appropriate.

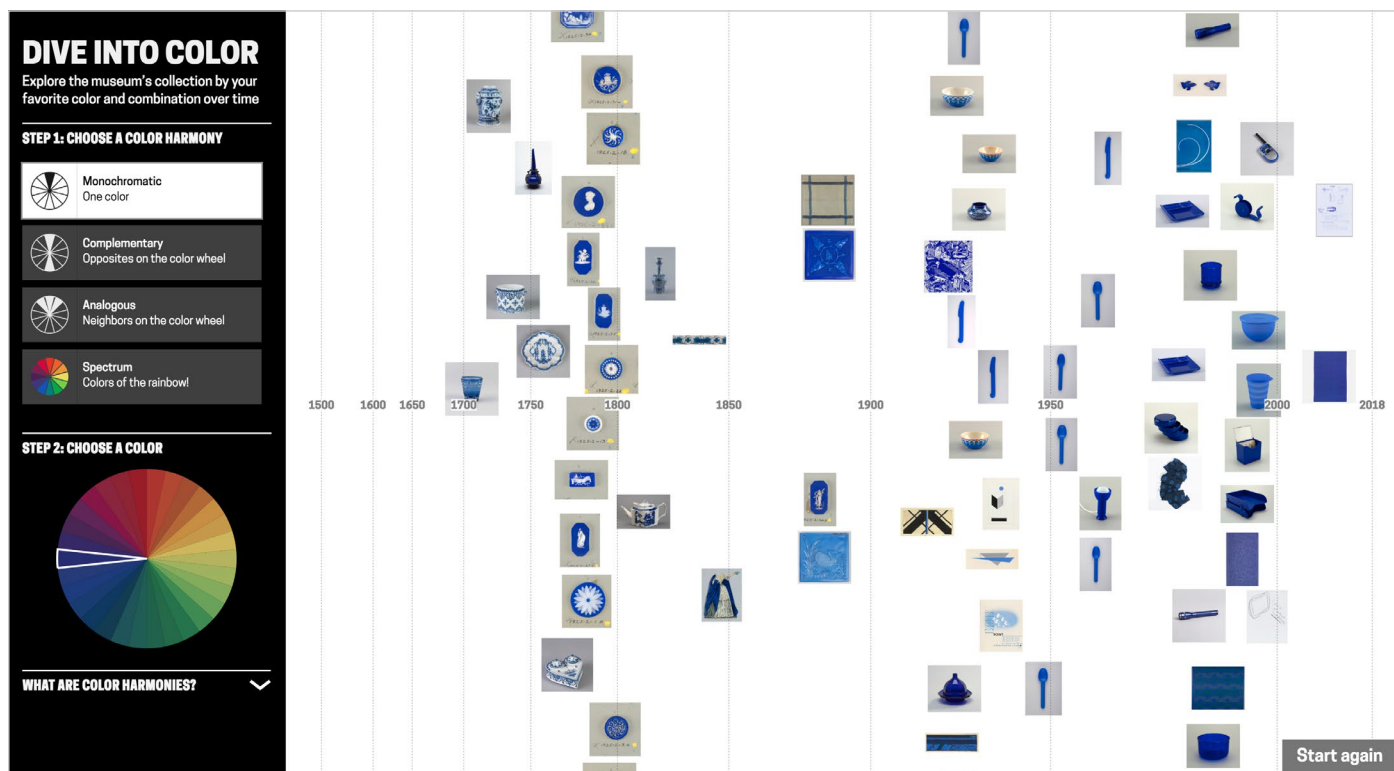


Figure 209. Visualisation of Cooper Hewitt collection data: Monochromatic, blue.

In some ways, this might be seen as in opposition to the need for honesty/transparency, as stressed by the interviews conducted with historians for the 'Stextext' project. But aesthetics are important in cultural data visualisation especially, as I discuss in 'Methods: Evaluating cultural visualisations', in bridging between values in computing and the humanities.

Pragmatic decisions

At times I also had to weigh up decisions for pragmatic reasons. Sometimes there was a lack of transparency in the data system I was using. When I was building 'Dive into Color', no one was able to explain why the colour dataset size corresponded to only about 60% of the total number of collection items. I decided to embark on prototyping with the data available since my design intention was not a quantitative research tool. Also, if the tool could be demonstrated as valuable on the data available, it would be relatively straightforward to include the absent items later if the missing data could be accessed.

In some of these projects, I manually edited my copy of the data. For example, in 'Dive into Color' I manually edited some of the colour data to remove obvious errors, and excluded textile sample book pages which had multiple designs in one image. In 'Faces of Sweden', I manually excluded items where the photograph did not show a portrait image. In these cases, the exclusions felt appropriate for the visualisation aims and none required expertise.

Evaluations

It can be difficult to bridge between digital humanities and visualisation/human-computer interaction research. There are different expectations for what counts as a valuable research contribution, what constitutes a visualisation and what kind of evaluation is appropriate. As Jänicke (2016) describes, it can be difficult for researchers working on cultural visualisations, to make a contribution to both computer science and humanities domains:

When an approach is too complex – which counts as a strong argument for a publication in a visualization realm – it might get ... [unusable] for humanities scholars due to problems of comprehension. On the other hand, if a clean, easily comprehensible visualization is valuable for a humanities scholar, the missing novelty most likely impedes a computer science publication (Jänicke, 2016).

I experienced this first-hand. Submitting the 'Stext' technique to a computer-science-domain visualisation conference, one reviewer commented: "the technique is simple and elegant, but that also raises questions as to the significance of the contribution and whether it is even a 'visualization' at all". Further, presenting the 'Stext' prototype to a cohort of other PhD students at a Human-Computer Interaction conference, it was recommended I evaluate the tool by running a lab study with historians, counting how many new insights they achieve with the tool. As I discuss in 'Methods: Evaluating cultural visualisations', there are problems with such an approach.

This PhD has used qualitative evaluations with prospective users to reveal their impressions of designs, interpretations of visualisations produced in them and point to their values and requirements. To better understand the potential role of 'Stext' in researcher's activities, though, case study evaluations would be required carried out over a longer period with a more robustly developed prototype. (This is the approach used by McCurdy et al. (2015), as recommended by Lam et al. (2011:p.6)). While this would be an interesting direction in future work, it has not been the investigation focus of this PhD.

Interviews

There were sometimes practical constraints around recruiting interviewees. This was one of the significant benefits of being embedded within Cooper Hewitt. It not only made it easier to arrange interviews with many of the museum curators, but also to interview museum *visitors*. Capturing impressions from visitors this way meant: I was engaging with those already interested in the museum; interviewing them in person allows a rich exchange and makes it easier to quickly introduce what the project is all about; and also did not require the prototype to be as developed, polished or robust as it would need to be to publish it online and seek feedback through the web. Similarly, I was fortunate in the 'Dive into Color' project that, although I had finished my fellowship within the museum, this PhD was conducted within an art and design university. This opened up recruiting history of design students for interviews and help contacting historical colour experts. In contrast, there were obstacles in 'Faces of Sweden', as my project partner was the Swedish National Heritage Board rather than the Nordic Museum itself. Interviewing a museum curator about the visualised collection in-person was only possible once, midway through the project on a scheduled Stockholm visit.

In interviews, one of the questions I received particularly insightful responses to was: what it meant, to the interviewee, to 'make sense' of a collection like the one visualised and whether they felt my prototype helped 'make sense'. I found the ambiguity of the question provoked really interesting answers as different individuals interpreted the question in different ways: making sense of collections in general, making sense for others, making sense of digitised collections. Answers related

back to their interests in, outlooks on, and interactions with collections. I found this diversity of perspectives enriched my understanding of the various roles these resources can play, and interactions users wish to have with them.

I also found that asking interviewees whether they would use the visualisation in their own work and, if so, how, prompted insightful responses. Rather than just focussing on what the visualisation in front of them is doing, this extended discussion to unanticipated use cases and questions, speculating what could be achieved perhaps with a different dataset or adjusted functionality. It also helped me develop a better understanding of where the visualisation might fit in wider activities. Although the designs developed in this PhD were motivated by individual datasets, interviewees often suggested other collections that would also be well suited to the prototype design. Collections may be diverse and heterogeneous, but there are still commonalities.

When I was evaluating ‘Tag-timeline’ with curators at Cooper Hewitt, I wondered if it would also be a good opportunity to get more feedback on the Steptext design from a slightly different audience. While the texts visualised are not themselves relevant, I was interested to capture impressions from the curators, who also work with historical texts, to the technique generally. Responses were mixed. While I had a very enthusiastic reaction from some (“That’s incredible... that’s so much data...I want this for my research right now...My hours of going through the documents...this [would] just cut hours”), others were left confused at the sudden change in direction of the session and did not know what to make of it. I decided it was not productive to do this in future sessions.

Limitations and future work

The PhD’s time constraint meant evaluations for the ‘Time-tiers’ design and, related, ‘Faces of Sweden’ are more preliminary and only with curators internally at those institutions. In future work it would be interesting to explore this visualisation technique with casual users, as this was the intended audience. It was also flagged ‘Faces of Sweden’ may be fitting in an education context. This could also be explored.

Is this work future-proof? There is growing interest in automated tagging of collections, for example. I would argue the prototypes developed in and outlook of this PhD investigation are future focussed. I have engaged with new kinds of and techniques for augmenting cultural data, and with institutions interested in pushing what digitised collections are and can do.

While the prototypes in this PhD were developed with real cultural datasets, they do not represent the universality of digitised collections; there is a bias towards visual art/design collections, and towards museums over other institution types. Further work with more diverse datasets and institutions would widen perspectives. Winters (2017b), for example, points to challenges for national archives and libraries archiving the web, with its own set of issues around completeness.

Although combining different cultural datasets using linked data structures (Hyvönen, 2012) has recently gained a lot of attention in the cultural heritage domain, it has not been the focus here; this PhD, in fact, has stressed the importance of data context in effectively designing cultural timeline visualisations. Further, these projects do not visualise data structures beyond flat ones; hierarchical data, for example, like the Babbage Papers archival records mentioned in the Background chapter, has not been considered. Both would be interesting topics to extend this work to.

These projects add to understanding of the many problems of dating in archives and museums. For example, one of the date information complications that came up in this PhD, for image data,

is that of date depicted vs. date made. There have been recent moves to make records deal with dating issues more subtly, comprehensively and transparently, for example the Extended Date/Time Format (EDTF) (Library of Congress, 2019). As support in data structuring for date depicted vs. made increases, the issue shows potential for further investigation with timeline visualisation.

Research contribution

The outputs of this PhD are primarily of interest to cultural heritage institutions and the field of digital humanities. They are also relevant to the fields of design, information visualisation and human-computer interaction.

Practical research contribution

A survey of the state of the art in timeline visualisation of cultural data is presented, exploring the various ways narratives and patterns in a collection can be presented—implicitly and explicitly—from handcrafted to automated examples. This PhD presents a portfolio of novel timeline designs exploring the qualitative visualisation space, customised to cultural datasets. By designing and developing prototypes with real, currently-available cultural datasets, and by seeking feedback with a range of users, I have been able to develop prototype designs responding to the realities of digitised collections, and supporting the types of interactions users wish to make with them. While these visualisations were developed around individual datasets, they could also be applied to other suitable cultural datasets.

Theoretical research contribution

This thesis argues that cultural data is changing, impacting what cultural visualisations can do and for who. I add clarity in relation to the question of what cultural visualisation is for, synthesising existing literature and including my own experiences. One of the problems raised at the beginning of this thesis is that we need to better evidence the value of digitised collections; digitisation projects are expensive to undertake and maintain, and for some there is a perception we need to get ‘more’ out of them. The visualisations developed in this PhD exemplify visualisation as added value to cultural data: for users of various kinds, and for the data-holding institutions.

Through the portfolio of visualisation designs developed in this research, this thesis has explored the relationship between procedural visualisation and narrative in cultural data, identifying and discussing factors that determine what is possible and its success. There is a tendency for quantitative visualisation to have a much bigger literature and more scholarly debate than the visualisation of qualitative or nominal data. I have helped redress this balance and shown how important, difficult and subtle the issues are in the latter. This thesis provides some preliminary examples extending the study of trust in data visualisation to the humanities context. I also make a research methods contribution by adding to the exemplars of visualisation for/as research, and practice-based research more generally.

References

Aigner, W., Miksch, S., Schumann, H. & Tominski, C. (2011) *Visualization of Time-Oriented Data*. Human-Computer Interaction Series. Springer.

Albano, C. (2014) Exhibition. In: *Encyclopedia of Aesthetics*. Oxford, Oxford University Press, pp.533–538.

Alexander, J., Barton, J. & Goeser, C. (2013) Transforming the Art Museum Experience: Gallery One. *MW2013: Museums and the Web 2013*. Available from: <https://mw2013.museumsandtheweb.com/paper/transforming-the-art-museum-experience-gallery-one-2/> [Accessed 21st March 2019].

