2015

Pathways through online museum collections: designing serendipitous user experiences using formal concept analysis

Timothy Daniel Wray
University of Wollongong

Pathways through Online Museum Collections: Designing Serendipitous User Experiences using Formal Concept Analysis

Timothy Daniel Wray, B InfoTech (Hons)

Supervisor:
Professor P. Eklund

This thesis is presented as part of the requirements for the conferral of the degree:

Doctor of Philosophy

The University of Wollongong
School of Computing and Information Technology

November 2015

Declaration

I, Timothy Daniel Wray, B InfoTech (Hons), declare that this thesis submitted in partial fulfilment of the requirements for the conferral of the degree Doctor of Philosophy, from the University of Wollongong, is wholly my own work unless otherwise referenced or acknowledged. This document has not been submitted for qualifications at any other academic institution.

Timothy Daniel Wray, B InfoTech (Hons)

November 24, 2015

Abstract

Museums are no longer seen as warehouses of objects, but as institutions of knowledge. A growing number of museums are putting their collections online in an effort to make them accessible, navigable and meaningful to their audiences. The digital medium offers the ability to highlight the narratives and relationships shared by their objects through the creation of exploratory interfaces and data visualisations.

In recognition of this – and inspired by perspectives on information seeking that incorporate elements of serendipity, play and aesthetics – this research explores the concept of the *pathway* as a design theory for visualising, navigating and exploring digitized museum collections. The thesis presents a series of concept designs where users follow, branch and navigate through a narrative of interconnected pathways and discover the rich connections that objects have with one another. These connections are formed through the use of Formal Concept Analysis: a technique that harvests the embodied knowledge within these collections.

Based on a series of case studies and user experience evaluations, the research highlights the importance of showing the context and relationships that surround the objects, the ability to meaningfully engage with high quality content, issues that concern knowledge acquisition and representation, and the role of aesthetics and visual momentum in the design of experiences that promote exploration within digitized museum collections. The design principle of *pathways* as a means to structure content and afford creative exploration is realised through the creation, shaping, evaluation and critique of three concept designs.

Acknowledgments

First and foremost, I would like to thank my supervisor Prof. Peter Eklund for his support, whose door was always open on matters of technical questions, and perhaps more importantly, for his sense of humour and camaraderie. Peter made what would otherwise have been a very tedious journey a fun and enjoyable experience. I would also like to thank Prof. Karlheinz Kautz for his boundless energy and enthusiasm for my research, and for his contribution and solid advice on my manuscript preparation for the ICIS2013 conference.

I extend my gratitude to Timothy Zelinsky and Adam Muhlbauer for their technical work on *A Place for Art*, making it look and feel *just* the way I wanted to, and to Phillippa Webb and Prof. Amanda Lawson for supporting the project and assisting on curatorial matters. I would also like to thank the Centre for Digital Ecosystems for providing travel and funding support for the *A Place for Art* user experience evaluations.

A warm heart-felt thank you extends to my friend Ben Gooden, whom I lived with for over a year, a fellow PhD student and partner in crime. I also extend my gratitude for my friends and family who have tolerated the last few months of my writing in solitude.

Contents

Chapter 1

Introduction

There are a growing number of cultural institutions, such as galleries, libraries, archives and museums that are making their collections available online in an effort to make them more accessible and engaging to their audiences [1]. [1] Online aggregators, such as Europeana [2], the Google Art Project [3] and the Flickr Commons [4] showcase massive aggregations of artworks, photography and material culture from hundreds of institutions. Not only do these sites vastly increase the accessibility of the works through digitisation and online access, but they also provide engaging interfaces that allow for personal exploration and meaning-making to take within their collections, as was reported in the case of the Flickr Commons [2].

Many of these sites also provide access to their collections data, or metadata, so that it can be explored by both humans and machines. [5] Collection metadata can be mined and presented in ways that allow for new use scenarios and meanings to emerge. For instance, Whitelaw et al. [3] build rich and immersive 'generous' interfaces that encourage users to explore cultural collections, and Manovich et al. [4] propose the practice of *cultural analytics* as a set of tools, techniques and visualisations that seek to elicit new meanings from cultural data.

Much of the meaning that can be interpreted from museum collections lies, at least in part, in the way individual objects are shown in relation to other individual objects [5]. Further, the context that an artefact is shown is considered to be just as important – if not more important – than the object itself [6–8]. An understanding of the importance of context stems from examining the role of the modern museum as a knowledge organisation [8–10]. By putting their collections online, museums

[1] A list of some of these online collections can be found at The Museum of Online Museums, available at : `http://www.coudal.com/moom/`

[2] `http://www.europeana.eu/`

[3] `https://www.google.com/culturalinstitute/project/art-project`

[4] `https://www.flickr.com/commons`

[5] A list of museum APIs can be found at :
`http://museum-api.pbworks.com/w/page/21933420/Museum%C2%A0APIs`

can connect with audiences that would be otherwise limited by geographical distance. In addition to providing greater and more equitable access to collections, the digital medium also allows other possibilities: the enrichment of collections via commenting and social media, the sharing of objects and re-appropriation of archival materials [2] and the redistribution of museum knowledge via social tagging [5, 11]. Online collections also have the vast potential to represent the embodied knowledge and rich connections that objects have with one another. It is this concept that inspires and motivates this research: the potential for new design possibilities – driven by emerging perspectives in information seeking – that can empower visitors to visualise knowledge, seek connections, and frame online collections as repositories of knowledge, rather than as warehouses of objects.

1.1 Research Objectives

The aim of this research is two-fold: a) to realise the design principle of a *pathway* as a design metaphor that, in its instantiation, encourages exploration and meaning-making to take place in online collections, and b) to employ the use of Formal Concept Analysis as a means of mining and exploring this landscape of knowledge.

The metaphor of the *pathway* can be defined under two perspectives: as a means of *semantically structuring content*, where linked artworks and artefacts can be represented as a traversable map, graph or structure; and as a means of *affording creative exploration* that embodies a curious and exploratory mindset of information seeking that encourages serendipity, unexpected connections and a learning experience that connects objects, concepts, ideas and people. [12]. The *pathway* metaphor is drawn from user-centered design and is inspired by emerging perspectives on information seeking that incorporate dimensions of aesthetics [13, 14], serendipity [15] and the notion of an "information landscape" [16]. Likewise, the use of Formal Concept Analysis is predicated on its close connections to the philosophical logic of human thought [17] and its ability to construct conceptual spaces and represent knowledge.

In this thesis, a survey of the literature in museology and information seeking derives a set of design principles. These design principles include theories that describe creative and curious information seeking [18], visuo-tactile interaction [14], and the use of pathways as a means to structure content and afford creative exploration [12]. The design principles are explored and realised in the following way:

1. An analysis of the theory of Formal Concept Analysis and how it can be used to create information spaces that allow users to explore and create their own paths within museum collections.

2. The creation of a framework, called *CollectionWeb*, that allows websites and apps to harvest the knowledge mined by Formal Concept Analysis from museum collection metadata.

3. The development of three different interfaces that allow users to explore pathways, each with their own collection.

4. A user experience evaluation based on two of these interfaces that considers visitors' perspective on the *pathway* metaphor as a means of browsing, along with an interpretation on how the interface would facilitate their engagement with the collection.

The three interfaces developed from this research are *The Virtual Museum of the Pacific*, a Web-based browsing interface developed in collaboration with the Australian Museum that showcases objects from the Pacific cultural collections; *A Place for Art*, an iPad app that allows users to explore the University of Wollongong Art Collection via the creation of pathway structures and the *Brooklyn Museum Canvas*, an app that loads data from the Brooklyn Museum's API to create a canvas of pathways structures that can be controlled by a movable and pannable 'camera.' These interfaces are evaluated with respect to the set of design principles. A separate analysis in Chapter 5 also summarises the key findings of the user experience evaluations.

While this thesis takes a strong focus on the user experience of *online* collections, the surveyed literature and case studies also briefly explore how that experience interplays with the physicality of the objects and the museum itself. While such observations contribute to the discussion and analysis of the case studies outlined in this work, a dedicated analysis of the relationship between the online and 'real-world' experience of a museum collection falls beyond the scope of this thesis.

1.2 Methods

The research follows the approaches of design science research [19] and concept-driven interaction design [20]. In design science research, theory is built through design and implementation of working artefacts that seek to solve organisational problems [21]. The design of artefacts can also be derived from existing principles and kernel theories [22]. Concept-driven interaction design follows a similar approach where artefacts are carefully crafted to stand for underlying theoretical ideas.

Design science is well recognised within the information systems and computer science disciplines [23] with several approaches, [19, 24–27], all adhering to the following general structure:

- Identify a problem.
- Propose a design, or implement the design as a working artifact.
- Evaluate the design or artifact using acceptable evaluation methods.
- Generalise and/or theorise the findings.

Design science, particularly within the field of information systems research, generally places great emphasis on solving relevant, well-defined organisational problems [19]. A variation of this approach, however, places more emphasis on exploring new design possibilities without an organisational problem to solve [22]:

- Explore *design challenges* and *derive design principles* from existing knowledge.
- Build a working artifact.
- Evaluate the artifact in an acceptable, analytical manner [21, 28] in line with the design principles.

In design science research, while it is common to employ focus groups and ethnographic research to inform its problem identification phase, this alternative approach uses existing theories and phenomena to guide new design explorations [21]. The approach is similar to *concept-driven interaction design*, where design exploration begins from a theoretical grounding rather than a user-centered approach, and artifacts are deliberately crafted as future-use scenarios that stand as arguments for their underlying theoretical ideas [20]. As such, the need to articulate the theoretical underpinnings and design principles, and how artifacts explore new design spaces in their instantiation, form an important part of this research approach.

The approach recognizes design science as a search process [19] and is suitable for exploring new design ideas that may fall counter to existing and predominant user-centered design approaches. Such approaches, while valuable for identifying problems and incremental improvements in existing design conventions, are criticized for their potential to quash radical or uncontested interface innovations [29]. Similarly, Norman and Verganti [30] assert that while human-centered design is highly suitable to improving the quality and utility of an existing product space, it limits the exploration of radical and divergent design trajectories that accompany meaning or technology change, and that such meaning-driven innovation can be potentially driven through concept-driven interaction design research. In a more divergent perspective that views information technologies are viewed as a cultural medium – and software as a design material [31] – Janet Murray argues that in

order to establish new conventions and advance the composition of new forms and potential use scenarios, it may be necessary to think beyond current conventions in interaction design or circumvent users existing perceptions of interface technologies and their affordances [16, p .6]:

> ... though we have a well-developed design protocol for user needs analysis and user testing of industrial products, users cannot tell us how to resolve problems that require new design strategies. Consumers of digital artifacts cannot often think past the familiar conventions of existing devices and applications, and may even claim to prefer more limited functionality because it is familiar. Industrial design and the social sciences have methodologies for framing design questions by analyzing the needs and expectations of user communities. But if a needs assessment team had asked people what they wanted a computer to do in 1970 they would not have elicited a desire for the personal computer, let alone the iPod or GPS-enabled cell phone.

In the disciplines of information systems research and human-computer interaction, new designs and interfaces are commonly validated via *observation* and *experimentation* with real world users [29], mostly in the form of usability or user acceptance evaluations. However, Hevner et al. [21] also propose descriptive and analytical methods as legitimate forms of evaluation that are well suited for innovative or radical designs. These methods present an evaluation of the artifact via an informed argumentation of an artifact's utility or how it fits within the design space, and in the case of concept-driven interaction design research, how designs are used to augment their underlying theoretical ideas. Descriptive and analytic evaluations are legitimate approaches for evaluating IS artifacts [28] that are also recognized for building IS theory [32] in that they link the designed artifact back to its theoretical underpinnings. The final contribution of this thesis, presented in Chapter 6, is a design theory that describes the pathway metaphor for navigating and exploring online museum collections.