Alexander, M., Dallachy, F., Piao, S., Baron, A. & Rayson, P. (2015) Metaphor, popular science, and semantic tagging: Distant reading with the historical thesaurus of english. *Digital Scholarship in the Humanities*, 30(Supplementary 1), pp.i16-i27.

Anderson, W. (2018) Foreword. In: *Spitzmaus Mummy in a Coffin and other Treasures*, p.3. Available from: https://press.khm.at/fileadmin/content/KHM/Presse/2018/Wes_Anderson/KHM_ArtistsChoice_Booklet_en_FINAL_1028_14.pdf [Accessed 17th April 2019].

Anne Frank House. (2010) *Anne Frank Timeline*. Available from: <https://web.archive.org/web/20101113192829/http://www.annefrank.org/timeline#!/web/20101113192829/http://www.annefrank.org/timeline> [Accessed 31st March 2019].

Anthony, L. (2019) AntConc. Available from: <http://www.laurenceanthony.net/software/antconc/> [Accessed 31st March 2019].

Archer, B. (1995) The Nature of Research. *Co-design, interdisciplinary journal of design*, 2(11), pp.6-13.

Arnold, T. & Tilton, L. (2019) Distant viewing: analyzing large visual corpora. *Digital Scholarship in the Humanities*, [fzq013].

Arrigoni, G. & Schofield, T. (2015) Understanding artistic prototypes between activism and research. In: *Computation, Communication, Aesthetics and X (XCoAx 2015)*, Glasgow, Scotland, pp.24-39. Available from: <http://2015.xcoax.org/pdf/xcoax2015-Arrigoni.pdf> [Accessed 10th March 2019].

Ashmolean: Eastern Art Online. (2013) *Pilgrimage*. Available from: http://jameelcentre.ashmolean.org/collection/6/10239/0/all/sort_by/seqn [Accessed 31st March 2019].

Baca, M. (ed.) (2016) Introduction to Metadata. 3rd ed. Los Angeles, Getty Research Institute. Available from: <http://www.getty.edu/publications/intrometadata/> [Accessed 31st March 2019].

Bailey, J. & Pregill, L. (2014) Speak to the eyes: the history and practice of information visualization. *Art Documentation: Journal of the Art Libraries Society of North America*, 33(2), pp.168-191.

Bal, M. (1996) *Double Exposures: The Subject of Cultural Analysis*. New York and London, Routledge.

Barnett, J. R., Miller, S. & Pearce, E. (2006) Colour and art: A brief history of pigments. *Optics & Laser Technology*, 38(4-6), pp.445-453.

Barrett-Small, R. (2014) *Colour Lens*. Available from: <http://colourlens.org/> [Accessed 31st March 2019].

Bateman, S., Mandryk, R. L., Gutwin, C., Genest, A., McDine, D. & Brooks, C. (2010) Useful junk?: the effects of visual embellishment on comprehension and memorability of charts. In: *Proceedings of the 2010 SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, pp. 2573-2582.

Bennett, T. (1988) The British experience with heroin regulation. *Law and contemporary problems*, 51(1), pp.299-314.

Bernhard, M. (2016) Gugelmann Galaxy: an unexpected journey through a collection of Schweizer Kleinmeister. *DAH-Journal: International Journal for Digital Art History*, 2, pp.94-114. Available from: <https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/23250/27208> [Accessed 21st March 2019].

Bishop, C. (2017) Against Digital Art History. *Humanities Futures*, Franklin Humanities Institute. Available from: <https://humanitiesfutures.org/papers/digital-art-history/> [Accessed 4th March 2019].

Boyd Davis, S. (2017) Early visualisations of historical time: "To see at one glance all the centuries that have passed". In: Black, A., Luna, P., Lund, O. & Walker, S. *Information design: research and practice*. Abingdon, Routledge, pp.3-22.

Boyd Davis, S. (2012) History on the Line: time as dimension. *Design Issues*, 28(4), pp.4-17.

Boyd Davis, S., Bevan, E. & Kudikov, A. (2010) Just In Time: defining historical chronographics. In: *Proceedings of the Conference on Electronic Visualisation and the Arts, EVA 2010, 5th-7th July 2010, London, UK*. British Computer Society, pp.355-362.

Boyd Davis, S. & Kräutli, F. (2015) The Idea and Image of Historical Time: Interactions between Design and Digital Humanities. *Visible Language*, 49(3). Available from: <http://visiblelanguagejournal.com/issue/172/article/1202> [Accessed 7th March 2019].

Boyd Davis, S. & Kräutli, F. (2014) Scholarly chronographics: can a timeline be useful in historiography? In: *Proceedings of 10th European Social Science History Conference, 23rd-26th April, Universität Wien, Vienna*. Available from: http://researchonline.rca.ac.uk/1623/1/BoydDavis_ScholarlyChronographicsESSHC_2014DRAFT.pdf [Accessed: 12th April 2019].

Boyd Davis, S., Vane, O. & Kräutli, F. (2016) Using Data Visualisation to tell Stories about Collections. In: *Proceedings of the Conference on Electronic Visualisation and the Arts, EVA 2016, July 2016, London, UK*. British Computer Society. pp.221-228.

Brehmer, M., Lee, B., Bach, B., Riche, N.H. & Munzner, T. (2017a) Timelines revisited: A design space and considerations for expressive storytelling. *IEEE transactions on visualization and computer graphics*, 23(9), pp.2151-2164.

Brehmer, M., Lee, B., Riche, N.H., Edge, D., White, C., Lytvynets, K., & Tittsworth, D. (2017b) Microsoft Timeline Storyteller. Available from: <https://timelinestoryteller.com> [Accessed 19th May 2019].

Brenner, S. & Adang, L. (2016) Winning (and losing) hearts and minds of museum staff: Administrative interfaces at Cooper Hewitt. *MW2016: Museums and the Web 2016*. Available from: <https://mw2016.museumsandtheweb.com/paper/winning-and-losing-hearts-and-minds-of-museum-staff-administrative-interfaces-at-cooper-hewitt/> [Accessed 19th March 2019].

Brett, M.R. (2012) Topic modeling: a basic introduction. *Journal of Digital Humanities*, 2(1). Available from: <http://journalofdigitalhumanities.org/2-1/topic-modeling-a-basic-introduction-by-megan-r-brett/> [Accessed 31st March 2019].

Bringéus, N. A. (2008) Massmedium i miniformat: En översikt över den svenska kistebrevsproduktionen. *Rig*, 91(3), pp.129-139.

British Museum & Google Cultural Institute. (2015) *The Museum of the World*. Available from: <https://britishmuseum.withgoogle.com/> [Accessed 31st March 2019].

Burdick, A. (2009) Design Without Designers. *Keynote for a conference on the future of art and design education in the 21st century. University of Brighton, England, 7th February 2009*. Available from: http://anneburdick.com/Design-wo-Designers/Burdick_Design_wo_Designers_sm.pdf [Accessed 4th March 2019].

Burke, P. (ed.) (2001) *New perspectives on historical writing*. 2nd Edition. Polity Press.

Cairns, S. (2013) Mutualizing museum knowledge: Folksonomies and the changing shape of expertise. *Curator: The Museum Journal*, 56(1), pp.107-119.

Cameron, F. & Robinson, H. (2007) Digital knowledgescapes: Cultural, theoretical, practical, and usage issues facing museum collection databases in a digital epoch. In: Cameron, F. & Kenderdine, S. (eds.). *Theorizing digital cultural heritage: A critical discourse*. MIT Press, Cambridge, MA, pp.165-191.

Card, S., Mackinlay, J. & Shneiderman, B. (1999) *Readings in Information Visualization: Using Vision to Think*. San Francisco, Morgan Kaufmann Publishers Inc.

Carpendale, S. (2008) Evaluating information visualizations. In: Kerren, A., Stasko, J., Fekete, J. D. & North, C. (eds.). (2008) *Information Visualization: Human-Centered Issues and Perspectives*, 4950. Springer, pp.19-45.

Caviglia, G., Coleman, N. & Ciuccarelli, P. (2012) Communication Design and Digital Humanities: Visualizations and Interfaces for Humanities Research. In: *4th International Forum of Design as a Process, within the programme of activities of the 4th Bienal Brasileira de Design*, Universidade do Estado de Minas Gerais, Brazil, 19th-22nd September 2012.

Chen, K. L., Dörk, M. & Dade-Robertson, M. (2014) Exploring the Promises and Potentials of Visual Archive Interfaces. In: *Proceedings of the 2014 iConference*, pp.735-741.

Cherry, A. & Mukunda, K. (2015) A case study in indigenous classification: Revisiting and reviving the Brian Deer Scheme. *Cataloging & Classification Quarterly*, 53, pp.548-567.

Chosen. (2011) Available from: <https://harvesthq.github.io/chosen/> [Accessed 22nd March 2019].

Chuang, J., Ramage, D., Manning, C. & Heer, J. (2012) Interpretation and trust: Designing model-driven visualizations for text analysis. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, May 2012*. ACM, pp.443-452.

City of Amsterdam. (2018) *Below the Surface*. Available from: <https://belowthesurface.amsterdam/en/vondsten> [Accessed 16st May 2019].

Cleveland Museum of Art. (2014) The Collection Wall, *Museums and the Web 2014, MW2014*. Available from: <http://mw2014.museumsandtheweb.com/bow/collection-wall/> [Accessed 31st March 2019].

Cockburn, A., Karlson, A.K. & Bederson, B.B. (2008) A Review of Overview+Detail, Zooming, and Focus+Context Interfaces. *ACM Computing Surveys*, 41(1).

Coles, K. (2017) Think Like a Machine (or not). In: *2nd Workshop on Visualization for the Digital Humanities (VIS4DH), held as part of VisWeek'17, Phoenix, United States, 2nd October 2017*. Available from: <http://vis4dh.dbvis.de/2017.html> [Accessed 2nd April 2019]

Coles, K. (2016) Show Ambiguity. In: *1st Workshop on Visualization for the Digital Humanities (VIS4DH), held as part of VisWeek'16, Baltimore, Maryland, USA, 24th October 2016*. Available from: <http://vis4dh.dbvis.de/papers/2016/Show%20Ambiguity%20Collaboration,%20Anxiety,%20and%20the%20Pleasures%20of%20Unknowing.pdf> [Accessed 2nd April 2019].

Collections Trust. (2019) *Primus*. Available from: <https://collectionstrust.org.uk/software/primus/> [Accessed 8th May 2019].

Collins, C., Viégas, F. B., & Wattenberg, M. (2009) Parallel tag clouds to explore and analyze faceted text corpora. In: *Symposium on Visual Analytics Science and Technology, VAST 2009*. IEEE, pp. 91-98.

Collection manager. (2018) *Recorded interview with Olivia Vane*, 1st May 2018.

Costante, E., Den Hartog, J. & Petkovic, M. (2011) On-line trust perception: What really matters. In: *1st Workshop on SocioTechnical Aspects in Security and Trust (STAST)*. IEEE, pp.52-59.

Cook, T. (2009) The archive(s) is a foreign country: Historians, archivists, and the changing archival landscape. *The Canadian Historical Review*, 90(3), pp.497-534.

Cooper Hewitt. (2019a) *About Cooper Hewitt*. Available from: <https://www.cooperhewitt.org/about-cooper-hewitt/> [Accessed 18th March 2019].

Cooper Hewitt. (2019b) *The Collection*. Available from: <https://collection.cooperhewitt.org/> [Accessed 18th March 2019].

Cooper Hewitt. (2019c) *Application Programming Interface*. Available from: <https://collection.cooperhewitt.org/api/> [Accessed 18th March 2019].

Cooper Hewitt. (2019d) *Bless you chair, 1999*. Available from: <https://collection.cooperhewitt.org/objects/18714235/> [Accessed 26th March 2019].

Cooper Hewitt. (2019e) *Color Wheel, 1936–37; Hilaire Hiler*; Oil paint on composition board. Available from: <https://collection.cooperhewitt.org/objects/18433055/> [Accessed 1st April 2019].

Cooper Hewitt. (2019f) *Saturated: The Allure and Science of Color*. Available from: <https://www.cooperhewitt.org/channel/saturated/> [Accessed 2nd April 2019].

Cooper Hewitt. (2019g) *Lotus-shaped Cup (Egypt)*. Available from: <https://collection.cooperhewitt.org/objects/18433657/> [Accessed 5th April 2019].

Cooper Hewitt. (2019h) *Scarf Sample (France)*. Available from: <https://collection.cooperhewitt.org/objects/18615843/> [Accessed 5th April 2019].

Cooper Hewitt. (2018) *Colors!*. Available from: <https://web.archive.org/web/20180407235218/https://collection.cooperhewitt.org/objects/colors/> [Accessed 1st April 2019].

Cooper Hewitt. (2017) Collection Data for Cooper Hewitt: Smithsonian Design Museum, *Github*. Available from: <https://github.com/cooperhewitt/collection> [Accessed 1st April 2019].

Cooper Hewitt. (2016) *Hewitt Sisters Collect*. Available from: <https://www.cooperhewitt.org/events/current-exhibitions/hewitt-sisters-collect/> [Accessed 25th March 2019].

Cooper Hewitt. (2014) A Toast to Life, Death, and Rebirth. Available from: <https://www.cooperhewitt.org/2014/04/26/a-toast-to-life-death-and-rebirth/> [Accessed 5th April 2019].

Cooper Hewitt curator. (2018) *Recorded interview over video call with Olivia Vane*, 6th March 2018.

Cooper Hewitt curator. (2017a) *Recorded interview with Olivia Vane*, 11th July 2017.

Cooper Hewitt curator. (2017b) *Recorded interview with Olivia Vane*, 12th July 2017.

Cooper Hewitt curator. (2017c) *Recorded interview with Olivia Vane*, 13th July 2017.

Cooper Hewitt Labs. (2017) *Exploring the Cooper Hewitt collection with timelines and tags: guest post by Olivia Vane*. Weblog. Available from: <https://labs.cooperhewitt.org/2017/exploring-the-cooper-hewitt-collection-with-timelines-and-tags-guest-post-by-olivia-vane/> [Accessed 28th March 2019].

Cooper Hewitt Labs. (2014) *Dataclimber explores colors in the Cooper Hewitt collection*. Weblog. Available from: <https://labs.cooperhewitt.org/2014/dataclimber-explores-colors-in-the-cooper-hewitt-collection/> [Accessed 1st April 2019].

Cooper Hewitt Labs. (2013a) *A Timeline of Event Horizons*. Weblog. Available from: <https://labs.cooperhewitt.org/2013/a-timeline-of-event-horizons/> [Accessed 11th March 2019].

Cooper Hewitt Labs. (2013b) *“All your color are belong to Giv”*. Weblog. Available from: <https://labs.cooperhewitt.org/2013/giv-do/> [Accessed 1st April 2019].

Cottrell, S. (2017) Understanding textual uncertainty in dates using interactive timelines. In: *Proceedings of the conference on Electronic Visualisation and the Arts, EVA, July 2017*. British Computer Society, pp.68-73.

- Crymble, A. (2016) Digital library search preferences amongst historians and genealogists: British History Online user survey. *Digital Humanities Quarterly*, 10(4). Available from: <http://www.digitalhumanities.org/dhq/vol/10/4/000270/000270.html> [Accessed 31st March 2019].
- Crymble, A. (2013) Review of paper machines, produced by Chris Johnson-Roberson and Jo Guldi. *Journal of Digital Humanities*, 2(1). Available from: <http://journalofdigitalhumanities.org/2-1/review-papermachines-by-adam-crymble/> [Accessed 31st March 2019].
- Cui, W., Liu, S., Tan, L., Shi, C., Song, Y., Gao, Z., Qu, H. & Tong, X. (2011) Textflow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), pp.2412-2421.
- Cui, W., Liu, S., Wu, Z. & Wei, H. (2014) How hierarchical topics evolve in large text corpora. *IEEE Transactions on Visualization and Computer Graphics*, 20(12), pp.2281-2290.
- Culler, J. (2010) The closeness of close reading. *ADE Bulletin*, 149, pp.20–25.
- Curator. (2018) *Recorded interview with Olivia Vane*, 21st May 2018.
- D3 API. (2019a) d3-force. *Github*. Available from: <https://github.com/d3/d3-force> [Accessed 10th April 2019].
- D3 API. (2019b) Power Scales. *Github*. Available from: <https://github.com/d3/d3-scale/blob/master/README.md#power-scales> [Accessed 25th March 2019].
- Darwin Correspondence Project. (2017a) *Darwin's Timeline*. Available from: <https://www.darwinproject.ac.uk/learning-resources/timeline#/> [Accessed 16th May 2019].
- Darwin Correspondence Project. (2017b) *Darwin's letters: a timeline*. Available from: <https://www.darwinproject.ac.uk/letters/darwins-letters-timeline> [Accessed 16th May 2019].
- Data-Driven Documents. (2019) Available from: <https://d3js.org/> [Accessed 5th March 2019].
- Diagne, C., Barradeau, N. & Doury, S. (2018) *t-SNE Map*. Available from: <https://experiments.withgoogle.com/t-sne-map> [Accessed 21st March 2019].
- Diagne, C. & Barradeau, N. (2017) *Free Fall*. Available from: <https://experiments.withgoogle.com/free-fall> [Accessed 21st March 2019].
- Diagne, C. & Doury, S. (2017) *Curator Table*. Available from: <https://experiments.withgoogle.com/curator-table> [Accessed 21st March 2019].
- Digital Public Library of America. (2019) Available from: <https://dp.la/> [Accessed 11th March 2019].
- DigitaltMuseum. (2019a) *Tavla*. Available from: <https://digitaltmuseum.se/011023417016/tavla> [Accessed 13th May 2019].
- DigitaltMuseum. (2019b) *Ring*. Available from: <https://digitaltmuseum.se/011023390180/ring> [Accessed 13th May 2019].

DigitaltMuseum. (2019c) *Folkkonst: By Nordiska museet*. Available from: <https://digitaltmuseum.se/011103077756/folkkonst?o=0&n=80> [Accessed 22nd May 2019].

DigitaltMuseum. (2019d) *Nordiska museet*. Available from: <https://digitaltmuseum.se/owners/S-NM> [Accessed 7th June 2019].

DigitaltMuseum. (2017) *Timeline*. Available from: <https://digitaltmuseum.org/search/?sv=timeline&aq=> [Accessed 16th May 2019].

D'Ignazio, C. & Klein, L. (2019) Chapter Five: The Numbers Don't Speak for Themselves. *Data Feminism*. MIT Press. [Manuscript draft]. Available from: <https://bookbook.pubpub.org/pub/6ui5n4vo> [Accessed: 8th April 2019].

Dörk, M., Pietsch, C. & Credico, G. (2017) One view is not enough: High-level visualizations of a large cultural collection. *Information Design Journal*, 23(1), July 2017, pp.39-47.

Douglass, J., Huber, W. & Manovich, L. (2011) Understanding scanlation: how to read one million fan-translated manga pages. *Image and Narrative*, 12(1). Available from: http://softwarestudies.com/Understanding_scanlation_2011.pdf [Accessed 22nd March 2019].

Doury, S. & Ferrier, E. (2018) *Art Palette*. Available from: <https://experiments.withgoogle.com/art-palette> [Accessed 1st April 2019].

Drucker, J. (2011) Humanities approaches to graphical display. *Digital Humanities Quarterly (DHQ)*, 5(1), pp.1-21.

Drucker, J. (2010) Graphesis: Visual knowledge production and representation. *Poetess Archive Journal*, 2(1), pp.1-50. Available from: <https://journals.tdl.org/paj/index.php/paj/article/view/4/50> [Accessed 4th March 2019].

Drucker, J. (2009) *SpecLab: digital aesthetics and projects in speculative computing*. University of Chicago Press.

Drucker, J. & Nowviskie, B. (2004) Speculative computing: Aesthetic provocations in humanities computing. In: Schreibman, S., Siemens, R. & Unsworth, J. (eds.) *A companion to digital humanities*. Blackwell Publishing Ltd, pp.431-447.

Drucker, J. & Nowviskie, B. (2003) The Temporal Modelling Project. Available from: <http://www2.iath.virginia.edu/time/project/> [Accessed 24th April 2019].

Eames Office. (2014) *A Computer Perspective*. Available from: <http://www.eamesoffice.com/the-work/a-computer-perspective/> [Accessed 31st March 2019].

Eklund, A. (2016) *beeswarm: The Bee Swarm Plot, an Alternative to Stripchart*. R package version 0.2.3. Available from: <https://cran.r-project.org/web/packages/beeswarm/beeswarm.pdf> [Accessed 25th March 2019].

Elasticsearch. (2019a) Available from: <https://www.elastic.co/products/elasticsearch> [Accessed 5th March 2019].

Elasticsearch. (2019b) *Nested datatype*. Available from: <https://www.elastic.co/guide/en/elasticsearch/reference/current/nested.html> [Accessed 1st April 2019].

European Commission. (2019) *Cultural Heritage: Digitisation, online accessibility and digital preservation. Consolidated Progress Report on the implementation of Commission Recommendation (2011/711/EU) 2015-2017*. Available from: https://ec.europa.eu/digital-single-market/en/news/european-commission-report-cultural-heritage-digitisation-online-accessibility-and-digital?utm_source=Cultural+Digital&utm_campaign=91a25a8a18-culturaldigital&utm_medium=email&utm_term=0_f5c318bb03-91a25a8a18-151297413 [Accessed 18th June 2019].

Europeana. (2019) Available from: <http://www.europeana.eu/portal/en> [Accessed 11th March 2019].

Fathom. (2019) *Porfiry*. Available from: <https://fathom.info/porfiry/> [Accessed 25th March 2019].

Fischer, M. (1966) The KWIC index concept: A retrospective view. *Journal of the Association for Information Science and Technology*. *JASIST*, 17(2), pp.57-70.

Foo, B. (2016) *NYPL Public Domain Visualization*. Available from: <http://publicdomain.nypl.org/pd-visualization/> [Accessed 20th March 2019].

Frayling, C. (1993) Research in Art and Design. *Royal College of Art Research Papers*. 1(1), 1-5. Available from: http://researchonline.rca.ac.uk/384/3/frayling_research_in_art_and_design_1993.pdf [Accessed 4th March 2019].

Friendly, M. (2008) The golden age of statistical graphics. *Statistical Science*, 23(4), pp.502-535.

Galey, A. & Ruecker, S. (2010) How a Prototype Argues. *Literary and Linguistic Computing*, 25(4), pp.405-424.

Gibbs, F. (2016) New forms of history: Critiquing data and its representations. *The American Historian*, 7. Available from: <https://tah.oah.org/february-2016/new-forms-of-history-critiquing-data-and-its-representations/> [Accessed 20th May 2019].

Gibbs, F. & Owens, T. (2012) Building better digital humanities tools: toward broader audiences and user-centered designs. *Digital Humanities Quarterly (DHQ)*, 6(2). Available from: <http://www.digitalhumanities.org/dhq/vol/6/2/000136/000136.html> [Accessed 31st March 2019].

Gist. (2018) *Knoll Archive Timeline*. Available from: <https://knoll.gistapp.com/archive> [Accessed 3rd June 2019].

Glinka, K., Pietsch, C. & Dörk, M. (2017) Past Visions and Reconciling Views: Visualizing Time, Texture and Themes in Cultural Collections. *DHQ: Digital Humanities Quarterly*, 11(2).

Goldschmidt, G. (1991) The Dialectics of Sketching. *Creativity Research Journal*, 4(2), pp.123-143.

Google Arts & Culture. (2019) Available from: <https://artsandculture.google.com/> [Accessed 31st March 2019].

Gortana, F., von Tenspolde, F., Guhlmann, D. & Dörk, M. (2018) Off the Grid: Visualizing a Numismatic Collection as Dynamic Piles and Streams. *Open Library of Humanities*, 4(2), p.30.

Graham, E. (2017) Introduction: Data visualisation and the humanities. *English Studies*, 98(5), pp.449-458.

- Grenci, J. (2014) *Around the World with the Library of Congress Poster Collection: An Overview of Travel Posters*. Available from: <https://www.loc.gov/rr/print/coll/travel-posters.html> [Accessed 26th March 2019].
- Groeger, L. & Currier, C. (2012) Stacking Up the Administration's Drone Claims, *ProPublica*, 13th September, 2012. Available from: <https://projects.propublica.org/graphics/cia-drones-strikes> [Accessed 31st March 2019].
- Grossner, K. & Meeks, E. (2013) *Topotime v0.1 gallery & sandbox*. Available at: <http://dh.stanford.edu/topotime/> [Accessed 24th April 2019].
- Guldi, J. & Johnson-Roberson, C. (2019) *Paper Machines*. Available from: <http://papermachines.org/> [Accessed 31st March 2019].
- Hai-Jew, S. (ed.) (2017) *Data analytics in digital humanities*. Springer.
- Harkness, N. (2006) The colour wheels of art, perception, science and physiology. *Optics & Laser Technology*, 38(4-6), pp.219-229.
- Haskiya, D. (2019) An Evaluation of Generous Interfaces. *EuropeanaTech Insight*, 11. Available from: <https://pro.europeana.eu/page/issue-11-generous-interfaces#an-evaluation-of-generous-interfaces> [Accessed 11th June 2019].
- Havre, S., Hetzler, B., & Nowell, L. (2000) ThemeRiver: Visualizing theme changes over time. In: *Proceedings of Symposium on Information Visualization, InfoVis 2000*. IEEE, pp. 115-123.
- Hjørland, B. (2008) What is Knowledge Organization (KO)? *Knowledge Organization*, 35(2-3), pp.86-101.
- Hinchcliffe, G. (2016) *Tate Explorer*. Available from: <http://gravitron.com.au/tate/> [Accessed 31st March 2019].
- Hinrichs, U., Alex, B., Clifford, J., Watson, A., Quigley, A., Klein, E. & Coates, C. M. (2015) Trading consequences: A case study of combining text mining and visualization to facilitate document exploration. *Digital Scholarship in the Humanities*, 30(Supplement 1), pp.i50-i75.
- Hinrichs, U., El-Assady, M., Bradely, A. J., Forlini, S., & Collins, C. (2017) Risk the drift! Stretching disciplinary boundaries through critical collaborations between the humanities and visualization. In: *2nd Workshop on Visualization for the Digital Humanities (VIS4DH), held as part of VisWeek'17*, Phoenix, United States, 2nd October 2017.
- Hinrichs, U., Forlini, S. & Moynihan, B. (2018) In Defense of Sandcastles: Research Thinking through Visualization in Digital Humanities. *Digital Scholarship in the Humanities (DSH), Advance Articles*, [fqy051].
- Historian. (2017) *Recorded interview with Olivia Vane*, 26th January 2017.
- Hitchcock, T. (2014) Big Data, Small Data and Meaning. *Historyonics*. Weblog. Available from: http://historyonics.blogspot.com/2014/11/big-data-small-data-and-meaning_9.html [Accessed 7th July 2019].