1.3 Thesis Structure

With respect to the above methods and research objectives, the thesis is structured as follows: Chapter 2 presents related work from multiple perspectives: a museology perspective that recognises the role of the museum as a knowledge organisation and an information seeking perspective that introduces constructs of the *information landscape*, *the information flaneur*, *pliability* and the *pathway*, all which inform the design principles described in Section 2.5.

In Chapter 3 the theoretical background of Formal Concept Analysis is introduced as a formal framework and data analysis technique that mathematises the design principle of *pathways as a mechanism to semantically structure content*. It describes how the technique can be used to generate the conceptual structures that allow users to create and explore pathways within museum collections.

Chapter 4 then presents the *CollectionWeb* framework, along with the *Virtual Museum of the Pacific*, *Brooklyn Museum Canvas* and *A Place for Art* case studies, followed by a summary of their analytic evaluations in Section 5.2, linking each artifact back to their design principles. User experience evaluations for the *Virtual Museum of the Pacific* and the *A Place for Art* case studies were also conducted, with their evaluation methodologies each described in their respective case studies. The results from these user experience evaluations further support and validate the analytic evaluations.

Chapter 5 then revisits the design principles outlined in Section 2.5, describing how they were realised through the design, implementation and evaluation of the artefacts. The chapter also presents a summary of key findings from the user experience evaluations. Finally, Chapter 6 concludes the research by summarising key outcomes of the research as a design theory, and provides additional points of discussion, reflection and opportunities for future work.

Chapter 2

Related Work

This chapter presents a review of related work from multiple perspectives. A museological perspective, presented in Section 2.1, describes the knowledge practices of museums, the characteristics of their collections, and the importance of showing the context that surrounds the objects. Section 2.2 introduces virtual museums, online collections and the underlying notion of *connectedness*.

Section 2.3 provides a brief overview of computer supported cooperative work and how *shared feedback* and *shared experiences* can be used as perspectives to understand how two or more people interact with an online or physical exhibition. Section 2.4 then presents a perspective of information seeking that incorporates elements of creativity, curiousity, play serendipity, aesthetics and visual design along the metaphor of *landscapes* and *pathways* as a way of framing our interaction with information spaces.

These perspectives, when combined, inform the design principles that are realised through the creation and evaluation of the interactive visualisations presented in Chapter 4.

2.1 Museums as Institutions of Knowledge

The primacy and importance of the individual artefact is reflected by the key functions of a traditional museum: the collection, study, preservation and interpretation of objects. This view not only risks museums being seen as warehouses of disparate objects but of institutions of power, where classification and access to material history is governed by the museum itself, containing "relics of a dead past, amenable only to the intellectually and aesthetically elite." [8, p.1]. A more progressive view of the museum is one the that posits the institution as a provider to the general public, helping audiences effectively use its information resources for their own quest of knowledge [6, 8, 33, 34]. In this context, knowledge is distinguished from information in that it imparts a learning experience for the visitor, where knowledge organisation activities are manifested through the way objects are classified and contextualised in public exhibitions. A similar movement – driven by political and economic changes since the 1970s – has forced museums to shift their emphasis to visitors rather than their collections [10] – a movement that has also long recognised the discretionary, self-directed role of museum visitors who seek participatory roles in shaping their experiences [35]. Such a movement – one that shifts the museum as information service providers rather than gatekeepers, and one that positions the institution as a repository of knowledge rather than a repository of objects, has been described as the New Museology movement [10, 34, 36, 37] or the *reinvented museum* [9].

Given the museum's role in collecting, creating and shaping knowledge, the *context* of an object has become an increasingly important part of its analysis, interpretation and communication. Context can refer to an object's materials, construction, design, ornamentation, provenance, history, environment, connection to people and human society [6, 34]. This focus towards context reflects a shift from a *classical* worldview, where objects were classed in terms of order, hierarchy and taxonomy, to a modern perspective where objects are analysed in terms of links to other objects, people, social and cultural histories [34]. Integral to the study of the artefact is its typological groupings among similar artefacts: for example, Pearce [6] describes how the process of contextualisation with other objects – such as comparing a newly acquired silver spoon to other spoons of similar material and construction, or a newly discovered portrait that shares similar lines and brush work with other portraits – remains fundamental to the process of dating, provenancing and interpretation. Placing objects in context via the process of grouping, as demonstrated by the case studies in this thesis, has been observed in influencing the selection and curatorial processes for a online anthropological museum, as reported in the *Virtual Museum of the Pacific* case study (Section 4.2.2). The process of contextualisation also augmented the visitor's understanding and interpretation of a collection as reported in

the *A Place for Art* case study (Section 4.3.4), and can also be used to construct a narrative of themes and visual tropes that describe an artist's body of work, as reported in the *Brooklyn Museum Canvas* case study (Section 4.4.3).

Individual objects are selected by curators because they contain information value through their uniqueness, historical appeal or aesthetic value [6]. Based on the totality of knowledge from these objects, the role of the museum is to select, organise and represent their objects, collections, exhibits and educational programs [8]. In recognition of this, Skov [37, p. 39] describes the following characteristics of museum collections:

- Collections are highly heterogeneous and can contain many different types of objects, such as photographs, sculptures, anthropological and archaeological specimens, text, sound and video – each with their own storage processes and curatorial practices.
- Objects are often described using highly subjective terms, where the cataloguing or description of an object often relies on a curator's own judgement. [1]
- All artefacts, including copies or replicas of existing artefacts, are inherently unique. They are often described in relation to, or as influenced by, the collection.
- The human, social and cultural contexts of their artefacts are just as important, or more important, than the artefact itself. [34]

These characteristics present major challenges in the design of virtual museums and online collections given that their purpose is to share or disseminate meaningful contextual knowledge about the objects and the relationships between them. As online collections are often implemented as direct gateways to museum databases, the problem is compounded by the fact that most museum documentation is either written by experts or specialists that reflect their own perspectives rather than visitor perceptions or public interests [38, 39]. However, case studies in social tagging [11, 39] and the use of social media platforms that encourage content sharing and appropriation [2] are promising avenues that can resolve the accessibility and knowledge gaps between museums and their audiences.

[1] The subjective tagging and cataloguing of works from the Brooklyn Museum's collection resulted in knowledge structures that reflected, in part, the subjective and interpretive knowledge of the 'curator', as was reported in the *Brooklyn Museum Canvas* case study. (Section 4.4.3)

2.2 Virtual Museums and Online Collections

The rise of the 'virtual museum' was partly born out of the New Museology movement, realising the potential for Web technologies to enrich the content and context of museum collections and make them attractive and accessible to the wider public [33]. Virtual museums employ a number of technologies to showcase their collections, such as advanced imaging formats, virtual reality exhibitions, haptic-based systems and hand-held devices [36]. According to this research, online virtual museums are usually characterised by: the ability for users to interact with the collection under a multiplicity of contexts; the encouragement of active rather than passive visitor engagement; and the presentation of the works in a primarily visual, rather than textual manner. Similarly, studies that examine best practices in creating user experiences for virtual museums report that visitors expect the following: the ability to access online collections and navigate based on personal interests [40]; information that is accurate and shareable; an expectation that they will learn from the experience; the ability to communicate and share knowledge with other visitors (e.g. via the use of chat rooms and comment posts) and a desire to freely experiment, play and explore [41].

Research efforts have also investigated how semantic enrichment of museum metadata could be used to better the convey the content and context of their objects [42] and provide personalised tours of virtual museums based on user preferences [43, 44]. Other efforts, such as those instigated by the Powerhouse Museum in Sydney, Australia, [11] have investigated how social tagging and the creation of subsequent 'folksonomies' can enhance access to collections and allow user and creator communities to ascribe their own interpretations to the object: an approach that can potentially allow audiences to shape both the content and context of the collection [39] and encourage the sharing and co-creation of knowledge [5]. Virtual museums have also experimented with the use of 3-dimensional displays to visualise physical museum buildings [45], maps, sites and archaeological installations [46] and present abstract data visualisations [47]. The use of data visualisation to create exploratory and browsing environments of online collections is also an active area of research [3, 4, 48–51], and other work has also experimented with the use of semantic clustering to connect objects and generate thematically coherent 'exhibition rooms' within a virtual museum [52].

The idea that computing can be used show context and relationships within collections be traced back to early visions of the "virtual museum" from the 1990s [33]. In one of these visions, Hoptman [53] introduced the notion of *connectedness*: the virtual museum as one that transcends traditional systems of information organization and classification and, in its place, introduces unparalleled levels, perspectives

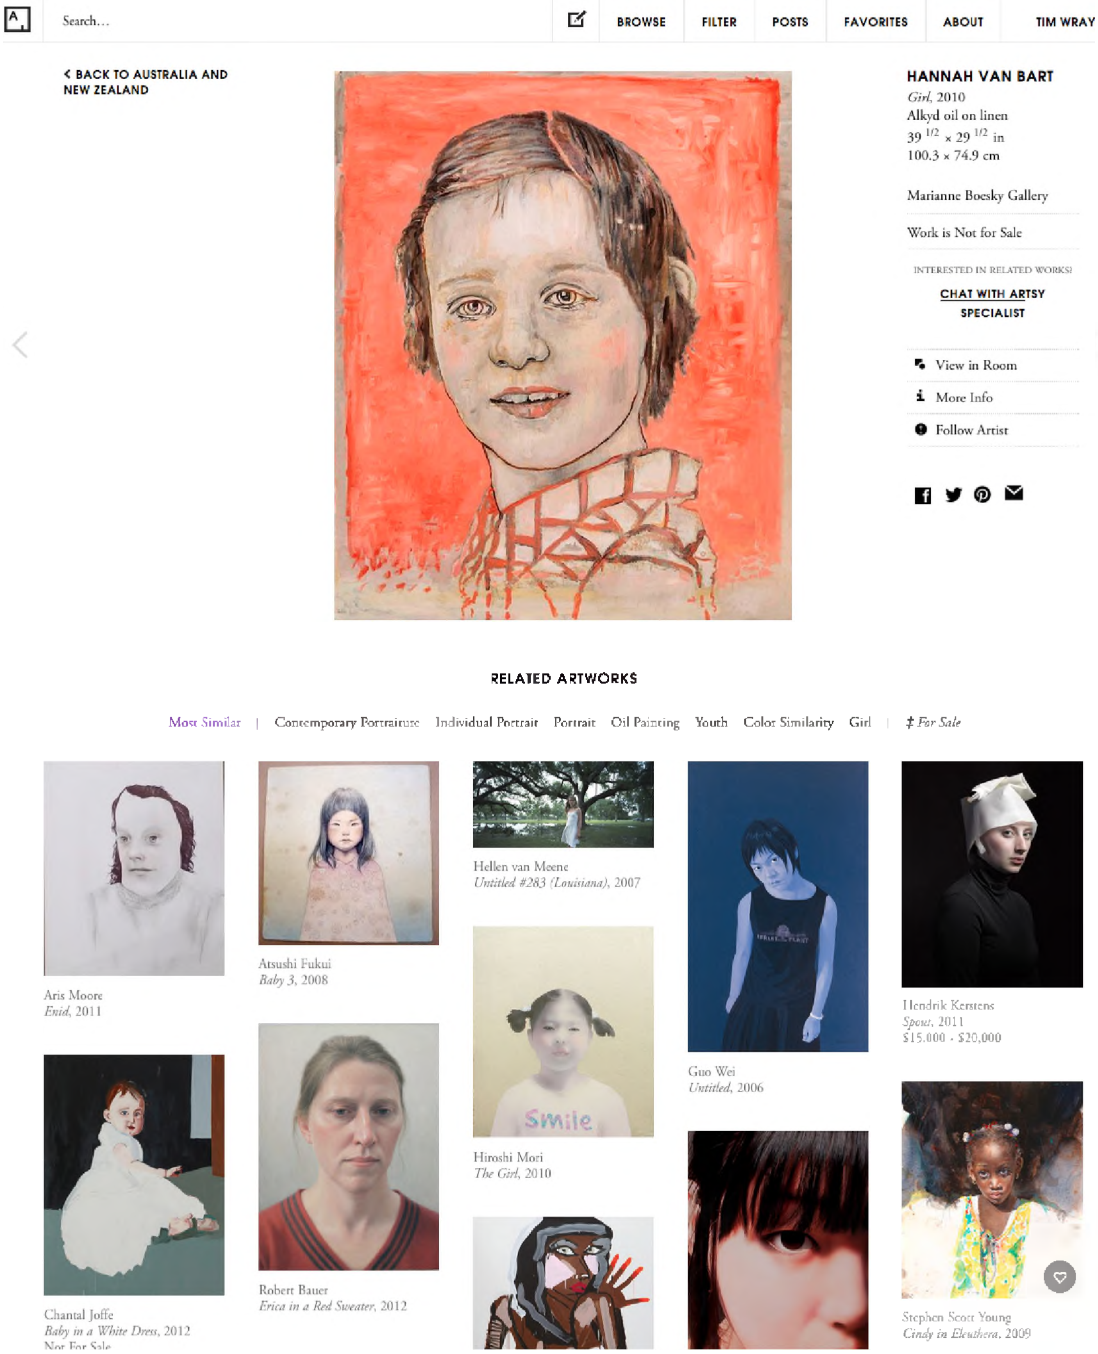

Figure 2.1: Artsy depicts semantically linked object through their ‘Related Artworks’ feature, allowing the user to serendipitously explore the collection and view the objects within a number of contexts.