Hochman, N. & Manovich, L. (2013) Zooming into an Instagram City: Reading the local through social media. *First Monday*, 18(7). Available from: <https://firstmonday.org/ojs/index.php/fm/article/view/4711/3698> [Accessed 22nd March 2019].

Hochman, N. & Manovich, L. (2014) A View from Above: Exploratory Visualizations of the Thomas Walther Collection. In: Abbaspour, M., Daffner, L.A. & Hambourg, M.M. (eds.) *Object:Photo. Modern Photographs: The Thomas Walther Collection 1909–1949. An Online Project of The Museum of Modern Art*. New York, The Museum of Modern Art. Available from: https://www.moma.org/interactives/objectphoto/assets/essays/Manovich_Hochman.pdf [Accessed 31st March 2019].

Hockey, S. (2007) The History of Humanities Computing. In: Schreibman, S., Siemens, R., & Unsworth, J. (eds.) *A Companion to Digital Humanities*. Blackwell Publishing, Oxford, UK, pp.3-19.

Hooper-Greenhill, E. (2000) Changing values in the art museum: Rethinking communication and learning. *International Journal of Heritage Studies*, 6(1), pp.9-31.

Hooper Greenhill, E. (1992) *Museums and the Shaping of Knowledge*. Routledge.

Huckle, K. (2018) *GLAMi Nomination: Algorithms + Artists: New Ways Into The RA Collection Online*. Available from: <https://mw18.mwconf.org/glami/algorithms-artists-new-ways-into-the-ra-collection-online/> [Accessed 21st March 2019].

Hughes, L. M. (2004) *Digitizing collections: strategic issues for the information manager*. Facet Publishing.

Hullman, J., Adar, E. & Shah, P. (2011) Benefitting infovis with visual difficulties. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), pp.2213-2222.

Hullman, J. & Diakopoulos, D. (2011) Visualization Rhetoric: Framing Effects in Narrative Visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 17(12), pp.2231-2240.

Hyvönen, E. (2012) Publishing and using cultural heritage linked data on the semantic web. In: Ding, Y. & Groth, P. (eds.) *Synthesis Lectures on the Semantic Web: Theory and Technology*, 2(1), pp.1-159.

ImageNet. (2019) Available from: <http://www.image-net.org/> [Accessed 10th May 2019].

International Organization for Standardization. (2019) *ISO 8601-2:2019*. Available from: <https://www.iso.org/standard/70908.html> [Accessed 24th April 2019].

Jänicke, S. (2018a) Time-based Impact Mosaics. In: *Workshop on Visualization for Communication (Viscomm), held as part of VisWeek'18, Berlin, Germany, 21st October 2018*. Available from: <https://viscomm.io/posters/timemosaics.pdf> [Accessed 15th May 2019].

Jänicke, S. (2018b) Timages: Enhancing Time Graphs with Iconographic Information. In: *Leipzig Symposium on Visualization in Applications, Leipzig, Germany, 18-19th October 2018*. Available from: https://pdfs.semanticscholar.org/8cdb/4c4a3a9cfd04c04b3844955ddac9323b1bfd.pdf?_ga=2.196554984.539264501.1557248545-1766844548.1550053919 [Accessed 15th May 2019].

- Jänicke, S. (2016) Valuable Research for Visualization and Digital Humanities: A Balancing Act. In: *1st Workshop on Visualization for the Digital Humanities (VIS4DH), held as part of VisWeek'16, Baltimore, United States, 24th October 2016.*
- Jänicke, S., Franzini, G., Cheema, M.F. & Scheuermann, G. (2015) On close and distant reading in digital humanities: A survey and future challenges. In: *Eurographics Conference on Visualization (EuroVis) - STARS.* The Eurographics Association.
- Janßen, S. (2015) Vom motiv bis zur wirkung. *Städelblog.* Weblog. Available from: <http://blog.staedelmuseum.de/vom-motiv-bis-zur-wirkung-schlagwoerter-fuer-die-digitale-sammlung/> [Accessed 21st March 2019].
- K-samsök. (2019a) Available from: <http://www.ksamsok.se/> [Accessed 23rd April 2019].
- K-samsök. (2019b) *API.* Available from: <http://www.ksamsok.se/in-english/api/> [Accessed 12th May 2019].
- Kassell, L., Hawkins, M., Ralley, R., Young J., Edge, J., Martin-Portugues, J. Y., & Kaoukji, N. (eds.). (2019) 'Casebooks', *The casebooks of Simon Forman and Richard Napier, 1596–1634: a digital edition.* Available from: <https://casebooks.lib.cam.ac.uk> [Accessed 4th March 2019].
- Kelton, K., Fleischmann, K. R. & Wallace, W. A. (2008) Trust in digital information. *Journal of the American Society for Information Science and Technology*, 59(3), pp.363–374.
- Kenderdine, S. & Hart, T. (2014) mARChive: Sculpting Museum Victoria's Collections. *MW2014: Museums and the Web 2014.* Available from: <https://mw2014.museumsandtheweb.com/paper/marchive-sculpting-museum-victorias-collections/> [Accessed 21st March 2019].
- Keras. (2019) *VGG16.* Available from: <https://keras.io/applications/#vgg16> [Accessed 6th May 2019].
- King, L., Stark, J. F. & Cooke, P. (2016) Experiencing the digital world: The cultural value of digital engagement with heritage. *Heritage & Society*, 9(1), pp.76-101.
- Klein, L. F. (2013) The Image of Absence: Archival Silence, Data Visualization, and James Hemings. *American Literature*, 85(4), pp.661-688.
- Knight, B. (2016) *kuler-d3/colorwheel.js*, *Github.* Available from: <https://github.com/benknight/kuler-d3/blob/master/colorwheel.js> [Accessed 1st April 2019].
- Koenig, B. (2003) *Color workbook.* Upper Saddle River, N.J, Prentice Hall.
- Kosara, R. & Mackinlay, J. (2013) Storytelling: The next step for visualization. *Computer*, 46(5), pp.44-50.
- Kräutli, F. (2016a) *Britten's Poets.* Available from: <http://www.kraeutli.com/index.php/2016/02/15/brittens-poets/> [Accessed 25th March 2019].
- Kräutli, F. (2016b) *Visualising Cultural Data: Exploring Digital Collections Through Timeline Visualisations.* PhD Thesis, Royal College of Art, London. Available from: <http://researchonline.rca.ac.uk/1774/> [Accessed 4th March 2019].

Kräutli, F. (2016c) *Timeline Tools*. Available from: <http://www.kraeutli.com/index.php/2016/04/08/timeline-tools/> [Accessed 31st March 2019].

Kräutli, F. & Boyd Davis, S. (2016) Digital Humanities Research Through Design. In: *DH Early Career Conference 'Mapping the Scope and Reach of the Digital Humanities'*, King's College London, 20 May 2016. Available from: <http://researchonline.rca.ac.uk/1787/1/Krautli%20Digital%20Humanities%20Research%20Through%20Design%202016.pdf> [Accessed 4th March 2019].

Kräutli, F. & Boyd Davis, S. (2013) Known Unknowns: representing uncertainty in historical time. In: *Proceedings of the Conference on Electronic Visualisation and the Arts, EVA 2013, London, UK*. British Computer Society. pp.61-68.

Krug, K., Lerner, M. & Wucher, S. (2018) *Forgotten Heritage*. Available from: <https://www.forgottenheritage.eu/artworks> [Accessed 17th May 2019].

Lam, H., Bertini, E., Isenberg, P., Plaisant, C. & Carpendale, S. (2011) Seven guiding scenarios for information visualization evaluation. *Technical Report, 2011-992-04*, Department of Computer Science, University of Calgary.

Lamqaddam, H., Brosens, K., Truyen, F., Beerens, J., De Prekel, I. & Verbert, K. (2018) When the Tech Kids are Running Too Fast: Data Visualisation Through the Lens of Art History Research. To be published in *IEEE Transactions on Visualization and Computer Graphics*. [Preprint] Available from: https://www.researchgate.net/publication/328191886_When_the_Tech_Kids_are_Running_Too_Fast_Data_Visualisation_Through_the_Lens_of_Art_History_Research [Accessed: 31st March 2019].

Lee, B., Riche, N. H., Isenberg, P. & Carpendale, S. (2015) More than telling a story: Transforming data into visually shared stories. *IEEE Computer Graphics and Applications (CG&A)*, 35(5), pp.84-90.

Lester, R. (2001) The convergence of museums and libraries?. *Alexandria*, 13(3), pp.183-191.

Library of Congress. (2019) *Extended Date/Time Format (EDTF) Specification*. Available from: <https://www.loc.gov/standards/datetime/edtf.html> [Accessed 23rd April 2019].

Lindley, S., Corish, R., Kosmack Vaara, E., Ferreira, P. & Simbelis, V. (2013) Changing perspectives of time in HCI. In: *CHI'13 Extended Abstracts on Human Factors in Computing Systems, April 2013*. ACM, pp. 3211-3214.

Liu, S., Zhou, M. X., Pan, S., Song, Y., Qian, W., Cai, W. & Lian, X. (2012) Tiara: Interactive, topic-based visual text summarization and analysis. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(2), 25.

Lloyd, D. & Dykes, J. (2011) Human-Centered Approaches in Geovisualization Design: Investigating Multiple Methods Through a Long-Term Case Study. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), pp.2498-2507.

London, M. (2005) History of addiction: a UK perspective. *American Journal on Addictions*, 14(2), pp.97-105.

Lorber-Kasunic, J. & Sweetapple, K. (2015) Visualising text-based data: Identifying the potential of visual knowledge production through design practice. *Studies in Material Thinking*, 13, pp.3-19.

- Loukissas, Y. A. (2019) *All Data Are Local: Thinking Critically in a Data-Driven Society*. MIT Press.
- Loukissas, Y. A. (2017). Taking Big Data apart: local readings of composite media collections. *Information, Communication & Society*, 20(5), pp.651-664.
- Loukissas, Y. A. (2014) *The Life and Death of Data*. Online interactive publication. Available from: <http://lifeanddeathofdata.org/> [Accessed 2nd April 2019].
- Lourdi, I. & Nikolaidou, M. (2009) Guidelines for Developing Digital Cultural Collections. In: Theng Y. L., Foo, S., Goh, D. & Na, J. C. (eds.) *Handbook of Research on Digital Libraries: Design, Development, and Impact*. IGI Global. pp.198-205.
- Maaten, L. V. D. & Hinton, G. (2008) Visualizing data using t-SNE. *Journal of machine learning research*, 9, pp.2579-2605.
- Manovich, L. (2011a) What is Visualization? *Visual Studies*, 26(1), pp.36-49. Available from: http://manovich.net/content/04-projects/064-what-is-visualization/61_article_2010.pdf [Accessed 21st March 2019].
- Manovich, L. (2011b) Media visualization: Visual techniques for exploring large media collections. In: Gates, K. (ed.) *Media Studies Futures*. New York, Blackwell. Available from: <http://manovich.net/content/04-projects/069-media-visualization-visual-techniques-for-exploring-large-media-collections/66-article-2011.pdf> [Accessed 20th March 2019].
- Manovich, L. (1999) Database as a symbolic form. *Convergence*, 5(2). Sage Publications. pp.80-99.
- Manovich, L. & Douglass, J. (2009) TimeLine. Available from: <http://manovich.net/index.php/exhibitions/timeline> [Accessed 20th March 2019].
- Manovich, L., Douglass, J. & Huber, W. (2010) *Mapping Time: Visualizations of temporal patterns in media and art*. Available from: <http://lab.softwarestudies.com/2010/10/mappingtime.html> [Accessed 31st March 2019].
- Marche, S. (2012) Literature is not data: Against digital humanities. *LA Review of Books*, 28th October 2012. Available from: <https://lareviewofbooks.org/article/literature-is-not-data-against-digital-humanities/> [Accessed 11th March 2019].
- Marsh, S. & Dibben, M. R. (2003) The role of trust in information science and technology. *Annual Review of Information Science and Technology* 37, 1, pp.465-498.
- Mayr, E., Hynek, N., Salisu, S. & Windhager F. (2019) Trust in Information Visualization. In: Kosara, R., Lawonn, K., Linsen, L. & Smit, N. (eds.), *EuroVis Workshop on Trustworthy Visualization (TrustVis) 2019*. The Eurographics Association, Porto, pp.25-29.
- McCurdy, N., Dykes, J. & Meyer, M. (2016) Action Design Research and Visualization Design. In: *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization (BELIV 2016)*, Baltimore, United States, 24th October 2016, pp.10-18.
- McCurdy, N., Lein, J., Coles, K. & Meyer, M. (2015) Poemage: Visualizing the sonic topology of a poem. *IEEE transactions on visualization and computer graphics*, 22(1), pp.439-448.