and dimensions on topics of interest. Schweibenz, on his early vision of the virtual museum [33], describes how *connectedness* can be used to elicit new experiences in providing meaningful context to works of art [33, p. 3]:

> Connectedness is the quality that allows the "virtual museum" to transcend the abilities of the traditional museum in presenting information. This new quality can be realised in different ways, for example, in displaying digital representations of works of art next to comparative works by the same artist, artists who have influenced him or her, or works of the same style or period that are exhibited in museums at various geographic locations or that are otherwise not normally accessible together.

Artsy [2] – arguably a modern implementation of the "virtual museum" – exploits this idea to encourage discovery and exploration within its collection. The site, consisting of over 30,000 artworks from over 500 galleries and museums, shows how they are related to one another across various dimensions, such as their medium, provenance, style and subject matter. It also 'recommends' works based on the similarity of the works that the visitor has already viewed. For example, the work shown in Fig. 2.1, Hannah Van Bart's *Girl*, is shown in relation to other works that also depict girls, works that also depict youth and other portraits with similar stylistic tendencies. Artsy uses an unpublished algorithm to rank and show connections to other works along various facets and themes. The connections highlight opportunities for the visitor to further explore and engage with the collection. Further, it also allows the visitor to examine the context that these works are shown in, even if those contexts are algorithmically curated without human oversight. These implications provide fruitful opportunities for exploring other ways of harvesting knowledge, drawing connections and highlighting equivalencies among objects within museum collections.

Although virtual museums employ use of the digital medium to display and connect objects, they can also be used to blur the line between the digitised objects and the physical spaces that they occupy. For example, the CHIP project [43] allows visitors to rate artworks from the Rijksmuseum's collection. These ratings are used to generate personalised recommendations of other artworks from the collection, which can then be used to create customised, location-based tours within the physical museum space. Conversely, visitors to the The Museum of Old and New Art are provided with a mobile device – "The O" [3] – that tracks their location in the museum and collects their personal ratings of artworks. After they visit the museum, the visitor can then access and share a map of their journey, allowing them

[2] `https://www.artsy.net/`
[3] `https://www.mona.net.au/theo`

to track what parts of the museum they have seen, and what parts they haven't. By recording the details of their trip and creating a shareable 'memory', the visitor experience is extended beyond the interface of the mobile device and the walls of the museum. The use of such devices to provide context and location sensitive content to physical exhibitions has also been investigated for the ability to facilitate active learning and engagement for school-aged children [54], and other studies have investigated alternative input modalities and experience designs for mobile devices in physical museum spaces [55, 56].

Museums have extensively used technology to shape their visitor experiences and augment their exhibitions [57–59]. Some of the experiences mediated by these technologies combine the *connectedness* aspect of the virtual museum with the experiential qualities of a physical museum visit [43]. Although perspectives that discuss the role of technology in mediating a meaningful experience with installation works [60] or its overall role in the service design of a museum exhibition [57, 59] fall beyond the scope of this thesis, the *A Place for Art* case study (Section 4.3) reports on the participants' experiences with the digital representations of the artwork and the connection to their physical presence among the built and natural surrounds of the university campus. Further, the user experience evaluation of *A Place for Art* was conducted in a novel manner: in recognition of its exploratory yet unfamiliar interface, participants were asked to share a single device and interact with it as a pair while discussing ideas on what or how the interface should behave. The idea that participants interact constructively and build a *common ground* – allowing them to discuss details that they otherwise may have ommitted – is a key factor that's well recognised within the literature of computer supported cooperative work (CSCW). [61–63]. Further, CSCW can draw insights into the design and evaluation of technology mediated exhibitions, given that the experience of visiting a museum is often a shared, rather than a solitary experience [57].

2.3 Computer Supported Cooperative Work

Research in the field of computer supported cooperative work (CSCW) investigates the way two or more people collaboratively interact with devices and systems. One such perspective outlines how CSCW can be used to encourage users to share their experiences or discover problems with a complex system in a highly dynamic environment. In one study on collaborative creative writing [64], the authors describe a system that employs a *shared feedback* approach to collaboration where all content production and feedback is made available to all team members who are working on a particular problem. The study outlined three benefits of sharing feedback:

- It reduces the cost to participants in their own information production through passive collection of activities. Users could produce, or withhold, commentary regarding their use of the system based on observation of their paired partner(s).
- It allows users to look for and extract the information that is most relevant for them. This leads to a greater diversity of feedback and opinion among participants.
- It allows the awareness of information so that users can find relevant information through the shared object and gain information from the other participant(s) and the object simultaneously.

Flor [61] notes that a key component in the success of any interaction is that the two participants must share a *common ground*: without this the pair is unable to coordinate their ideas and thoughts and sequence their ideas. Common ground refers to the way two people converse with similar sets of ideas in a shared language. In a more passive and ludic sense, the notion of the *shared experience* is explored in a museum setting where, in one study, paired participants were given mobile devices that allowed them to not only listen to an audio guidebook of the exhibition, but 'eavesdrop' on the guidebook of their companion as well [57]. This allowed further conversations about the work to take place and it also encouraged a social bonding between the two companions. Further, Kjedskov and Stage [65] identified that mobile systems such as these are usually used in highly dynamic contexts, and given that participants are in a museum or otherwise relaxed and casual setting – rather than a task-focused work environment – these insights can be used to shape the experience design and observation protocols of visitors using such interfaces.

2.4 Exploration and Information Seeking

There are two emerging trends that characterise our interaction with online information spaces. The first highlights an awareness of the increasing availability, cultural significance and audience sentiment of open data [66]. In the case of museums, it can be seen in the vast digitization of online collections and the use of online media to share knowledge and provide avenues for exploration and meaning-making [5, 37]. The massive scale and federated availability of online collections creates opportunities for audiences to explore these collected works and find new ways to create meaning. This also encourages personal exploration, knowledge sharing and re-appropriation of collections from cultural institutions to their communities of interest. [37, 67] The meaning-making and immersive qualities of interactions with such collections – to the point of "losing track of time" [68] – calls for new information seeking principles beyond "the utilitarian goals of overcoming information needs, knowledge gaps, uncertainty and problems" [15, p. 1].

The second trend lies in the recognition of pleasure, play [69] and aesthetics [13] in human computer interaction. McCarthy and Wright [70] argue that any experience with technology can be broken down into compositional, spatio-temporal, sensual and emotional dimensions. These qualities are explicitly recognized and observed, for example, in our interaction with public interactive cinema displays [71] and reactive table top surfaces [72], reinforcing the idea of *Homo Ludens* [69] – humans as innately playful. In recognition of this, Gaver [73] advocates that our interactions with technology, or a measure of the worth of those interactions, should go beyond preoccupations with utility or task efficiency and embrace approaches that recognise our playful nature, advocating for design principles that address recreational, diversionary and pleasure-seeking activities [73, p. 1]:

> The idea of a Homo Ludens – humans defined as playful creatures [69] – is an antidote to assumptions that technology should provide clear, efficient solutions to practical problems. From this perspective, we are not just characterized by our thinking or our achievements, but by our playfulness : our curiousity, our love of diversions, our explorations, invention and wonder. An aimless walk in the city centre, a moment of awe, a short-lived obsession, a joke – all are defining and valuable facets for our humanity.

These ideas of serendipity and play are explicitly acknowledged as facets of interactions with information retrieval and browsing systems [74] – a phenomena observed in our interactions with physical library spaces [75], and their digital counterparts [76]. Section 2.4.4 explores these concepts further, along with the concept

of *pliability*: a quality of human computer interaction described by Löwgren [14] that is considered to be malleable, responsive and tightly coupled to the hands and the eyes, and the degree that the interaction facilitates exploration, serendipity, and playfulness in use.

Taking into the account the exploratory mindset of the information seeker, or the expansive information landscapes that they may immerse themselves in, Section 2.4.1 briefly introduces the notion of exploratory search, then Section 2.4.2 explores the concept of the *information landscape* by drawing analogies to the physical world that we live in. Section 2.4.3 then describes the *information flaneur* [15]: a literary personification of a creative, curious and critical information seeker that likens the information seeking experience in a digital space to that of an explorer of a large, modern city. From this, Section 2.4.5 draws on the notion of a *pathway* as a mean of traversing these information landscapes.

2.4.1 Exploratory search

Exploratory search is a form of information seeking where initial information needs are unclear, and that interaction with the information space is rich and serendipitous. Qu and Furnas [77] developed a model-driven approach to exploratory search where literature from human behaviour is used to inform the evaluation process of such systems. Their premise is that, in exploratory search, the participant lacks knowledge of their topic of interest – instead of basic searching, the participant employs techniques of multiple complex searches and information gathering activities. In exploratory search, success is dependent on the ability of the user to undertake their exploratory process and the complex interactions that occur as they make sense of the information that they are viewing while using a 'sense-making model' as a guide for the evaluation process. The objective of the study by Qu and Furnas [77] "was to understand where and how people gather representation ideas during an exploratory search process in a sense-making context, and how familiar keyword-based search systems support this process." Their study drew a number of implications for designing a user experience based on exploratory search:

- *Support for expressing structure needs*: refers to the way an interface conveys structure within its information space, and the way that it supports the user in understanding its structure. For example, a 'table of contents' style view that provides an overview of the main themes represented in a collection, or a concise data-visualisation that represents some of its key relationships.

- *Support for finding useful representation structure in the world*: based on the user's structure and understanding of its information space, an interface should

employ interaction approaches that support exploration and convey knowledge and the way it's linked and represented.

- *Support for task management*: an interface should be able to provide tools and aid efficiency in completing information seeking or information gathering activities over an extended period of time.

2.4.2 Information landscapes

Interaction design conventions, as argued by Murray [16], are shaped and influenced by conventions and human experiences that have existed long before any digital artifact. In doing this, the author recognizes our allure to physical geography and the natural path-finding instincts of our ancestors, and introduces the metaphor of a *landscape* as a useful notion for framing interactions with technology that elicit navigable, open and connected spaces. The author argues that the way we view the world as a landscape is manifest through our use of language: the "information spaces" of the "World Wide Web"; mastering a "domain" or "expanse" of knowledge; viewing "landmark" works, "crossing discipline boundaries". Landscapes are considered to be a fundamental framework that humans use to make sense of the world, both in natural spatial and non-spatial ways. Following from this, Dörk et al. [15] argue that the design of navigable spaces should establish conventions that determine and consistently uphold rules for spatial interaction: these can range from the consistency of a multi-level menu presented on a news website, the use of "breadcrumbs" to show paths traversed, and the visual and semantic distinction between nodes and paths that are shown in many data visualisations.

The idea that information seekers 'orient' themselves in an *information landscape* was a concept posited by O'Day and Jeffries [78] in describing the way library professionals use multiple, interconnected searches to explore topics in an undirected fashion. This closely ties with the notion of exploratory search [79]: the recognition of discovery and learning as a significant part of the search process where information goals may be fuzzy or unknown. The idea of *information landscapes* as a metaphor for interaction and data visualisation have also been described in [48, 80], although Murray's conceptualisation will be used for the remainder of this thesis [16]. In this thesis, *information landscapes* are defined as a form of networked knowledge that allows an information seeker to navigate and orient themselves within semantically linked data, commonly manifested via the use of exploratory interfaces and data visualization [3, 15, 18, 76].

Some of the most readily observed examples of *information landscapes* takes place in the form of the 3D games and in cartographic software. As these depict environments that are attuned with our innate path-finding instincts [16], technology

products such as Google Earth and Google Street View shift the meaning behind mapping and cartography software from primarily functional tools for navigation and way-finding to interfaces that drive more leisurely pursuits of curious exploration. For example, using integrated and expansive 3-dimensional displays of street and aerial photography, the ability to explore foreign cities and landscapes based on photography of buildings, cities and terrain have been known to induce cartographic and emotional experiences in the form of mediated tourism [81]. A similar drive for curious exploration, at least within the realm of cartography, is reflected in a concept design in the form of a coffee table that displays drifting aerial imagery [72] where, for some users, it was used to explore new areas or provide aerial imagery of places that were meaningful to them. Other users, however, in an almost futile interaction with the technology, simply "[let it] wander and see what happens" [72, p. 896].

To further illustrate the concept, the *landscape* analogy can be contrasted with the more restrictive organizing framework of a *container*, a concept also introduced by Murray [16]. Whereas *landscapes* are fluid, expansive and open, *containers* are discrete, restrictive and categorical. The contrast in these paradigms highlight different qualities that shape the navigational experience of online collections: landscapes are expansive and open, yet employ boundaries, signposts and markers to denote meaning – a quality commonly attributed to data visualizations and topographical maps but also attributed to 3D games and non-linear narratives, whereas containers employ labels and nesting to employ a sense of position and hierarchy, such as that typically seen within a menu bar or within museum classification systems (Fig. 2.2).

In demonstrating a more concrete representation of the *landscape*, Andrews et al. [80] present a 3-dimensional visualization of hyperlinked data-entities in which the viewer can orient themselves within an expansive structure. A similar approach was adopted by Ruffaldi et al. [82] in presenting cultural heritage content, where the metaphor of an information landscape was fruitfully realized in its depiction of linked knowledge, textual data and rich media through the use of 3-dimensional elements and deep visual perspective. Goldfarb et al. [47] also adopt the 3-dimensional landscape paradigm in their presentation of semantically linked cultural heritage content. As an exploratory data visualisation, the interface was designed to entice users to explore artists and their influencers, presented through time as meandering 'rivers' across a broad landscape, as shown in Fig. 2.3. It is evident from these interfaces that they do not exist solely for purpose of aiding task efficiency or for the purpose of resolving technical, task-based problems, but as artefacts in their own right that are open to interpretation, reflection and critique. This calls for information seeking frameworks that recognise the landscapes embedded within these information spaces that describe a more playful [73] and serendipitous [74], approach to information seeking.

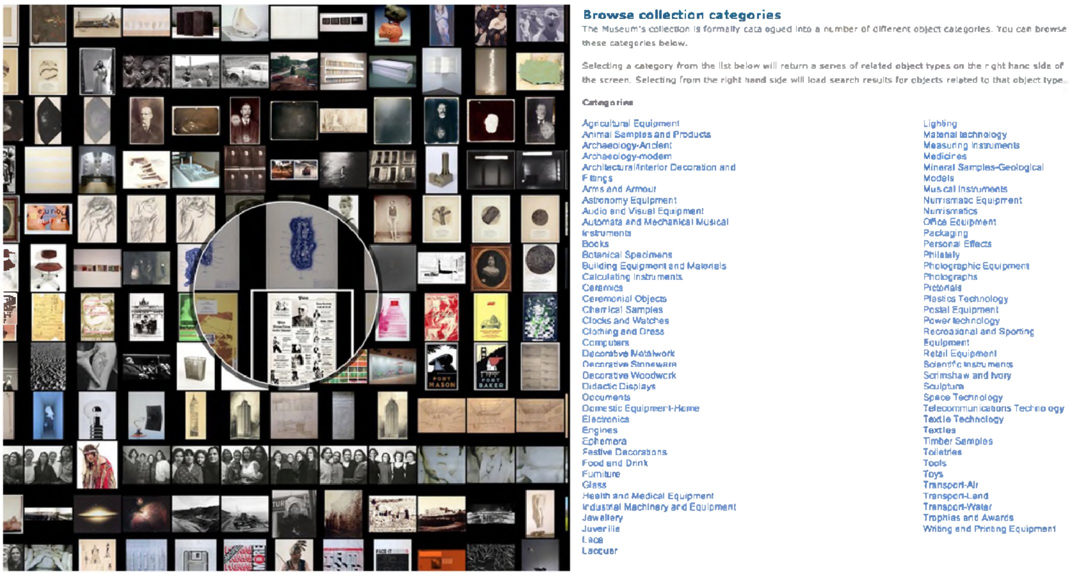

Figure 2.2: The *landscape* and *container* metaphor as represented by two different interfaces for online museum collections: SFMOMA's ArtScope, on the left, depicts an expansive *landscape* of collection content laid out in a pannable, zoomable visual browser, whereas the Powerhouse Museum's Online Collection, on the right, depicts categorical *containers* as hyperlinks.

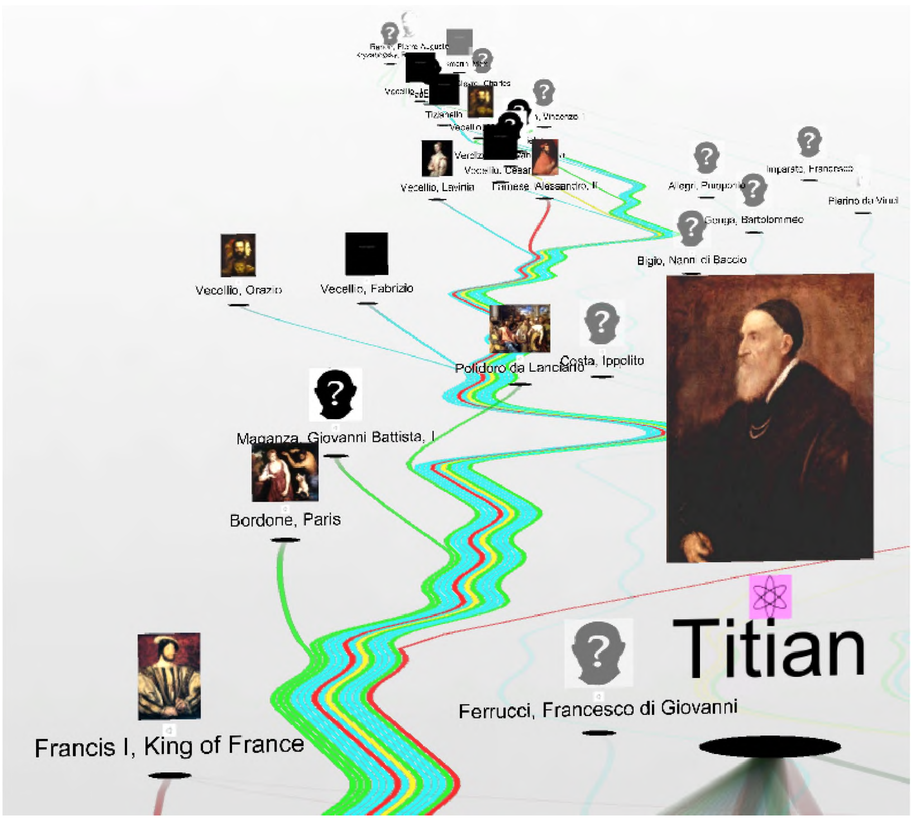

Figure 2.3: A representation of an artist and his successors, presented as 3-dimensional information landscape. cf. [47]

2.4.3 The information flaneur

Dörk et al. argue the need for new information seeking metaphors that characterize a particular kind of interaction that is primarily based on exploration, serendipity and discovery – one that is under-represented by existing interaction models and methodologies that primarily focus on externally defined user goals and task efficiency [15]. In particular, the authors draw parallels between the immersive qualities of the interwoven streetscapes of modern cities and the increasingly linked and enriched information spaces of the modern Web and propose the *information flaneur* as a persona that depicts the curious, creative and critical information seeker. This is in antithesis to the personas traditionally used in use case analysis that characterizes goal-driven interactions with computing, such as the attentive search agent who craves relevant information recall but demands precise results. The *information flaneur*, on the other hand, offers an almost recreational perspective of human-computer interaction: namely that of a wandering, curious information seeker. As described by Dörk et al. [15, p. 3]:

> "We are particularly interested in his exploratory mindset. In order to experience the city, the flaneur does not methodically navigate streets, checking each edifice like a building inspector in search of code violations. Nor does the flaneur hastily interrogate each city-dweller, like a police officer in search of a thief. Because the flaneur does not accurately scrutinize everything that crosses his path, he is able to sense what city life is about. The flaneur is the embodiment of exploration and serendipity, while the police officer and building inspector personify traditional search and browsing."

In addition to describing the explorative potential of emerging information spaces, Dörk et al. [15] also advocate interaction design conventions that promote exploration and curiosity. To entice the perceptual qualities of the *information flaneur*, they posit the use of techniques such as similarity-based suggestions and visual information surrogates [83]. The authors also advocate the use of *visual momentum*: referring to the way interfaces provide a smooth transition from one view to another. Interfaces that exhibit high levels of visual momentum require less cognitive effort to comprehend display changes, and employ techniques such as animated transitions, detail-on-demands and contextual zoomable views [84].

Other research, such as the work by McCay-Peet & Toms [74], have also revealed how digital interfaces can be designed to encourage exploration and serendipity. The authors state that serendipitous interfaces should enable connections, introduce the unexpected, present variety, trigger divergences and induce curiousity. These principles were adopted by Ennis-Butler et al. [3] in their work on exploratory data vi-

sualisations for museum collections as they recognize data itself as cultural medium and employ ideas of play as a way of engaging with complex data [50, 85]. Following the previously discussed museological notions of *context* and *connectedness* in Sections 2.1 and 2.2, Whitelaw [50] argues that the value that lies within cultural collections can be better harnessed by designing "generous" interfaces that promote immersive browsable views, prominent visual imagery of primary content, and depictions of context and relationships. These views parallel the design principles developed and empirically tested by McCay-Peet and Toms [74]: principles that allude to the perceptual needs of the *information flaneur*. While the information flaneur embodies creative, curious information seeking within information landscapes, another approach – pliability – articulates the *aesthetics* of interaction design.

2.4.4 Pliability

Pliability, as defined by Löwgren [14], refers to both the degree that an interaction facilitates exploration, serendipity, and playfulness in use, and the degree that the interaction feels malleable and tightly coupled in the sense that "the user is drawn into a sense of shaping the digital information with his/her fingertips, even though the actual artifact might employ standard, non-tactile interaction techniques such as mouse, keyboard, and a display monitor" [14, p. 71]. Löwgren [14] emphasizes pliability as an innate and sensuous quality that can only be perceived "in the moment" of the interaction: referring to both the tenuous nature of the work of interaction designers as they iteratively assess and "gain a feel" for designing interactions, and the ethereal and fleeting nature of these aesthetic interactions as perceived by end users [86]. This section examines two specific aspects of pliability: the tight coupling of action and response, and the idea that user interface and contents are the one and the same.