McCloskey, M. & Wei, S. (2014) Bud's Story, from the Records, *ProPublica*, March 6, 2014. Available from: <https://projects.propublica.org/graphics/mia-letters> [Accessed 31st March 2019].

Medical Officer of Health for Islington Metropolitan Borough. (1967) *Report of the Medical Officer of Health for Islington Metropolitan Borough*. Available from: <http://wellcomelibrary.org/moh/report/b1823947x> [Accessed 17th April 2019].

Medical Officer of Health for Kensington Borough. (1957) *Report of the Medical Officer of Health for Kensington Borough*. Available from: <https://wellcomelibrary.org/moh/report/b18236042> [Accessed 12th April 2019].

Meirelles, I. (2013) *Design for information: an introduction to the histories, theories, and best practices behind effective information visualizations*. Rockport publishers.

Metropolitan Museum of Art. (2000) *The Met's Heilbrunn Timeline of Art History*. Available from: <http://www.metmuseum.org/toah/> [Accessed 31st March 2019].

MOMA. (2012) *Inventing Abstraction*. Available from: <https://www.moma.org/interactives/exhibitions/2012/inventingabstraction/?page=connections> [Accessed 12th June 2019].

Moretti, F. (2005) *Graphs, maps, trees: abstract models for a literary history*. Verso.

Morse, C., Koenig, V., Lallemand, C. & Wieneke, L. (2019) Art in Rich-Prospect: Evaluating Next-Generation User Interfaces for Cultural Heritage. In: *Museums and the Web 2019 (MW19)*. Available from: <https://mw19.mwconf.org/paper/art-in-rich-prospect-evaluating-next-generation-user-interfaces-for-cultural-heritage-2/> [Accessed 6th March 2019].

Mucha Foundation. (2012) *Alphonse Mucha Timeline*. Available from: <http://www.muchafoundation.org/timeline> [Accessed 31st March 2019].

Münster, S. & Terras, M. (2019) The visual side of digital humanities: a survey on topics, researchers, and epistemic cultures. *Digital Scholarship in the Humanities*. [fqz022]. Available from: <https://academic.oup.com/dsh/advance-article/doi/10.1093/llc/fqz022/5485806> [Accessed 7th May 2019].

Munzner, T. (2014) *Visualization Analysis and Design*. New York, CRC Press.

Museum Computer Group. (2019) Collections online: the rationale. *MCG community discussion list*, 25th Feb 2019. Available from: <https://www.jiscmail.ac.uk/cgi-bin/webadmin?A1=ind1902&L=MCG#8> [Accessed 2nd May 2019].

Museums and the Web. (2019) *Bibliography*. Available from: <https://www.museweb.net/bibliography/> [Accessed 18th May 2019].

Nasjonalmuseet. (2017a) *Machine vision on the collection*. Available from: http://www.nasjonalmuseet.no/en/collections_and_research/our_collections/Machine+vision+on+the+collection.b7C_wjWXh.ips [Accessed 31st March 2019].

Nasjonalmuseet. (2017b) *Project: 'Principal Components'*. Available from: http://www.nasjonalmuseet.no/en/collections_and_research/collection_management/digital_collection_management/Project%3A+%C2%ABPrincipal+Components%C2%BB.b7C_wjU4L.ips [Accessed 31st March 2019].

- Nordic Museum curator. (2019) *Email correspondence with Olivia Vane*, 3rd June 2019.
- Nordic Museum curator. (2018) *Recorded interview with Olivia Vane*, 19th November 2018.
- Olah, et al. (2018) The Building Blocks of Interpretability, *Distill*. Available from: <https://distill.pub/2018/building-blocks/> [Accessed 6th May 2019].
- Olson, H.A. (2001) The power to name: Representation in library catalogs. *Signs*, 26(3), pp.639-668.
- Oomen, J. & Aroyo, L. (2011) Crowdsourcing in the cultural heritage domain: opportunities and challenges. In: *Proceedings of the 5th International Conference on Communities and Technologies, C&T'11, 29th June–2nd July 2011, QUT, Brisbane, Australia*. ACM. pp.138-149.
- Owens, T. (2011) Defining Data for Humanists: Text, Artifact, Information or Evidence? *Journal of Digital Humanities*, 1(1). Available from: <http://journalofdigitalhumanities.org/1-1/defining-data-for-humanists-by-trevor-owens/> [Accessed 11th March 2019].
- Owens, T. (2011) Please Write it Down: Design and Research in Digital Humanities. *Journal of Digital Humanities*, 1(1).
- Parvaneh, G. (2013) *RoyGBiv*. Available from: <https://github.com/givp/RoyGBiv> [Accessed 1st April 2019].
- Petersen, M. G., Iversen, O. S., Krogh, P. G. & Ludvigsen, M. (2004) Aesthetic Interaction — A Pragmatist's Aesthetics of Interactive Systems. In: *Proceedings of the Conference on Designing Interactive Systems, DIS 2004, 1st-4th August 2004, Cambridge, Massachusetts, USA*. ACM. pp.269-276.
- PhotoSwipe. (2019) <https://photoswipe.com/> [Accessed 12th May 2019].
- Pietsch, C. (2017) *VIKUS Viewer*. Available from: <https://vikusviewer.fh-potsdam.de/> [Accessed 19th May 2019].
- Pim, H. (2018) Exploring Wellcome Collection with computer vision. *Wellcome Collection*. Weblog. Available from: <https://stacks.wellcomecollection.org/exploring-wellcome-collection-with-computer-vision-7513dff8126d> [Accessed 2nd May 2019].
- Post, R. (2018) *TextContext*. Available from: <http://setfilter.net/> [Accessed 12th April 2019].
- Pousman, Z., Stasko, J. & Mateas, M. (2007) Casual information visualization: Depictions of data in everyday life. *IEEE transactions on visualization and computer graphics*, 13(6), pp.1145-1152.
- Preziosi, D. (2004) Brain of the earth's body: Museums and the framing of modernity. In: Carbonell, B. M. (ed.) *Museum Studies: An anthology of contexts*. Malden, Mass., Blackwell. pp.71-84.
- Rabinowitz, A. (2014) It's about time: historical periodization and Linked Ancient World Data. *ISAW Papers*, 7(22). Available from: <http://dlib.nyu.edu/awdl/isaw/isaw-papers/7/rabinowitz/> [Accessed 3rd May 2019].
- Rabinowitz, A. & Shaw, R. (2019) *PeriodO*. Available from: <https://perio.do/en/> [Accessed 23rd April 2019].

Ramsay, S. & Rockwell, G. (2012) *Developing Things: Notes toward an Epistemology of Building in the Digital Humanities*. In: Gold, M. & Klein, L. (eds.) *Debates in the Digital Humanities*. Minneapolis: University of Minnesota Press. pp. 75-84.

Resig, J. (2014) *Node.js Year Range Parser*. Available from: <https://github.com/jeresig/node-yearrange> [Accessed 21st March 2019].

Reyes, O. (2015) *Paul Klee viz*. Available from: <http://ereyes.net/kleeviz/> [Accessed 1st April 2019].

Riley, J. (2017) *Understanding Metadata: What is Metadata, and What is it For?* NISO (National Information Standards Organization) Primer Series. Available from: https://groups.niso.org/apps/group_public/download.php/17443/understanding-metadata [Accessed 11th March 2019].

Robinson, H. (2014) Knowledge utopias: An epistemological perspective on the convergence of museums, libraries and archives. *Museum and Society*, 12(3), pp.210-224.

Robinson, H. (2012) Remembering things differently: Museums, libraries and archives as memory institutions and the implications for convergence. *Museum Management and Curatorship*, 27(4), pp.413-429.

Rosenberg, D. & Grafton, A. (2010) *Cartographies of time: A history of the timeline*. Princeton Architectural Press.

Royal Photographic Society. (2017) *RPS Collection expands the V&A collection of photography*. Weblog. Available from: <http://www.rps.org/news/2017/april/rps-collection-expands-the-vanda-collection-of-photography> [Accessed 6th May 2019].

Rudd, J., Stern, K. and Isensee, S. (1996) Low vs. high-fidelity prototyping debate. *interactions*, 3(1). pp.76-85.

Ruecker, S. (2015) A Brief Taxonomy of Prototypes for the Digital Humanities. *Scholarly and Research Communication*, 6(2).

Ruecker, S., Radzikowska, M. & Sinclair, S. (2011) *Visual interface design for digital cultural heritage: A guide to rich-prospect browsing*. Ashgate Publishing.

Sarkar, M. & Brown, M. H. (1992) Graphical fisheye views of graphs. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, pp. 83-91.

Schmidt, B. (2018) Modelling Time. In: Flanders, J. & Jannidis, F. (eds.) *The Shape of Data in Digital Humanities: Modeling Texts and Text-based Resources*. Routledge.

Schmidt, B. (2012) Words alone: Dismantling topic models in the humanities. *Journal of Digital Humanities*, 2(1). Available from <http://journalofdigitalhumanities.org/2-1/words-alone-by-benjamin-m-schmidt/> [Accessed 31st March 2019].

Schnapp, J., Presner, T., Lunenfeld, P. & Drucker, J. (2009) The Digital Humanities Manifesto 2.0. In: *UCLA Mellon Seminar on the Digital Humanities ("What Is(n't) Digital Humanities?")*. Available from: http://jeffreyschnapp.com/wp-content/uploads/2011/10/Manifesto_V2.pdf [Accessed 31st March 2019].

Schofield, T., Whitelaw, M. & Kirk, D. (2017) Research through design and digital humanities in practice: What, how and who in an archive research project. *Digital Scholarship in the Humanities*, 32(Issue suppl_1), pp.i103–i120.

Science Museum Group. (2019) *Babbage, Charles 1791 - 1871*. Available from: <https://collection.sciencemuseum.org.uk/people/ap8/babbage-charles> [Accessed 15th May 2019].

scikit-learn. (2019a) *Principal component analysis (PCA)*. Available from: <https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.PCA.html> [Accessed 11th May 2019].

scikit-learn. (2019b) *t-distributed Stochastic Neighbor Embedding (t-SNE)*. Available from: <https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html> [Accessed 11th May 2019].

SciPy. (2019) *scipy.spatial.distance.cosine*. Available from: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.cosine.html#scipy.spatial.distance.cosine> [Accessed 11th May 2019].

Segel, E. & Heer, J. (2010) Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 16(6), pp.1139-1148.

Shneiderman, B. (1996) The eyes have it: A task by data type taxonomy for information visualizations. In: *Proceedings of the Symposium on Visual Languages, 3rd-6th September 1996, Boulder, USA*. IEEE.

Silva, S. F. & Catarci, T. (2000) Visualization of linear time-oriented data: a survey. In: *Proceedings of the First International Conference on Web Information Systems Engineering, WISE'2000, Hong Kong, June 2000, 1*. IEEE, pp.310-319.

Simonyan, K. & Zisserman, A. (2015) Very Deep Convolutional Networks for Large-Scale Image Recognition. In: *Proceedings of the 3rd International Conference on Learning Representations, ICLR2015*.

Sinclair, S. & Rockwell, G. (2019) *Voyant Tools*. Available from: <https://voyant-tools.org/> [Accessed 31st March 2019].

Sinclair, S., Ruecker, S. & Radzikowska, M. (2013) Information Visualization for Humanities Scholars. *Literary Studies in the Digital Age: An Evolving Anthology*. Available from: <https://dlsanthology.mla.hcommons.org/information-visualization-for-humanities-scholars/> [Accessed 4th March 2019].