One of the key features of a 'pliable' design is the direct connection between action and response. In an early example of the idea, Löwgren demonstrates the direct data manipulation employed in Ben Shneiderman's Film Finder [87]. The Film Finder, shown in Fig. 2.4, presents a data visualisation of films represented as dots on a graph where the user can directly filter what dots are displayed by manipulating a series of sliders that apply criteria such as title, actor and length. The visualisation changes in direct response to the sliders, conveying the effect that the user is 'shaping' the information directly: every manipulation of the filters, however subtle or minute, is instantly reflected in the visualisation. Moreover, this form of manipulation introduces a *temporal* element of user interaction: not only is the user capable of interpreting rich data over a two-dimensional plot, but they

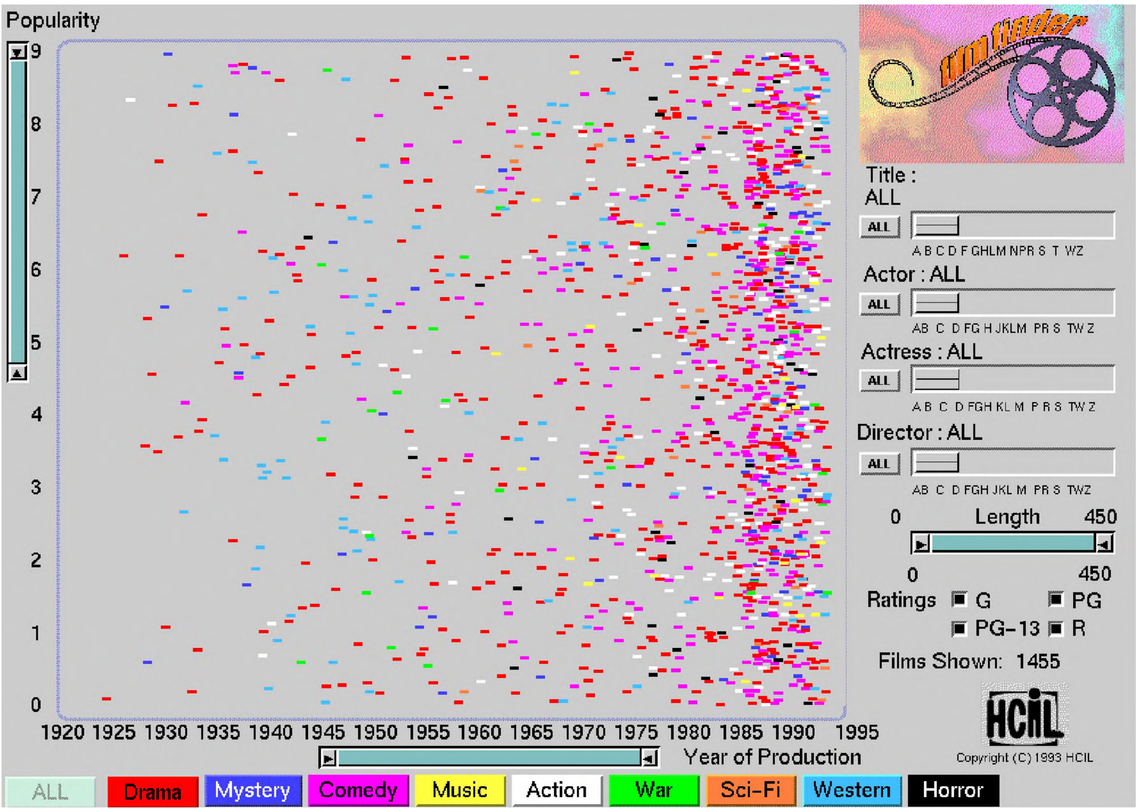

Figure 2.4: Ben Shneiderman's Film Finder demonstrates a concept of *pliability* using direct and fine-grained manipulation. The visualisation of films, represented as a scatter plot, changes in direct response to the sliders on the right.

are also capable of manipulating that data in real-time as the data-set shifts and changes in response to the sliders.

A similar form of interaction is conveyed by the Design Museum's app for the iPad. [4] The app, shown in Fig. 2.5, showcases a collection of objects from London's Design Museum. The objects are represented as a series of tiles randomly arranged as a grid. Users can apply filters to the collection, such as type of object, the artist and its medium, and the collection dynamically changes in response to the filters. As the filters change, the app's interface responds in real-time, as the collections grows and shrinks, the tiles 'flip' and hide in response to the criterion applied. Like the Film Finder, the presented data-set grows and expands in direct response to user input, that in turn, provide a direct and visual response to the way the collection is classified. The visuo-tactile nature of the interaction – the direct action of applying filters and its animated, skeuomorphic response – characterises it as strongly *pliable*: an interaction that goes beyond the instrumentality of applying filters and querying a data-set, to one that employs direct manipulation and animations to evoke a direct, seamless response to fine-grained interactions.

[4]The Design Museum App is available for the iPad at: `https://itunes.apple.com/au/app/design-museum-collection-for/id510964197?mt=8`

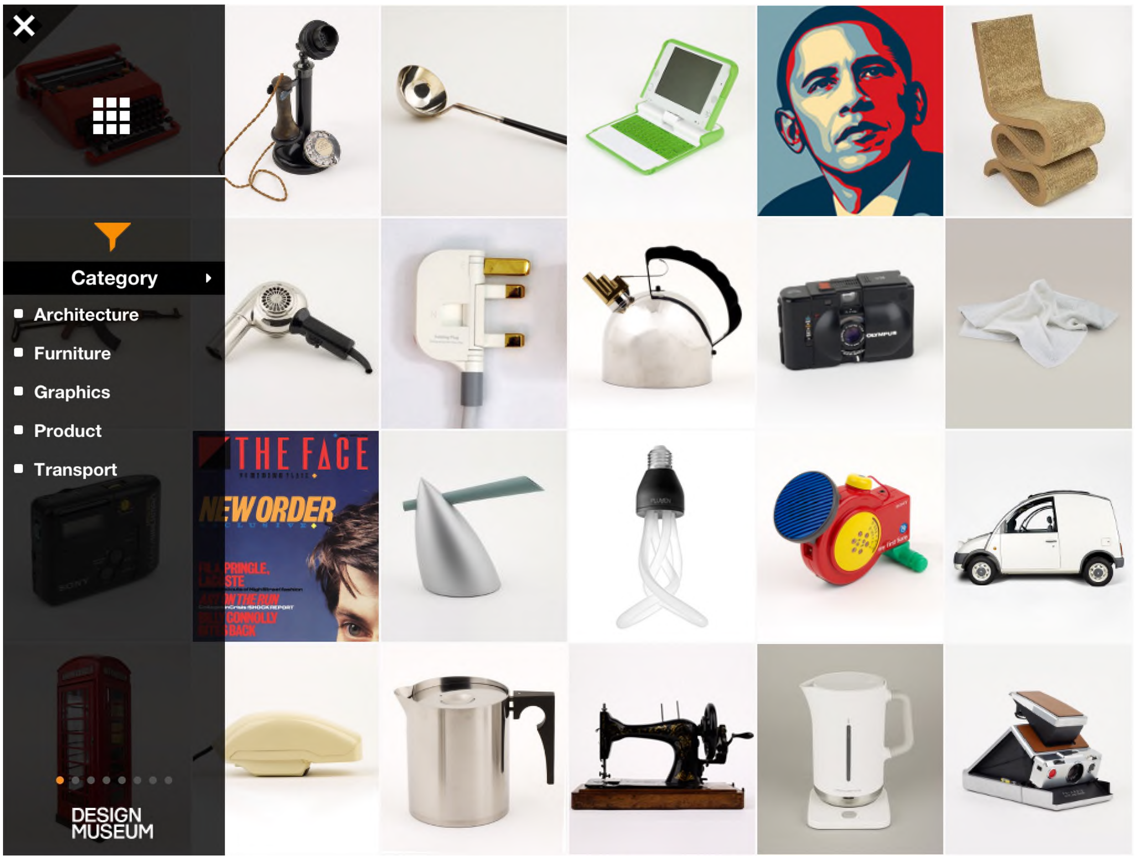

Figure 2.5: The Design Museum App for iPad displays a list of filters that, when applied, dynamically shrinks and expands its collection. *Pliability* specifically refers to its direct, visual and tactile response, rather than its functional purpose or instrumentality.

Pliability incorporates another dimension which can be realized through direct manipulation as well, one that "considers 'interface' and 'contents' [as] not different levels of abstractions but merely two aspects of the same experience" [14, p. 79]. Again, Löwgren demonstrates by example: in Google Earth, a user can directly drag and pan across the globe by manipulating it directly and adjusting a virtual camera, and any discrete transformation, such as responding to a search query to navigating to a specific place name, is dealt with through the use of a 'flyover' animation that smoothly transitions the user from one place to another. Although the program provides controls and buttons to manipulate its 3D map, most of the interaction could be accomplished by directly clicking and dragging on the map itself – the interface and its content are therefore the one and the same.

In another, more contemporary demonstration of this concept, Figure 2.6 presents two note-taking apps for iOS. The first, on the left, is the in-built Notes app for Apple's mobile operating system, and the second, by a third-party, is a to-do list app called Clear.[5] Both apps offer a very similar level of functionality: they allow the user to create, edit and delete short to-do lists. However, their differences lie in their user interaction, and more specifically, the way Clear conflates interface and contents. In Notes, the user taps on buttons to create, edit or delete notes: each action is rather discrete, and a portion of the interface is devoted towards showing the buttons. However, in Clear, the manipulation of to-do list items is performed directly on the interface itself: the user adds a list item by swiping down on the screen with their finger that 'pulls' a new item in, and deletes an item by swiping it left or right off-screen. It is this direct connection between interface and response, between gesture and action, rather than the lack of traditional user interface elements, that demonstrates this aspect of *pliability.* Clear achieves this not through using touch gestures per-se, but through a responsive interface; fluid animations that convey user actions; and the adoption of similar design conventions found elsewhere on the mobile operating system. Given the proliferation of touch-screen devices, it is likely that the elements of pliability – the tight coupling of both action and response and the blurring of interface and contents – continue to influence emerging possibilities of interaction for both touch and non-touch devices.

2.4.5 Pathways in information landscapes

Skov [37] found that online visitors demonstrated exploratory behaviors such as serendipity and exhibited meaning making qualities when visiting online museums. The research stated that such users sought highly visual experiences, such as the use of large and prominent photographs, and noted that exploratory behaviors were

[5] The Clear app is available for iOS devices at: `https://itunes.apple.com/au/app/clear-tasks-reminders-to-do/id493136154?mt=8`

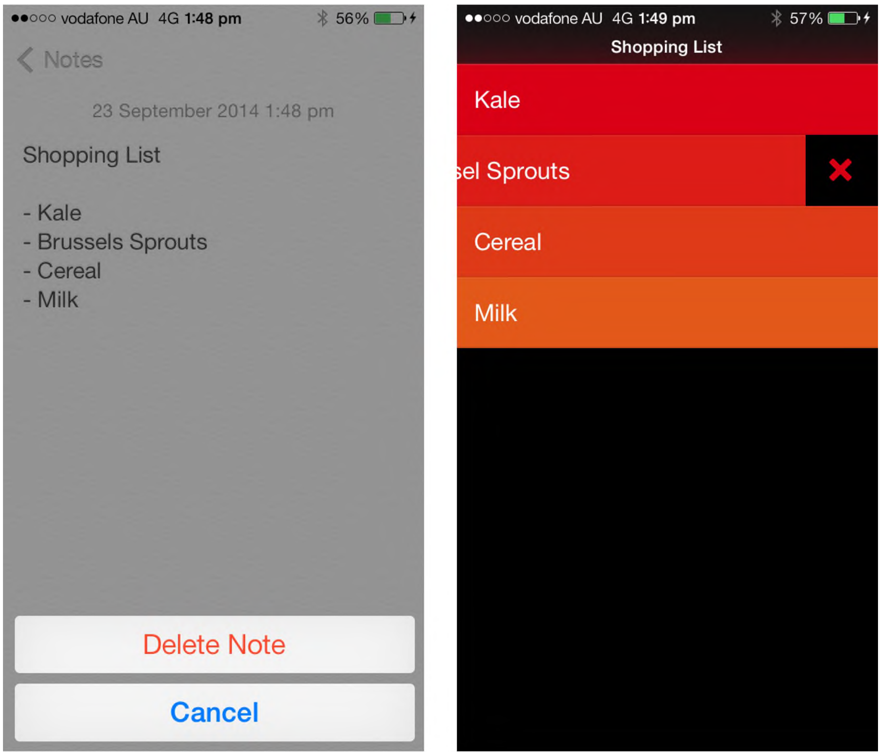

Figure 2.6: Two note taking apps on Apple's iOS demonstrating a 'delete' interaction. Notes, on the left, employs a standard button and confirmation dialog, whereas Clear employs a direct gesture on the data.

featured amongst a broader pattern of information seeking needs [67]. Goodale et al. [12] conceptualize the *pathway* as a guiding metaphor to characterize the design of digital artifacts that support creative and divergent exploration of cultural data-sets. Following from an empirical survey that examined information seeking requirements and behaviors of expert and non-expert users of museum collections, they elaborate the concept of pathways in two distinct dimensions. Firstly, pathways provide a mechanism that affords creative, non-linear interaction within a collection:

- A pathway can provide a starting point for a user's journey.
- A pathway in and itself is a route-through. A pathway can be new and exploratory, or can be "well-trodden" via the visualisation of other people's user journeys.
- A pathway is an augmented reality, allowing a user to travel through an offer an immersive experience of the collection, with a chance to lead to the unexpected.
- A pathway can offer a learning experience that connects objects, concepts, ideas and people.