Smithies, J. (2012) Evaluating Scholarly Digital Outputs: The Six Layers Approach. *Journal of Digital Humanities*, 1(4). Available from: <http://journalofdigitalhumanities.org/1-4/evaluating-scholarly-digital-outputs-by-james-smithies/> [Accessed 4th March 2019].

Smithsonian Digitization Program Office. (2019) *2019 Smithsonian Digitization Conference - Welcome*. Available from: <https://dpo.si.edu/2019-smithsonian-digitization-conference-welcome> [Accessed 6th July 2019].

Smithsonian Institution Archives. (2019) *Cooper Hewitt, Smithsonian Design Museum*. Available from: <https://siarchives.si.edu/history/cooper-hewitt-national-design-museum> [Accessed 25th March 2019].

Society of American Archivists. (2019) *A Glossary of Archival and Records Terminology: 'Collection'*. Available from: <https://www2.archivists.org/glossary/terms/c/collection> [Accessed 11th March 2019].

softwarestudies.com. (2009) Timeline: 4535 Time magazine covers, 1923-2009. *Flickr*. Available from: <https://www.flickr.com/photos/culturevis/3951496507/in/set-72157624959121129> [Accessed 31st March 2019].

software studies initiative. (2011) *Mondrian vs Rothko: footprints and evolution in style space*. Available from: <http://lab.softwarestudies.com/2011/06/mondrian-vs-rothko-footprints-and.html> [Accessed 1st April 2019].

Speakman, R., Hall, M.M. & Walsh, D. (2018) User Engagement with Generous Interfaces for Digital Cultural Heritage. In: Méndez, E., Crestani, F., Ribeiro, C., David, G. & Lopes, J. (eds.) *Digital Libraries for Open Knowledge, TPDL 2018, Lecture Notes in Computer Science, vol 11057*. Springer, pp.186-191.

Staniszewski, M. (1998) *The Power of Display: A History of Exhibition Installations at the Museum of Modern Art*. Cambridge, MA, MIT Press.

Stefaner, M. (2010) *revisit*. Available from: <http://truth-and-beauty.net/projects/revisit/> [Accessed 21st March 2019].

Terras, M. (2015) Opening Access to collections: the making and using of open digitised cultural content. *Online Information Review*, 39(5), pp.733–752.

Terras, M., Nyhan, J. & Vanhoutte, E. (2013) *Defining digital humanities: a reader*. London, Routledge.

Thompson, P., Batista-Navarro, R.T., Kontonatsios, G., Carter, J., Toon, E., McNaught, J., Timmermann, C., Worboys, M. & Ananiadou, S. (2016) Text mining the history of medicine, *PLoS ONE*, 11(1), e0144717.

Thorp, J. (2018) *A Library of Color*. Available from: <https://library-of-colors.glitch.me/> [Accessed 1st April 2019].

Thudt, A., Hinrichs, U. & Carpendale, S. (2012) The Bohemian Bookshelf : supporting serendipitous book discoveries through information visualization. In: *Proceedings of the Conference on Human Factors in Computing Systems, CHI 2012, May 2012*. ACM. pp.1461-1470.

Toon, E., Timmermann, C. & Worboys, M. (2016) Text-Mining and the History of Medicine: Big Data, Big Questions? *Medical History*, 60(2), pp.294-296.

Trant, J. (2009) Tagging, folksonomy and art museums: Early experiments and ongoing research. *Journal of Digital Information*, 10(1).

Tufte, E. (1983) *The Visual Display of Quantitative Information*. Cheshire, CT, Graphics Press.

Tunkelang, D. (2009) Faceted search. In: Marchionini, G. (ed.) *Synthesis lectures on information concepts, retrieval, and services*. Morgan & Claypool, pp.1-80. Available from: <http://www.iro.umontreal.ca/~nie/IFT6255/Books/FacetedSearch.pdf> [Accessed 31st March 2019].

- Urban Complexity Lab. (2017) *Coins – a journey through a rich cultural collection*. Available from: <https://uclab.fh-potsdam.de/coins/> [Accessed 31st March 2019].
- Urban Complexity Lab. (2016) *Past Visions: penned by Frederick William IV*. Available from: <https://uclab.fh-potsdam.de/fw4/en/> [Accessed 5th March 2019].
- Urban Complexity Lab. (2015) *Past Visions*. Available from: <https://uclab.fh-potsdam.de/fw4/> [Accessed 31st March 2019].
- Vander Wal, T. (2007) Folksonomy coinage and definition. *vanderwal.net*. Weblog. Available from: <http://vanderwal.net/folksonomy.html> [Accessed 22nd March 2019].
- V&A Collections. (2019) *Lacock Abbey towards Sharrington's Tower from the south-east*. Available from: <http://collections.vam.ac.uk/item/O1387805/lacock-abbey-towards-sharringtons-tower-photograph-talbot-william-henry/> [Accessed 10th May 2019].
- V&A Museum. (2019a) *Royal Photographic Society Collection*. Available from: <https://collections.vam.ac.uk/search/?limit=15&q=royal+photographic+society+collection&commit=Search&narrow=1&offset=0&slug=0> [Accessed 6th May 2019].
- V&A Museum. (2019b) *Photographic processes*. Available from: <https://www.vam.ac.uk/articles/photographic-processes> [Accessed 6th May 2019].
- V&A Museum. (2018) *Our new Photography Centre is now open*. Available from: <https://www.vam.ac.uk/blog/news/our-new-photography-centre-is-now-open> [Accessed 9th May 2019].
- V&A Museum. (2010) *Fabric Visualiser*. Available from: http://collections.vam.ac.uk/information/information_fabricvisualiser [Accessed 21st March 2019].
- V&A curator. (2019) *Recorded interview with Olivia Vane*, 4th April 2019.
- V&A curator. (2018) *Recorded interview with Olivia Vane*, 11th July 2018.
- V&A digital media staff. (2018) *Recorded interview with Olivia Vane*, 26th June 2018.
- Viégas, F. B. & Wattenberg, M. (2007) Artistic data visualization: Beyond visual analytics. In: *Online Communities and Social Computing (OCSC 2007), held as part of HCI International 2007, Beijing, China, July 2007*, pp.182-191.
- VIS4DH. (2019) Available from: <http://www.vis4dh.org/> [Accessed 10th July 2019].
- visualising data. (2011) *Introducing ImagePlot visualisation software*. Available from: <http://www.visualisingdata.com/2011/09/introducing-imageplot-visualisation-software/> [Accessed 31st March 2019].
- Wallace, A. (2001) Collections Management and Inclusion. In: Dodd, J. & Sandell, R. *Including museums: Perspectives on museums, galleries and social inclusion*. Research Centre for Museums and Galleries, University of Leicester, UK, pp.80-83.
- Wattenberg, M. & Viégas, F. B. (2008) The word tree, an interactive visual concordance. *IEEE Transactions on Visualization and Computer Graphics*, 14(6), pp.1221-1228.

Wellcome Library. (2019a) *London's Pulse: Medical Officer of Health reports 1848-1972*. Available from: <http://wellcomelibrary.org/moh/> [Accessed 31st March 2019].

Wellcome Library. (2019b) *London's Pulse: Medical Officer of Health reports 1848-1972, About the reports*. Available from: <https://wellcomelibrary.org/moh/about-the-reports/about-the-medical-officer-of-health-reports/> [Accessed 10th April 2019].

Wellcome Library. (2019c) *London's Pulse: Medical Officer of Health reports 1848-1972, Using the report data*. Available from: <https://wellcomelibrary.org/moh/about-the-reports/using-the-report-data/> [Accessed 10th April 2019].

Wellcome Library. (2019d) *London's Pulse: Medical Officer of Health reports 1848-1972, The changing face of London*. Available from: <https://wellcomelibrary.org/moh/about-the-reports/the-changing-face-of-london/> [Accessed 10th April 2019].

Wellcome Library. (2016) *Meet the digital data dabblers*. Available from: <http://blog.wellcomelibrary.org/2016/11/meet-the-digital-data-dabblers/> [Accessed 31st March 2019].

Westland, S., Laycock, K., Cheung, V., Henry, P. & Mahyar, F. (2007) Colour harmony. *Journal of the International Colour Association (JAIC)*, published as *Colour: Design and Creativity by The Society of Dyers and Colourists (UK)*, 1, pp.1-15. Available from: https://aic-color.org/resources/Documents/jaic_v1_review.pdf [Accessed 1st April 2019].

Whitelaw, M. (2015a) Generous Interfaces for Digital Cultural Collections. *Digital Humanities Quarterly*, 9(1), pp.1-16.

Whitelaw, M. (2015b) Representing Digital Collections. In: Carlin, D. & Vaughan, L. (eds.) *Performing Digital: Multiple Perspectives on a Living Archive*. Ashgate Publishing Limited, Farnham, pp.77-96.

Whitelaw, M. & Hinchcliffe, G. (2013) Discover the Queenslander. Available from: <https://www.slq.qld.gov.au/discover/exhibitions/past-exhibitions/discover-queenslander#/grid> [Accessed 17th May 2019].

Wikidata. (2019) *date depicted*. Available from: <https://www.wikidata.org/wiki/Property:P2913> [Accessed 6th May 2019].

Williford, J. (2011) Graphing Culture, *Humanities Magazine*, March/April 2011. Available from: <https://www.neh.gov/humanities/2011/marchapril/feature/graphing-culture> [Accessed 31st March 2019].

Windhager, F., Federico, P., Mayr, E., Schreder, G. & Smuc, M. (2016) A Review of Information Visualization Approaches and Interfaces to Digital Cultural Heritage Collections. In: Aigner, W., Schmiedl, G., Blumenstein, K. & Zeppelzauer, M. (eds.) *Proceedings of the 9th Forum Media Technology 2016, St. Pölten, Austria, 24th November 2016*. pp.74-81.

Windhager, F., Federico, P., Schreder, G., Glinka, K., Dörk, M., Miksch, S. & Mayr, E. (2018) Visualization of Cultural Heritage Collection Data: State of the Art and Future Challenges. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, April 2018.

Winters, J. (2017a) Tackling complexity in humanities big data: From parliamentary proceedings to the archived web. *Studies in Variation, Contacts and Change in English*, 19 *Big and Rich Data in*

English Corpus Linguistics: Methods and Explorations. Helsinki, VARIENG. Available from: <http://www.helsinki.fi/varieng/series/volumes/19/winters/> [Accessed 7th July 2019].

Winters, J. (2017b) Web archives for humanities research - some reflections. In: Brügger, N & Schroeder, R. (eds.) *The Web as History: Using Web Archives to Understand the Past and Present*. London: UCL Press. pp.238-248.

Wrubel, L. (2018) *Library of Congress Colors*. Available from: <https://loc-colors.glitch.me/> [Accessed 1st April 2019].

Yale University Library Digital Humanities Lab. (2017) Neural Neighbors. Available from: <http://dhlab.yale.edu/projects/neural-neighbors/> [Accessed 10th May 2019].

Zorich, D. M. (2013) Digital art history: a community assessment. *Visual Resources*, 29(1-2), pp.14-21.

List of Figures

Figure 1. Physical cultural collections: (top) National Museum of Natural History, Smithsonian, Entomology Collection, photo Chip Clark (public domain); (middle) Brooklyn Museum open storage, photo Mark Schlemmer (CC BY 2.0); (bottom) historic newspapers at the British Library, photo Luke McKernan (CC BY-SA 2.0).....	9
Figure 2. Example digital record: digital reproduction (left) and beginning of the associated metadata (right), in JSON format, from Cooper Hewitt collection. Fragment, ca. 1700. Public domain image. Data © Cooper Hewitt, Smithsonian Design Museum.....	10
Figure 3. Cooper Hewitt museum collection website. Sèvres dinner plates set. Screenshot 16/08/19. © Cooper Hewitt, Smithsonian Design Museum.....	11
Figure 4. MOMA. 2012. ‘Inventing Abstraction’. Content produced and provided by Second Story..	13
Figure 5. Mitchell Whitelaw. 2009. ‘The Visible Archive’.....	13
Figure 6. (above) Giovanna Ceserani, Nicole Coleman, Mark Braude, Giorgio Caviglia, Ethan Jewett. 2014. Palladio.....	13
Figure 7. (left) Nadav Hochman and Lev Manovich. 2014. Visualization of the Thomas Walther Collection.....	13
Figure 8. Franco Moretti. 2005. ‘The rise of the novel, 18th to 20th century’. From ‘Graphs, Maps, Trees: Abstract Models for Literary History’ (p.6).....	15
Figure 9. Science Museum Group Collection website. 2019. The Babbage Papers archive.....	17
Figure 10. Flavio Gortana, Franziska von Tenspolde and Daniela Guhlmann. 2017. ‘Coins – a journey through a rich cultural collection’, showing data mapped by date against historical period.....	20
Figure 11. The Royal Photographic Society Collection. Photograph. Date: 1860s (photographed) mid 20th century (printed). © Victoria and Albert Museum, London.....	21
Figure 12. Mariner’s Astrolabe (replica) (USA), original: 1602; replica: 1963. © Cooper Hewitt, Smithsonian Design Museum.....	21
Figure 13. (right) Bebyggelseregistret (Swedish Building Register). Rölunda church. Record has 46 dates attached for episodes in the building’s history. Screenshot shows beginning of this date information in the record.....	21
Figure 14. (left) Nordic Museum. Print. Production date ‘1828-1835’, and ‘tidsanda’ (time spirit): ‘1700-talets första hälft’ (first half of the 18th Century). (CC BY-NC-ND 4.0).....	22
Figure 15. (right) Jewellery design. ‘production date’: ‘1908-1917’, ‘style’: ‘20th century’. © Victoria and Albert Museum, London.....	22
Figure 16. (left) Girolamo Andrea Martignoni. 1718. ‘Imago Romani Imperii’. Detail. Collection: Institut Cartogràfic i Geològic de Catalunya. (Public domain).....	23
Figure 17. (right) Joseph Priestley. 1765. ‘Chart of Biography’. Detail. Collection: Chetham’s Library, Manchester. Photo: Stephen Boyd Davis.....	23
Figure 18. Mucha Foundation. 2012. Timeline.....	25
Figure 19. ProPublica. 2014. ‘Bud’s Story, from the Records’ (McCloskey & Wei, 2014).....	25

Figure 20. Metropolitan Museum of Art. 2000. ‘Heilbrunn Timeline of Art History’. Screenshot 16/08/19.	26
Figure 21. British Museum & Google Cultural Institute. 2015. ‘The Museum of the World’.	26
Figure 22. Cleveland Museum of Art. 2014. Gallery One: ‘Love and Lust’ timeline.	26
Figure 23. ‘Darwin’s letters: a timeline’. 2017. (Darwin Correspondence Project, 2017b). Created by Surface Impression Ltd.	27
Figure 24. ‘Pilgrimage’ timeline, Ashmolean: Eastern Art Online (2013). © Ashmolean Museum, University of Oxford.	27
Figure 25. Christian Bernhardt, Gabriel Credico, Christopher Pietsch, Marian Dörk. 2014. ‘Deutsche Digitale Bibliothek visualized’.	27
Figure 26. ‘Tate Explorer’ (Hinchcliffe, 2016).	27
Figure 27. Lev Manovich and Jeremy Douglass. 2009. Covers of every Time magazine issue published from 1923 to 2009, arranged in order of publication.	28
Figure 28. Brian Foo. 2016. ‘NYPL Public Domain Visualization’.	28
Figure 29. City of Amsterdam, Monuments and Archaeology, Belowthesurface.amsterdam. 2018. ...	28
Figure 30. Pietsch. 2017. ‘VIKUS Viewer’ visualising image data from a collection of drawings by Frederick William IV of Prussia (1795–1861) (Glinka, Pietsch & Dörk, 2017).	29
Figure 31. Kilian Krug, Markus Lerner and Severin Wucher. 2018. ‘Forgotten Heritage’: a project led by Arton Foundation and realised together with Luca School of Art, UzF Office for Photography and Kumu Art Museum. Supported by Creative Europe. © 2020 Plural, The Visual Archive: visualarchive.de www.forgottenheritage.eu	29
Figure 32. Florian Kräutli. 2016. ‘Timeline Tool’ visualising the Tate’s collection. (Top) colouring orange artworks created by Eric Gill and (bottom) the Eric Gill artworks separated from the remaining dataset (Kräutli, 2016b: p.178).	29
Figure 33. Stefan Jänicke. 2018. Timages. Visualising portraits of musicians from the MusiX database. (Jänicke, 2018b:fig.5).	30
Figure 34. Stefan Jänicke. 2018. Time-based Impact Mosaics. 211 paintings with a horizontal golden ratio format. (Jänicke, 2018a:fig.3a).	30
Figure 35. Flavio Gortana, Franziska von Tenspolde and Daniela Guhlmann. 2018. ‘Coins – a journey through a rich cultural collection’. Data mapped by earliest date and grouped by material.	30
Figure 36. Manovich, Douglass & Huber. 2010. Time magazines cover visualisation. Detail. X-axis: publication date. Y-axis: saturation (softwarestudies.com, 2009).	31
Figure 37. Manovich, Douglass & Huber. 2010. Time magazines cover visualisation. X-axis: publication date. Y-axis: brightness (visualising data, 2011).	31
Figure 38. Knoll Archive Timeline (Gist, 2018). Lines and opacity showing items related to Lilly Reich.	32
Figure 39. Knoll Archive Timeline (Gist, 2018). By selecting Lilly Reich, only related items are visualised and the date marker arrangement redraws.	32
Figure 40. DigitaltMuseum. 2017 Timeline view. (DigitaltMuseum, 2019c).	33

Figure 41. Tate ‘Timeline of Modern Art’. 2015. Images courtesy of Framestore, Tate.....	33
Figure 42. ‘Steptext’ visualisation of MOH reports data: ‘nurse’.....	43
Figure 43. Example cover and internal pages from Medical Officer of Health reports, from the Wellcome Library website (2019). Left to right: Paddington 1862, cover; Chelsea 1954, p.37; St Martin-in-the-Fields 1856, p.16.	44
Figure 44. Line graph showing number of MOH reports per year.....	44
Figure 45. Prototype visualisation of MOH reports data: ‘typhoid carrier’.....	46
Figure 46. Prototype visualisation of MOH reports data: ‘typhoid carrier’.....	46
Figure 47. Prototype visualisation of MOH reports data: ‘asparagus’.....	46
Figure 48. Detail from visualisation of MOH reports data: ‘nurse’, showing tooltip on hover.	47
Figure 49. Detail from visualisation of MOH reports data: ‘drinking water’ showing bracket, tooltip and underline indicating multiple snippets are from the same document.....	48
Figure 50. Detail from visualisation of MOH reports data: ‘typhoid’ showing snippets corresponding to table entries.	48
Figure 51. Detail from visualisation of MOH reports data: ‘opiates’.	49
Figure 52. Visualisation of MOH reports data: ‘blitz’.	49
Figure 53. Detail from visualisation of MOH reports data: ‘blitz’.	49
Figure 54. Visualisation of MOH reports data: ‘heroin’.	50
Figure 55. Detail from visualisation of MOH reports data: ‘heroin’.	50
Figure 56. Detail from visualisation of MOH reports data: ‘laudanum’, 1857-68.	51
Figure 57. Detail from visualisation of MOH reports data: ‘addiction’, 1950-62.....	51
Figure 58. Detail from visualisation of MOH reports data: ‘addiction’, 1969.	51
Figure 59. (left) Detail from visualisation of MOH reports data: ‘typhoid’, showing that when there are many hits the snippets become a list.....	52
Figure 60. (right) Zoomed out view of visualisation of MOH reports data: ‘typhoid’.	52
Figure 61. Prototype visualisation of MOH reports data: ‘ice cream’ & ‘poison’.	55
Figure 62. Mockup visualisation of MOH reports data including page photo crops: ‘inoculate’. ...	56
Figure 63. ‘Tag-timeline’ visualisation of Cooper Hewitt collection data: ‘Coffee and tea drinking’ ...	58
Figure 64. Poster, Miniwatt/Philips Radio; Client: Philips; lithograph on paper. Tags: <i>graphic design, electronics, commercial poster</i> . (20/03/19). © Cooper Hewitt, Smithsonian Design Museum.....	59
Figure 65. Pair Of Slippers (USA), 1889; soldered tin. Tags: <i>interior, decoration, personal, display, gifts, women, commemorative, ribbons, figurative, walking, decorative, celebration, metal, bows, pair, fashion accessory, fashion, metalwork, shoes, footwear, buckles</i> . (19/03/19). © Cooper Hewitt, Smithsonian Design Museum.	59
Figure 66. Prototype visualisation items tagged ‘triangles’ in the Cooper Hewitt Smithsonian Design Museum collection horizontally by date.....	61

Figure 67. Visualisation annotated to show possible alternative positions (blue rectangles) for an example object, with date span shown in yellow.....	62
Figure 68. Visualisation of Cooper Hewitt collection data: ‘Black and white’.....	62
Figure 69. Visualisation of Cooper Hewitt collection data, time running vertically.....	63
Figure 70. Visualisation of Cooper Hewitt collection data, detail: selection box.....	64
Figure 71. Visualisation of Cooper Hewitt collection data: ‘Timekeeping’, detail: selected object with tag information.....	64
Figure 72. Visualisation of Cooper Hewitt collection data: ‘Flowers’ with 300 images.....	64
Figure 73. Visualisation of Cooper Hewitt collection data: ‘Night’ with 9 images (below 10 as not all objects in the data had an image). Showing images scaled up.....	65
Figure 74. Visualisation of Cooper Hewitt collection data: ‘Coffee and tea drinking’.....	66
Figure 75. Visualisation of Cooper Hewitt collection data: ‘Chairs’.....	66
Figure 76. Visualisation of Cooper Hewitt collection data: ‘Checkerboard’.....	66
Figure 77. Visualisation of Cooper Hewitt collection data: ‘Bent’.....	67
Figure 78. Visualisation of Cooper Hewitt collection data: ‘Earrings’.....	67
Figure 79. Visualisation of Cooper Hewitt collection data: ‘Personal Environmental Control’.....	67
Figure 80. Visualisation of Cooper Hewitt collection data: ‘Art nouveau’.....	68
Figure 81. Visualisation of Cooper Hewitt collection data: ‘Travel posters’.....	68
Figure 82. Visualisation of Cooper Hewitt collection data: ‘Skyscrapers’.....	69
Figure 83. Visualisation of Cooper Hewitt collection data: ‘Space’.....	69
Figure 84. Visualisation of Cooper Hewitt collection data: ‘Water’.....	70
Figure 85. Visualisation of Cooper Hewitt collection data: ‘Handkerchiefs’.....	71
Figure 86. Visualisation of Cooper Hewitt collection data: ‘Timekeeping’.....	71
Figure 87. Visualisation of Cooper Hewitt collection data: ‘Portable’.....	71
Figure 88. User testing/interviewing museum visitors in the galleries: video recording stills.....	72
Figure 89. Visualisation of Cooper Hewitt collection data: ‘Lamp’, annotation in yellow showing similar minimal lamp designs.....	73
Figure 90. Visualisation of Cooper Hewitt collection data: ‘Chimeras’.....	76
Figure 91. Olivia Vane. 2018. ‘Dive into Color’. Photo: Caroline Koh Smith.....	78
Figure 92. Colour harmony examples.....	80
Figure 93. Light and shadow on 3D objects can introduce multiple, illusory colours when pixel values are used to extract colours for objects. Object IDs (left to right): 35460799; 18806261. © Cooper Hewitt, Smithsonian Design Museum.....	80
Figure 94. Tiles photographed with stickers. Tile wall facing (Netherlands), ca.1725. Object IDs (left to right): 18726001; 18725273; 18725393. © Cooper Hewitt, Smithsonian Design Museum.....	81

Figure 95. Buttons photographed with yellow paint. Object IDs (left to right): 18309573; 18309683; 18309545. © Cooper Hewitt, Smithsonian Design Museum.	81
Figure 96. Lace photographed against a dark background. Object IDs (left to right): 18569397; 18569395; 18318023. © Cooper Hewitt, Smithsonian Design Museum.	81
Figure 97. Colour search on the Cooper Hewitt website (Cooper Hewitt, 2018). Screenshot taken Aug 2018. © Cooper Hewitt, Smithsonian Design Museum.	82
Figure 98. Visualisation of Cooper Hewitt collection data: ‘Orangered’.	82
Figure 99. Visualisation of Cooper Hewitt collection data: ‘Steelblue’.	82
Figure 100. Visualisation of Cooper Hewitt collection data: ‘Olivedrab’.	83
Figure 101. Snapped-to colours in the Cooper Hewitt collection. CSS4 (left) and Crayola (right) palettes mapped by hue (in HSL). Angle = hue, radius = lightness.	83
Figure 102. CSS4 palette colours (left) mapped round a wheel (right). Hue (HSL) = angle, ordered by lightness.	84
Figure 103. (left) Colours mapped by hue (HSL) resulting in an RGB colour wheel. Angle = hue, radius = lightness.	84
Figure 104. (right) Colours mapped by hue (HSL), adjusted to more closely resemble a pigment (RYB) colour wheel. Angle = hue, radius = lightness.	84
Figure 105. Visualisation of Cooper Hewitt collection data: Purple & olive.	85
Figure 106. Visualisation of Cooper Hewitt collection data: Orangered & cyan/blue.	85
Figure 107. Different colour harmonies tried in prototyping.	85
Figure 108. (upper) Object with colour palette, (lower) palette mapped to a wheel, easy to see complementary harmony.	86
Figure 109. Colour wheel graphic examples.	86
Figure 110. Prototype visualisation exhibited at Royal College of Art’s January 2018 ‘Work in Progress’. Photo: Olivia Vane.	87
Figure 111. ‘Sample Book (France), 1938; Made by Textile Argus’. Object ID 18431057. © Cooper Hewitt, Smithsonian Design Museum.	87
Figure 112. Olivia Vane. 2018. ‘Dive into Color’ installed at Cooper Hewitt. Photo: Caroline Koh Smith.	88
Figure 113. Olivia Vane. 2018. ‘Dive into Color’ installed at Cooper Hewitt. Photo: Caroline Koh Smith.	89
Figure 114. ‘Color Wheel, 1936–37; Hilaire Hiler’. (Cooper Hewitt, 2019e). © Cooper Hewitt, Smithsonian Design Museum.	89
Figure 115. Detail from ‘Dive into Color’ showing Colour Wheel picker, inspired by Hiler’s design. ...	89
Figure 116. Visualisation of Cooper Hewitt collection data: Spectrum.	90
Figure 117. Visualisation of Cooper Hewitt collection data: Monochromatic, purple.	90
Figure 118. Visualisation of Cooper Hewitt collection data: Monochromatic, purple.	91
Figure 119. Visualisation of Cooper Hewitt collection data: Complementary, purple & green.	91