Pathways were also interpreted as a means of semantically structuring content and information seeking journeys:

- A pathway is a search history traversed through a map, graph or structure.
- A pathway represents an information seeking journey.
- A pathway represents linked meta-data.

The research conducted by Goodale et al. [12] also examined the expected characteristics of interfaces that implement the pathway metaphor. In recognition of the divergent nature of information seeking [88], linearity was not seen as a way of maximising the potential for pathways as devices for exploration and information seeking, and that their starting and end-points should be seen as flexible rather than fixed so that users can begin and end an information seeking journey in a number of different ways. As semantic structures, pathways were expected to organize and link content by theme, date, location, narrative and author, and were also expected to connect works across a number of dimensions, offering several different perspectives on the path creator – the entity or the person that 'creates' the paths.

The authors also described the kind of browsing features needed to facilitate pathway-based exploration within information spaces, many of which have been discussed previously in this section. This includes features such as contextual overviews [18, 84] with details-on-demand [87, 89]; navigational context that visualises the next, last and nearby nodes within a pathway navigation [15]; opportunities to branch into new pathways or converge onto existing pathways and the ability to easily visualise, access and display objects on the pathway structure [84]. Pathways, and the objects on them, can be augmented with context and interpretation so that they can convey narrative and meaning. Hence, pathways can be posited as a way of visualising knowledge, highlighting connections and creating structures that support exploration, browsing and learning. The study also alludes to the use of pathways as a means of providing a ludic, playful and leisurely interactive experience, an idea that closely follows the previously discussed concepts of playfulness [73], serendipity among museum visitors [35, 74] and the *information flaneur* [15]. As Goodale et al. write [12, p. 77]:

> In addition to learning, paths may also serve to deliver entertainment and an enjoyable interaction experience for more general audiences. In practical terms, paths may simply be used as a means of introducing people to a collection and its stories, and in due course, encouraging them to venture further in a more independent fashion.

The authors also stress the importance of the non-linear nature of pathway-based interaction in providing and explorative browsing experience, with connotations to previously discussed notion of *connectedness* [33] within virtual museums. Goodale et al. [12, p. 77] write:

> Many of these characteristics are seen in existing path systems, but limitations arise from the linearity that is commonplace. Exploration and deeper levels of engagement within collections requires more complex path structures, carefully curated content, interpretation and narrative, and interconnectedness of paths and other content within and outside of the system.

In their study, the authors envision a system that allows for both path *creation* and *exploration*, where curators, educators and end users can create paths for others to explore (see: [90]). The case studies in this thesis experiment with the use of Formal Concept Analysis as a means of *automatic* path creation, an idea further explored in Chapter 3.

2.5 Summary of Design Principles

In following a *concept-driven interaction design* [20] approach, the above literature forms the basis of the design principles for the case studies presented in this thesis. These case studies account for the dual meaning of pathways as a *metaphor for creative exploration within information spaces* to support non-linear interaction and information seeking behaviors and as a *mechanism that semantically structures content* to support connectedness [33], meaning making [2, 37] and context [6] within museum collections. These principles are articulated in the context of role of the modern museum as knowledge organisations; the importance of highlighting context and narrative among objects; emerging models of information seeking that incorporate elements of play and creativity along with empirical research that examines design principles of serendipitous interfaces and the way users interact with museum collections. The following design principles are drawn from this research:

- Museums as collectors, creators and disseminators of *knowledge*. [5, 6, 8, 34]
- Highlighting *context* [6] and *connectedness* [33] within museum collections.
- Pathways as a *mechanism that semantically structures content*. [12]
- Pathways as a *metaphor for creative exploration*. [12]
- Designing for the *information flaneur*: the curious, critical and creative information seeker. [15]
- *Pliability*: malleable, responsive, smooth and tightly coupled interaction design that facilitates exploration and serendipity in use, and the direct connection between action and response. [14]
- The use of *visual momentum* to convey movement, motion and state changes. [15]
- Designing interfaces that employ high amounts of *visual imagery and similarity-based recommendations* [15, 37, 91]

The remainder of this thesis focuses on the realisation of these design principles through the construction of artefacts. Chapter 3 describes how Formal Concept Analysis mathematizes the idea of pathways and information landscapes in terms of conceptual structures, and Chapter 4 describes various design artifacts and the case studies that implement these principles.

Chapter 3

Formal Concept Analysis

Formal Concept Analysis (FCA) was developed in the early 1980s as a mathematization of the human cognitive constructs of concepts and concept hierarchies [92]. Although commonly used as a technique for data analysis, knowledge representation and visualisation [93], the case studies in this thesis primarily use it as a means of creating a conceptual space that can be navigated or explored by an end user. This chapter presents Formal Concept Analysis in terms of how its implemented in the case studies in Chapter 4: as a way of allowing users to create pathways through collections of conceptually related content.

Section 3.1 introduces the formal background Formal Concept Analysis. Then, in Section 3.2, the construct the of the concept lattice is introduced, along with how a user could navigate its conceptual neighbourhoods to explore pathways through a collection. Next, in Section 3.3, the notion of concept similarity is introduced. Finally, Section 3.4 introduces a new construct – the concept pathway – that uses concept similarity to generate tree structures conceptually related content based on a user's navigation choices.

This chapter presents a theoretical overview of Formal Concept Analysis as applied to the case studies in Section 4. All visual representations arising from the use of FCA are presented individually in the case studies in Section 4, and conversely, all examples in this chapter draw from actual museum collection data used by these case studies.

3.1 Concepts and Contexts

Formal Concept Analysis is based on the philosophy of human thought and communication. Beginning with an understanding of its cognitive constructs, *concepts* can be understood as basic units of thought formed by observations of existing phenomena "formed in dynamic processes within social and cultural environments" [17, p. 2]. According to this definition [17], a *concept* consists of a set of objects as its *extension*, and all attributes, properties and meanings that apply to those objects as its *intension*. As an example, if one considers the descriptor of "abstract paintings with geometric patterns" (its intension) and the actual 7 paintings that fit that description (its extension), then a concept is defined as the simultaneous perception of that intension and extension – i.e., the qualities of those paintings as *attributes* (or *formal attributes*) and the actual paintings as *objects* (or *formal objects*) defined by those attributes. In Formal Concept Analysis, concepts are mathematized as *formal concepts* defined as a pair (G, M) with a set of objects G (its extension) and a set of attributes M that describe those objects (its intension). In context of a museum collection, a formal concept (G, M) can be used to circumscribe a set of objects G that possess attributes or meta-data M. For example, the following is a formal concept (A, B) [1] that describes a set of works from the University of Wollongong Art Collection, [2] where A represents a set of titles of 3 works from the collection, and B represents a set of attributes that describe those works:

A = { 'Bush Rocks After the Rain', 'Solar Boat', 'Port Kembla Landscape' }
B = { 'abstract', 'painting' }

Fig. 3.1 shows how a concept could be perceived by a human through a visual representation of its works with an accompanying description in natural language. Most of the case studies described in Chapter 4 present formal concepts in this way – works are presented as thumbnails grouped into clusters or pathways that employ a semantic layer and natural language pipeline to generate a description of those concepts.

Formal Concept Analysis typically works within a *context* with a fixed set of attributes and objects. In this way, a context could be thought of as the entire set of objects within a museum data-set, and the meta-data and attribute values that

[1] A formal concept is commonly presented as (G, M) when defined, but use of notations (A, B), (C, D), (E, F) etc. are commonplace in describing different formal concepts within the same example: A, C and E represent the set of objects of those formal concepts, and B, D and F represent the set of attributes of those formal concepts.

[2] Works from the University of Wollongong Art Collection are represented in the *A Place for Art* case study described in Section 4.3

Figure 3.1: A visual representation and natural language description of a formal concept.

apply to those objects. In FCA, the context is mathematized as a *formal context* as $K(G, M, I)$ where G and M respectively describe its set of objects and attributes and I describes the associations between them. Formally, each $g \in G$ is interpreted as an object of that context and each $m \in M$ is interpreted as an attribute of that context, and $I \subseteq G \times M$ is a binary relation where $(g, m) \in I$ is read “object g has attribute m” or gIm. Table 3.1 shows a formal context as a cross-table, with rows representing objects and columns representing attributes. The presence of $\times$ in row i and column j indicates that the object i has attribute j, or conversely that the attribute in j describes the object in i. Given a formal context $K(G, M, I)$, for $A \subseteq G$ and $B \subseteq M$, the following two derivation operators $(\cdot)'$ define a Galois connection between the powersets of G and M:

$$A' = \{m \in M | (\forall g \in A) gIm\} \text{ for } A \subseteq G$$

$$B' = \{g \in G | (\forall m \in B) gIm\} \text{ for } B \subseteq M$$

These two operators $\{(\cdot)', (\cdot)'\}$ put in relation elements of the lattices $(2^G, \subseteq)$ of objects and $(2^M, \subseteq)$ of attributes [94]. A Galois connection induces a closure operator $(\cdot)''$ and realises a one-to-one correspondence between all closed sets of objects and all closed sets of attributes that are referred to as *concept extents* and *concept intents*, respectively. This Galois connection underlies a duality in Formal Concept Analysis in that concepts with a higher number of attributes have fewer objects than concepts with a lower number of attributes, and conversely, concepts with a higher number of objects have fewer attributes than concepts with a lower number of objects [95].

For a formal concept (A, B), A' is the set of attributes common to all objects in A, and B' is the set of objects common to all attributes in B. Hence, a formal

Table 3.1: The cross-table of formal context $K(G, M, I)$ containing information about the objects (G) and their attributes (M).

$K(G, M, I)$	painting	abstract	coarse brush strokes	geometric patterns	print	sculpture
Waiting, Port Kembla	×		×			
Bush Rocks After the Rain	×	×	×			
Port Kembla Landscape	×	×		×		
Large Jug		×		×	×	
Gateway to Mt. Keira		×		×		×
Solar Boat	×	×	×	×		

concept (A, B) of context $K(G, M, I)$ is a pair where $A \subseteq G$, $B \subseteq M$, $A' = B$, $B' = A$, where A is its *concept extent* (or just *extent*) and B is its *concept intent* (or just *intent*). As such, a formal concept of a formal context can be formed by taking an attribute and collecting all objects that it describes, and then again collecting all common attributes shared by those objects. Via this process, one could assert that "attribute m implies attribute n". Using the context in Table 3.1 as an example, let m be attribute 'coarse brush strokes' and C be the objects described by m:

C = { 'Bush Rocks After the Rain', 'Waiting, Port Kembla', 'Solar Boat' }

Now let D be the attributes common to C:

D = { 'painting', 'coarse brush strokes' }

This forms the formal concept (C, D) from the objects described by m = { 'coarse brush strokes' }, induces the additional attribute n = { 'painting' }. Via this process, one could assert that "all works with coarse brush strokes are paintings":

C = { 'Bush Rocks After the Rain', 'Waiting, Port Kembla', 'Solar Boat' }
D = { 'painting', 'coarse brush strokes' }

Formal concepts do not necessarily group objects into exclusive categories, and any formal concept can contain objects from another formal concept. Dually, any objects, or group of objects represented by a formal concept can share attributes from other formal concepts. For example, within the context of the collection, the

formal concept (C, D) shares objects 'Bush Rocks after the Rain' and 'Solar Boat' and the attribute 'painting' with (A, B). To this end, the two concepts (A, B) and (C, D) are different formal concepts, but share a degree of *concept similarity*, an idea further explored in Section 3.3.