Figure 120. Visualisation of Cooper Hewitt collection data: Monochromatic, red.....	92
Figure 121. ‘Frieze (USA), 1890–1910; Manufactured by Hobbs, Benton & Heath’. Object ID: 18500019. © Cooper Hewitt, Smithsonian Design Museum.....	92
Figure 122. ‘Sidewall, Anemone, 1960–66; Designed by Phoebe Hyde’. Object ID: 18459631. © Cooper Hewitt, Smithsonian Design Museum.	92
Figure 123. Visualisation of Cooper Hewitt collection data: Monochromatic, blue. Annotated to show late 17th Century/early 18th Century Dutch blue & white ceramics.....	93
Figure 124. ‘Plate (Netherlands), 1675–1725; tin-glazed earthenware’. Object ID: 18621479. © Cooper Hewitt, Smithsonian Design Museum.	93
Figure 125. ‘Plaque (Netherlands), 1675–1725; tin-glazed earthenware’. Object ID: 18621427. © Cooper Hewitt, Smithsonian Design Museum.	93
Figure 126. ‘Obelisk (Netherlands), ca. 1700–25; tin-glazed earthenware’. Object ID: 18621483 © Cooper Hewitt, Smithsonian Design Museum..	93
Figure 127. Visualisation of Cooper Hewitt collection data: Monochromatic, red. Annotated to show French red and white textiles in the late 18 th / early 19 th Century.....	94
Figure 128. ‘Textile (France), late 18th Century; cotton’. Object ID: 18438769. © Cooper Hewitt, Smithsonian Design Museum.....	94
Figure 129. ‘Textile (France), ca. 1850; cotton’. Object ID: 18671007. © Cooper Hewitt, Smithsonian Design Museum.....	94
Figure 130. ‘Textile (France), 18th Century; cotton’. Object ID: 18482057. © Cooper Hewitt, Smithsonian Design Museum.....	94
Figure 131. Visualisation of Cooper Hewitt collection data: Monochromatic, blue. Annotated to show English blue & white, Wedgwood-type, late 18th Century ceramic buttons/medallions.....	95
Figure 132. ‘Medallion (England), late 18th Century; stoneware’. Object ID: 68730605. © Cooper Hewitt, Smithsonian Design Museum.....	95
Figure 133. ‘Medallion (England), late 18th Century; stoneware’. Object ID: 18711563. © Cooper Hewitt, Smithsonian Design Museum.....	95
Figure 134. ‘Button (England), late 18th Century; stoneware’. Object ID: 18310009. © Cooper Hewitt, Smithsonian Design Museum.....	95
Figure 135. Paper title on board reads: ‘Buttons and Mounts, Wedgwood type; England and Belgium, 18th and 19th Centuries’. Object ID: 69112317. © Cooper Hewitt, Smithsonian Design Museum.	95
Figure 136. Visualisation of Cooper Hewitt collection data: Complementary, blue & yellow.	96
Figure 137. Scarf Sample (France), mid-19th Century. (Cooper Hewitt, 2019h). © Cooper Hewitt, Smithsonian Design Museum.....	97
Figure 138. Samuel Bourne. 1866. ‘Evening in the Mountains’. Museum number: RPS.347:57-2017. © Victoria and Albert Museum, London.....	100
Figure 139. George Washington Wilson. Before 1891. ‘Rotten Row’. Museum number: RPS.1925-2018. © Victoria and Albert Museum, London.....	100

Figure 140. Camille-Léon-Louis Silvy. ca. 1860s. 'Woman in a dress'. Museum number: RPS.1728-2018. © Victoria and Albert Museum, London.....	100
Figure 141. Lewis Carroll. c. 1876. Xie (Alexandra) Kitchin. Museum number: RPS.2246-2017. © Victoria and Albert Museum, London.....	100
Figure 142. Paul Martin. 1970-2000. 'Trafalgar Square at Dusk'. Museum number: RPS.5176-2018. © Victoria and Albert Museum, London.....	100
Figure 143. Glass lantern slide. Museum number: RPS.1876-2018. © Victoria and Albert Museum, London.....	100
Figure 144. RPS data plotted horizontally by date, and at a random vertical position.....	101
Figure 145. RPS data plotted horizontally by date, and at a random vertical position.....	102
Figure 146. RPS data plotted horizontally by date, and vertically by photographic technique (techniques ordered from most to least items).....	102
Figure 147. Detail from RPS data plotted by time and technique: different techniques have different tone and image content.	103
Figure 148. Detail from RPS data plotted by time and technique: Julia Margaret Cameron carbon process photographs.	103
Figure 149. Annotated detail from RPS data plotted by time and technique showing anomalous data points.....	104
Figure 150. Detail from RPS data plotted by date and photographic technique: daguerreotypes. ..	104
Figure 151. RPS data plotted horizontally by date and vertically by 'category', excluding the categories 'photography' and 'RPS'.	105
Figure 152. Detail from RPS data plotted by date and category: 'portraits'.	106
Figure 153. Annotated detail from RPS data plotted by date and 'category'.	106
Figure 154. Detail from RPS data visualised by date and technique: lots of carte de visites.	107
Figure 155. Reverse image search results from the RPS collection. Julia Margaret Cameron portraits of women.	108
Figure 156. Reverse image search results from the RPS collection. Julia Margaret Cameron portraits.....	108
Figure 157. Reverse image search results from the RPS collection. Magic lantern slides from Mount Everest expedition of 1921 and Mount Everest expedition of 1922.....	108
Figure 158. Reverse image search results from the RPS collection. Photographs from The Royal Visit to Gwalior album, 1905.....	108
Figure 159. Reverse image search results from the RPS collection. Salted paper prints by William Henry Fox Talbot, all pages from the same album 1844.....	109
Figure 160. Reverse image search results from the RPS collection. Stereoscope daguerreotypes of displays at the Crystal Palace, Sydenham c.1855.....	109
Figure 161. Reverse image search results from the RPS collection. Cartes de visite: full-length, standing figures.....	109

Figure 162. Reverse image search results from the RPS collection. Cartes de visite: (mostly) sitting men.....	109
Figure 163. Reverse image search results from the RPS collection. Cartes de visite: head and shoulders portraits.	109
Figure 164. Reverse image search results from the RPS collection. Photographs of people mounted in a circle/oval, including [middle] The Head of John the Baptist on a Charger by Oscar Gustav Rejlander 1857.....	109
Figure 165. Reverse image search results from the RPS collection. Landscapes with trees: more cartes de visite, [middle] Agnes Warburg’s coloured Avenue of Trees, c.1920, and [right of middle] a lantern slide from Mount Everest expedition of 1921.	110
Figure 166. Reverse image search results from the RPS collection. Photographs of people, some in pairs, inside or by buildings.	110
Figure 167. (left) Reverse image search results from the RPS collection. Landscapes, including photographs [top, left, right of left] by James Ravilious 1970, and [right] Olives and almonds by Agnes Warburg 1930.	110
Figure 168. Reverse image search results from the RPS collection. Mixed images, [top] Goutte de rosée sur une feuille de capucine by Guyula Hálás Brassai, c. 1935 and [right] Lace by William Henry Fox Talbot, early 1840s.....	110
Figure 169. Reverse image search results from the RPS collection.	110
Figure 170. ‘t-SNE’ plot of the RPS data.....	111
Figure 171. Unknown photographer, St. Paul’s Cathedral, daguerreotype, mid-19th Century. Museum number: RPS.518-2017. © Victoria and Albert Museum, London.	112
Figure 172. Unknown photographer, Portrait of a Man and Boy, daguerreotype, mid-19th Century. Museum number: RPS.409-2017. © Victoria and Albert Museum, London.	112
Figure 173. Annotated ‘t-sne’ plot of the RPS data.....	112
Figure 174. Mockup of tiered timeline visualisation.	113
Figure 175. Prototype visualisation of RPS collection data.....	114
Figure 176. Prototype visualisation of RPS collection data.....	114
Figure 177. Prototype visualisation of RPS collection data.....	115
Figure 178. Prototype visualisation of RPS collection data.	115
Figure 179. Prototype visualisation of RPS collection data.....	116
Figure 180. Prototype visualisation of RPS collection data.....	116
Figure 181. Prototype visualisation of RPS collection data.....	117
Figure 182. Prototype visualisation of RPS collection data.....	117
Figure 183. Prototype visualisation of RPS collection data.	118
Figure 184. (left) Nordic Museum. Man’s portrait. 1620-1660 (Estimate). Identifier: NM.0060722. (CC BY-NC-ND).....	119

Figure 185. (right) Nordic Museum portrait. Woman's portrait. Production 1830s. Identifier: NM.0122896. (CC BY-NC-ND).	119
Figure 186. (left) Nordic Museum. Miniature portrait. Production 1600-1659. Identifier: NM.0300957. (CC BY-NC-ND).	119
Figure 187. (right) Nordic Museum portrait. Silhouette portrait. Production 1770-1820 (estimate). Identifier: NM.0046024. (CC BY-NC-ND).	119
Figure 188. 785 portraits in the Nordic Museum collection, ordered by date and chunked vertically into centuries—ordered earliest to latest by (beginning of date span for) production date.	120
Figure 189. Secondary timelines showcasing items from around 1610, and around 1800.	121
Figure 190. Bottom secondary timeline showing portraits from 1604-1887.	122
Figure 191. Additional secondary timelines.	122
Figure 192. Main timeline, showing when in time there are/are not items by the strength and thickness of the line.	123
Figure 193. Annotation on secondary timeline. Clicking the total, '180', generates a tile view below of all 180 portraits for this time span. Clicking the shuffle icon will display a different 7 images.	123
Figure 194. Tile view below timeline.	123
Figure 195. Lightbox viewer.	124
Figure 196. Example silhouettes in 'Faces of Sweden', all dated: estimate 1770-__ __ __ .	125
Figure 197. Pair of portraits in the Nordic Museum collection. (left) Identifier: NM.0236954. (CC BY-NC-ND) (right) Identifier: NM.0236955. (CC BY-NC-ND).	125
Figure 198. Trio of paintings including 2 portraits catalogued in 1 record. Identifier: NM.0030050A-C. (CC BY-NC-ND).	126
Figure 199. Lithographs in the Nordic Museum collection, 1850-60, depicting historical figures. (Example record identifier: NM.0269513). (All CC BY-NC-ND).	126
Figure 200. (left) Nordic Museum. Portrait miniature, 1850-1897, showing actor Edvard Mauritz Swartz in costume as Hamlet. Identifier: NM.0094514. (CC BY-NC-ND).	126
Figure 201. (right) Photograph, 1983, in frame, 1890-1920. Identifier: NM.0323027A-B. (CC BY-NC-ND).	126
Figure 202. Detail from visualisation of Nordic Museum data.	127
Figure 203. (left) Production date 1700-1799. Identifier: NM.0010732. (CC BY-NC-ND).	127
Figure 204. (right) Production date 1800-1899. Identifier: NM.0082269. (CC BY-NC-ND).	127
Figure 205. 'Faces of Sweden' visualisation.	128
Figure 206. Visualisation of Cooper Hewitt collection data: 'Water'.	131
Figure 207. Tiles photographed with stickers. Extracting colour data from these photographs includes the sticker's colour. The colour data is 'correct' in a simple sense, though a human would treat it as irrelevant. Object IDs (left to right): 18726001; 18725273; 18725393. © Cooper Hewitt, Smithsonian Design Museum.	133

Figure 208. Prototype visualisation of RPS collection data. Using computationally-extracted visual similarity data connects an image of blotchy paper (top)—a photograph’s back—with photographs of grass (bottom middle); this could be considered naive, though the images are texturally similar... 134

Figure 209. Visualisation of Cooper Hewitt collection data: Monochromatic, blue.....138