Formal concepts of a context have an ordering, where formal concepts can be more specific or general than other formal concepts. A formal concept (A_1, B_1) is more specific than another concept (A_2, B_2) if the objects of the first formal concept are a sub-set of the second formal concept, and the attributes of the first formal concept are a super-set of the second formal concept. Hence, an order relation over formal concepts can be defined as:

$$(A_1, B_1) \leq (A_2, B_2) \iff A_1 \subseteq A_2$$

which is equivalent to:

$$(A_1, B_1) \leq (A_2, B_2) \iff B_1 \supseteq B_2$$

For example, as a comparison to the previous formal concept (C, D) consider the following formal concept (E, F):

$E = \{$ 'Bush Rocks After the Rain', 'Solar Boat' $\}$
$F = \{$ 'painting', 'abstract', 'coarse brush strokes' $\}$

Given that $E \subset C$ and $F \supset D$, $(E, F) < (C, D)$ – i.e., formal concept (E, F) represents a more 'specific' concept than formal concept (C, D). Likewise, (E, F) is a subconcept of (C, D), and (C, D) is a superconcept of (E, F).

The complete set of a concepts for a context $K(G, M, I)$ is denoted S [3]. Given that there is an ordering relation over all concepts in S, a conceptual structure, known as a *concept lattice* can be constructed and visualised in the form of a line diagram, highlighting the spatial properties of concepts and in particular, subconcept and superconcept relationships. [4]

[3] In texts such as [93], the set of all formal concepts of a formal context is commonly referred to as C, however, S is used in this chapter as C already refers to the extent of formal concept (C, D).

[4] The theorem of Formal Concept Analysis and concept lattices can be found in [17, 92, 93].

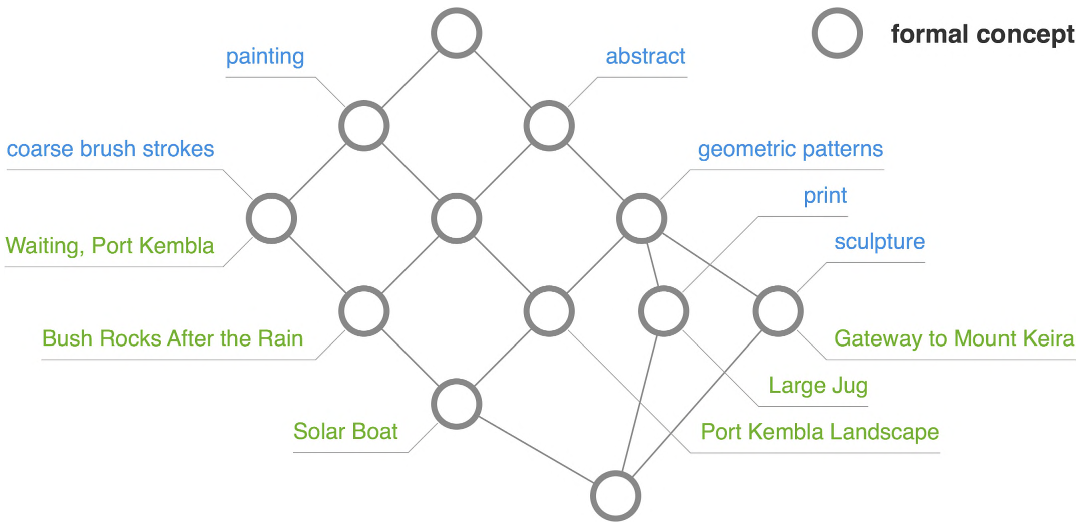

Figure 3.2: A concept lattice as a line diagram of the formal context shown in Table 3.1

3.2 Concept Lattices and Conceptual Neighbourhoods

A common method of constructing a knowledge space in Formal Concept Analysis is via the exploitation of the subconcept-superconcept relationship, i.e., the idea that certain formal concepts can be more general or more specific than other formal concepts. Within a given formal context, this can be represented via a *concept lattice*. [96]. [5] A line diagram, such as the one shown in Fig. 3.2, can be used to visualise concepts and their relationships within the concept lattice.

The concept lattice is generated from the formal context in Table 3.1. The nodes on the line diagram represent formal concepts, and the vertices represent their subconcept and superconcept relationships. [6] Using Fig. 3.2 as an example, the formal concept of each node can be derived by following the vertices and reading the labels on each node. Its objects can be derived by following the vertices and reading the green labels below each node traversed (without any upward connections), and its attributes can be derived by following the vertices and reading the blue labels above each node traversed (without any downward connections).

Fig. 3.3 shows the concept lattice annotated with labels of the formal concepts from Section 3.1. As an example, the objects of formal concept (A, B), as labelled in the figure, can be read by following the vertices below it so that you get the labels

[5]The definition of a lattice and a concept lattice, along with *The Basic Theorem on Concept Lattices* can be found in [93][p.13].

[6]In this thesis, concept lattices are presented with *reduced labelling* [97], where the labels of objects are shown in the extension of the most specific concept, and the labels of the attributes are shown in the intension of the most general concept. This ensures that the line diagram legible and easy to read while it still conveys information about the objects and attributes of each formal concept.

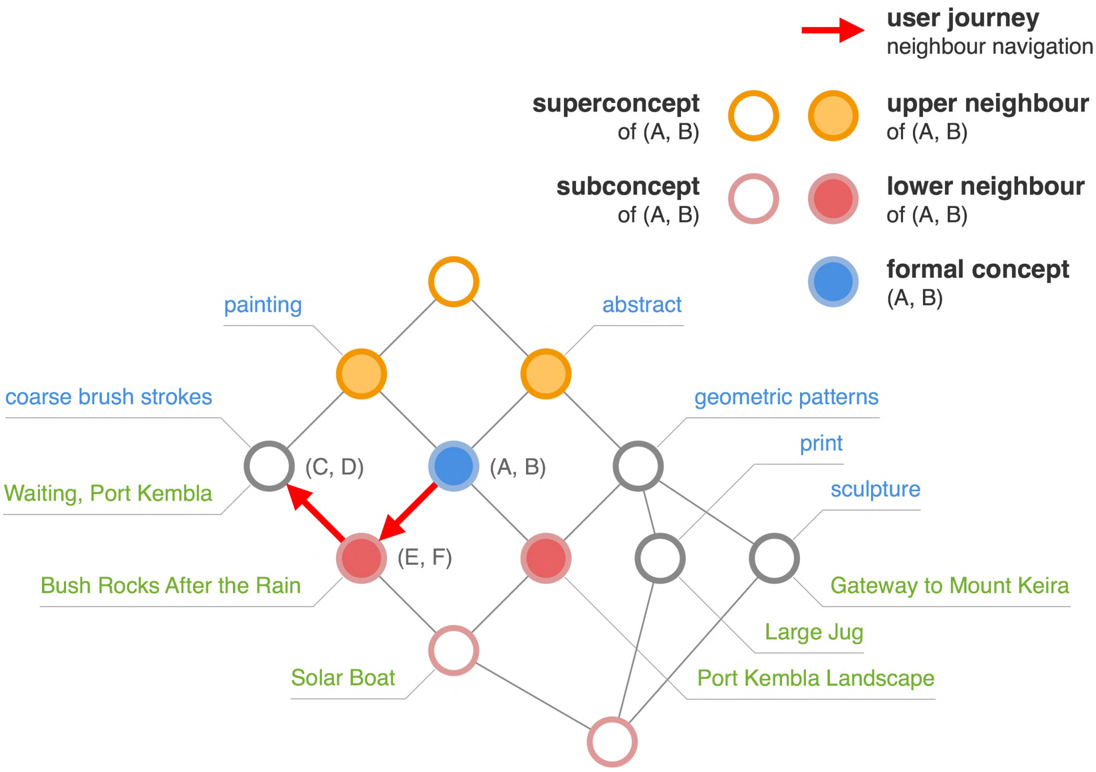

Figure 3.3: The concept lattice from Fig. 3.2 highlighting upper neighbours, lower neighbours, superconcepts, subconcepts, and a user's journey through neighbouring concepts.

'Bush Rocks After the Rain', 'Solar Boat' and 'Port Kembla Landscape'. Likewise, following the vertices above it returns its attributes, giving 'painting' and 'abstract'. In Fig. 3.3, concept (A, B) can be read from the line diagram as:

$$A = \{ \text{'Bush Rocks After the Rain', 'Solar Boat', 'Port Kembla Landscape'} \}$$
$$B = \{ \text{'abstract', 'painting'} \}$$

The other example concepts (C, D) and (E, F) are also annotated on the diagram, clearly showing their subconcept-superconcept relationships: $(E, F) < (A, B)$ and $(E, F) < (C, D)$, as indicated by the lower position of (E, F) on the diagram. More broader formal concepts exist at the top of the diagram, and more specific formal concepts exist at the bottom of the diagram. The top-most concept on the diagram is the formal concept that contains all objects in its context, along with the shared attributes of those objects (if any). Conversely, the bottom-most concept on the diagram is the formal concept that contains all attributes in its context, along with the shared objects of those attributes (if any). For the concept lattices in Figs. 3.2 and 3.3, since their formal contexts contain no attributes shared by all objects, or no objects shared by all attributes, their top-most and bottom-most concepts respectively contain all of the context's objects and none of its attributes,

and all of the context's attributes and none of its objects.

On the line diagram, the subconcepts of a formal concept can derived by following the vertices below it and reading each concept (without any upward connections) and its superconcepts can be derived by following the vertices above it and reading each concept (without any downward connections). For example, in Fig. 3.3, (A, B) has four subconcepts and three superconcepts, and (C, D) has three subconcepts and two superconcepts. [7]

The line diagram is also useful for determining a formal concept's upper and lower neighbours. For any formal concept on a line diagram, its *upper neighbours* are superconcepts that appear immediately above it, and its *lower neighbours* are subconcepts that appear immediately below it. In Fig. 3.3, for example, the formal concept (A, B) (shaded blue) has two upper neighbours (shaded orange) and two lower neighbours (shaded red). Formally, a concept (A_1, B_1) is said to be an upper neighbour of concept (A_2, B_2) if $(A_1, B_1) > (A_2, B_2)$ and there is no concept (A_3, B_3) that will give rise to $(A_1, B_1) > (A_3, B_3) > (A_2, B_2)$. Dually, a concept (A_1, B_1) is said to be a lower neighbour of concept (A_2, B_2) if $(A_1, B_1) < (A_2, B_2)$ and there is no concept (A_3, B_3) that will give rise to $(A_1, B_1) < (A_3, B_3) < (A_2, B_2)$.

Previous research efforts [98, 99] have explored the use of concept lattices to provide a representation of an information space where a user could view information about a formal concept's objects and attributes, and then navigate to that concept's lower or upper neighbour via query expansion or contraction. Via repeated steps of this user interaction, a user could effectively 'move through' and trace a pathway through the concept lattice. For example, as shown in Fig. 3.3, consider the scenario where one would 'move' from concept (A, B) to its lower neighbour (E, F).

$A =$ { 'Bush Rocks After the Rain', 'Solar Boat', 'Port Kembla Landscape' }
$B =$ { 'abstract', 'painting' }

$\Downarrow$

$E =$ { 'Bush Rocks After the Rain', 'Solar Boat' }
$F =$ { 'abstract', 'painting', 'coarse brush strokes' }

Moving from (A, B) to (E, F) adds the attribute 'coarse brush strokes' while further constraining the objects in the formal concept so that the object 'Port Kembla Landscape' is removed. This is a *specialisation* step that moves to a more specific formal concept with less objects. Likewise, consider the scenario where one would

[7] A more detailed explanation of the mathematical foundations that give rise to this can be found in [17, 97].

move from concept (E, F) to its upper neighbour (C, D):

$E = \{$ ‘Bush Rocks After the Rain’, ‘Solar Boat’ $\}$
$F = \{$ ‘abstract’, ‘painting’, ‘coarse brush strokes’ $\}$

$$\Downarrow$$

$C = \{$ ‘Bush Rocks After the Rain’, ‘Waiting, Port Kembla’, ‘Solar Boat’ $\}$
$D = \{$ ‘painting’, ‘coarse brush strokes’ $\}$

Moving from (E, F) to (C, D) excludes the attribute ‘abstract’ and therefore removes the constraint on the set of objects in the formal concept so that the object ‘Waiting, Port Kembla’ is added. This is a *generalisation* step that moves to a more general formal concept with more objects.

These *specialisation* and *generalisation* steps form part of an interaction approach where a user would navigate to neighbouring concepts in a concept lattice by progressively adding and excluding attributes from a collection of objects. This approach, called the *conceptual neighbourhood approach*, was used by a program, called *ImageSleuth* [98, 100] shown in Fig. 3.4. In this program, objects of a *focus concept* are presented as a series of thumbnails along with the option of moving to that concept’s upper neighbours via the removal of attributes presented at the top of the screen, or its lower neighbours via the inclusion of attributes presented at

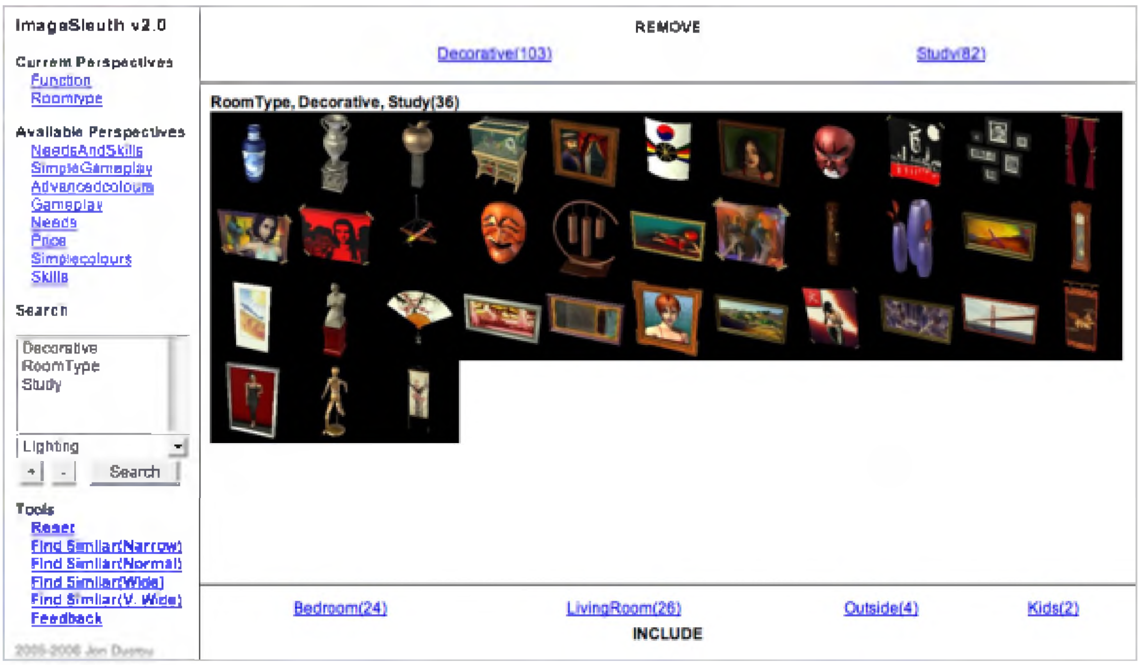

Figure 3.4: The user interface of *ImageSleuth*, showing the objects of a focus concept with links to upper and lower neighbours (shown top and bottom).

the bottom of the screen. When a user moves to an upper or lower neighbour, that formal concept becomes the focus concept, and its upper and lower neighbours are subsequently displayed. Over time, the user would create a pathway and traverse the concept lattice of its context. Although the website exposes sub- and superconcept relationships, the user would never see the concept lattice line diagram as presented in Fig. 3.2. In fact, the entire concept lattice is never computed – only the focus concept and its surrounding neighbours are computed each time the user moves from one focus concept to another. This approach, while suited for moving between upper and lower neighbours within a concept lattice, overcomes the scalability concerns of computing the entire concept lattice, along with the practical concerns of rendering a complex line diagram for a large, non-trivial formal context [93]. Although there are a number of algorithms for computing only the formal concept and its upper and lower neighbours [93], in all cases, the time to compute a formal concept and its upper and lower neighbours is significantly less than it would be to compute the entire concept lattice.

In Formal Concept Analysis, *conceptual scales* can be used to store views that partition the data being analysed [101]. [8] A conceptual scale could be thought of as a group of attributes that describe objects according to a type or sub-type (e.g. 'painting', 'drawing', 'line drawing'); a value or range of values (e.g. 'work dated before 1900, 'work dated between 1950 and 1960', 'work dated after 1960'); or a set of qualities with a natural ordering (e.g. 'mildly agree', 'agree', 'strongly agree') – such scales are respectively described as *nominal scales*, *interval scales* and *ordinal scales* [95, 101]. From an information science perspective, conceptual scales can be thought of as *facets* of a domain that can be applied as filters to produce *derived contexts* – parts of a formal context where objects are delimited by a subset of attributes or objects from its formal context. These derived contexts can be used to generate concept lattices – each generated from a single or combination of conceptual scales – that show concept hierarchies based on a selection of one or more facets as applied to the collection. For example, the formal context in Table 3.1, and its subsequent concept lattice shown in Fig. 3.1, is based on a *derived context* generated from the combination of the *medium* and *visual style* scales shown in Tables 3.2 and 3.3. Conceptual scales can be combined in other ways to provide custom views or 'lenses' on the collection: for example, the *Virtual Museum of the Pacific* case study (Section 4.2) – which uses the *conceptual neighbourhood* approach as a way of navigating its collection – provides 29 pre-defined conceptual scales that can be selected and combined by the user to filter the set of objects in view and influence the way neighbouring concepts are related.

[8] A formal definition of *conceptual scaling* in Formal Concept Analysis can be found in [101].

Table 3.2: A cross-table of objects and attributes representing the *medium* conceptual scale, which forms part of the derived context K shown in Table. 3.1

K_1	painting	print	sculpture
Waiting, Port Kembla	×		
Bush Rocks After the Rain	×		
Port Kembla Landscape	×		
Large Jug		×	
Gateway to Mt. Keira			×
Solar Boat	×		

Table 3.3: A cross-table of objects and attributes representing the *visual style* conceptual scale, which forms part of the derived context K shown in Table. 3.1

K_2	abstract	coarse brush strokes	geometric patterns
Waiting, Port Kembla		×	
Bush Rocks After the Rain	×	×	
Port Kembla Landscape	×		×
Large Jug	×		×
Gateway to Mt. Keira	×		×
Solar Boat	×	×	×

The *conceptual neighbourhood* approach of creating pathways in concept lattices as reported in the *Virtual Museum of the Pacific* case study (Section 4.2) is based on the interaction paradigm set by *ImageSleuth* (Figure. 3.4). Section 4.1.3 describes the algorithms used by its supporting framework, *CollectionWeb* to compute formal concepts of a context and implement the conceptual neighbourhood approach. Overall, this approach provides a useful means of expressing and navigating concept hierarchies – exposing users to the relationship among objects and allowing them to create pathways through the collection. In addition to relying on the ordering of concepts via subconcept and superconcept relationships, formal concepts can also be related in terms of *concept similarity*. Concept similarity is based on the notion that formal concepts, regardless of their position within the concept lattice, can be identified as being semantically close. The technique has generated a number of methods and applications towards understanding rich data-sets [102–104]. Section 3.3 presents a brief overview of a method of determining *concept similarity* as applied to the case studies within this thesis.

3.3 Concept Similarity

A technique for measuring *concept similarity*, as devised by Saquer and Deogun [104], is based on the common overlap of objects and attributes of two formal concepts. The *concept similarity* of two formal concepts (A, B) and (C, D) is defined as follows:

$$s((A, B), (C, D)) = \frac{1}{2}\left(\frac{|A \cap C|}{|A \cup C|} + \frac{|B \cap D|}{|B \cup D|}\right)$$

To demonstrate concept similarity by example, consider the formal concept (A, B) as reintroduced from Sections 3.1 and 3.2:

$A = \{$ ‘Bush Rocks After the Rain’, ‘Solar Boat’, ‘Port Kembla Landscape’ $\}$
$B = \{$ ‘abstract’, ‘painting’ $\}$

Table 3.4 presents a list of all formal concepts from the concept lattice in Fig. 3.2, along with their computed similarity measures with respect to (A, B), ordered by rank from ‘most similar’ to ‘least similar’. [9] In this way, concepts can be spatially represented based on their similarity to other formal concepts, as shown in Fig. 3.5.

For a given formal concept, finding its most similar concept within a context requires the set of formal concepts S from a formal context $K(G, M, I)$, which, as demonstrated further in Section 4.1.3, can be implemented as a pre-computed set that can be updated given incremental changes to a formal context. This, however, does not rely on computation of a complete or partial concept lattice, nor does it require computation of a concept’s upper and lower neighbours. This provides a practical advantage of using concept similarity as a means of representing spatial distances between formal concepts on large formal contexts where it may not be feasible to compute complete or partial concept lattices due to the fact that the worst case time complexities of these algorithms are based on the size of the formal context [93]. Concept similarity can be used to depict structures that are not based on complete or partial concept lattices and hence avoid the complexity costs associated with their generation. One such example, as demonstrated in the next section, is the notion of a *pathway*, a conceptual structure based on the similarity of concepts.

[9] Formal concepts that have an equal similarity measure with respect to formal concept (A, B) share the same rank. A concept similarity measure of 1 indicates that the two formal concepts share identical sets of objects and attributes. The formal concept with the empty attribute list $\{ \}$ is the top-most formal concept in the concept lattice in Fig. 3.2 that contains all objects but no attributes from its formal context. The formal concept with the attribute list $\{$ ‘abstract’, ‘geometric patterns’, ‘painting’, ‘coarse brush strokes’, ‘print’, ‘sculpture’ $\}$ is the bottom-most formal concept in the concept lattice in Fig. 3.2 that contains all attributes but no objects from its formal context.

Table 3.4: Similarity measures of all formal concepts in Fig. 3.2 with respect to formal concept (A, B) where $B =$ { 'abstract', 'painting' }

.

Rank	Intent (attributes) of formal concept	Similarity
1	{ 'abstract', 'painting' }	1.00
2	{ 'abstract', 'painting', 'geometric patterns' }	0.67
2	{ 'abstract', 'painting', 'coarse brush strokes' }	0.67
4	{ 'painting' }	0.63
5	{ 'painting', 'coarse brush strokes' }	0.58
6	{ 'abstract' }	0.55
7	{ 'abstract', 'geometric patterns' }	0.50
8	{ 'abstract', 'geometric patterns', 'painting', 'coarse brush strokes' }	0.42
9	{ }	0.25
10	{ 'abstract', 'geometric patterns', 'painting', 'coarse brush strokes', 'print', 'sculpture' }	0.17
10	{ 'abstract', 'geometric patterns', 'sculpture' }	0.17
10	{ 'abstract', 'geometric patterns', 'print' }	0.17

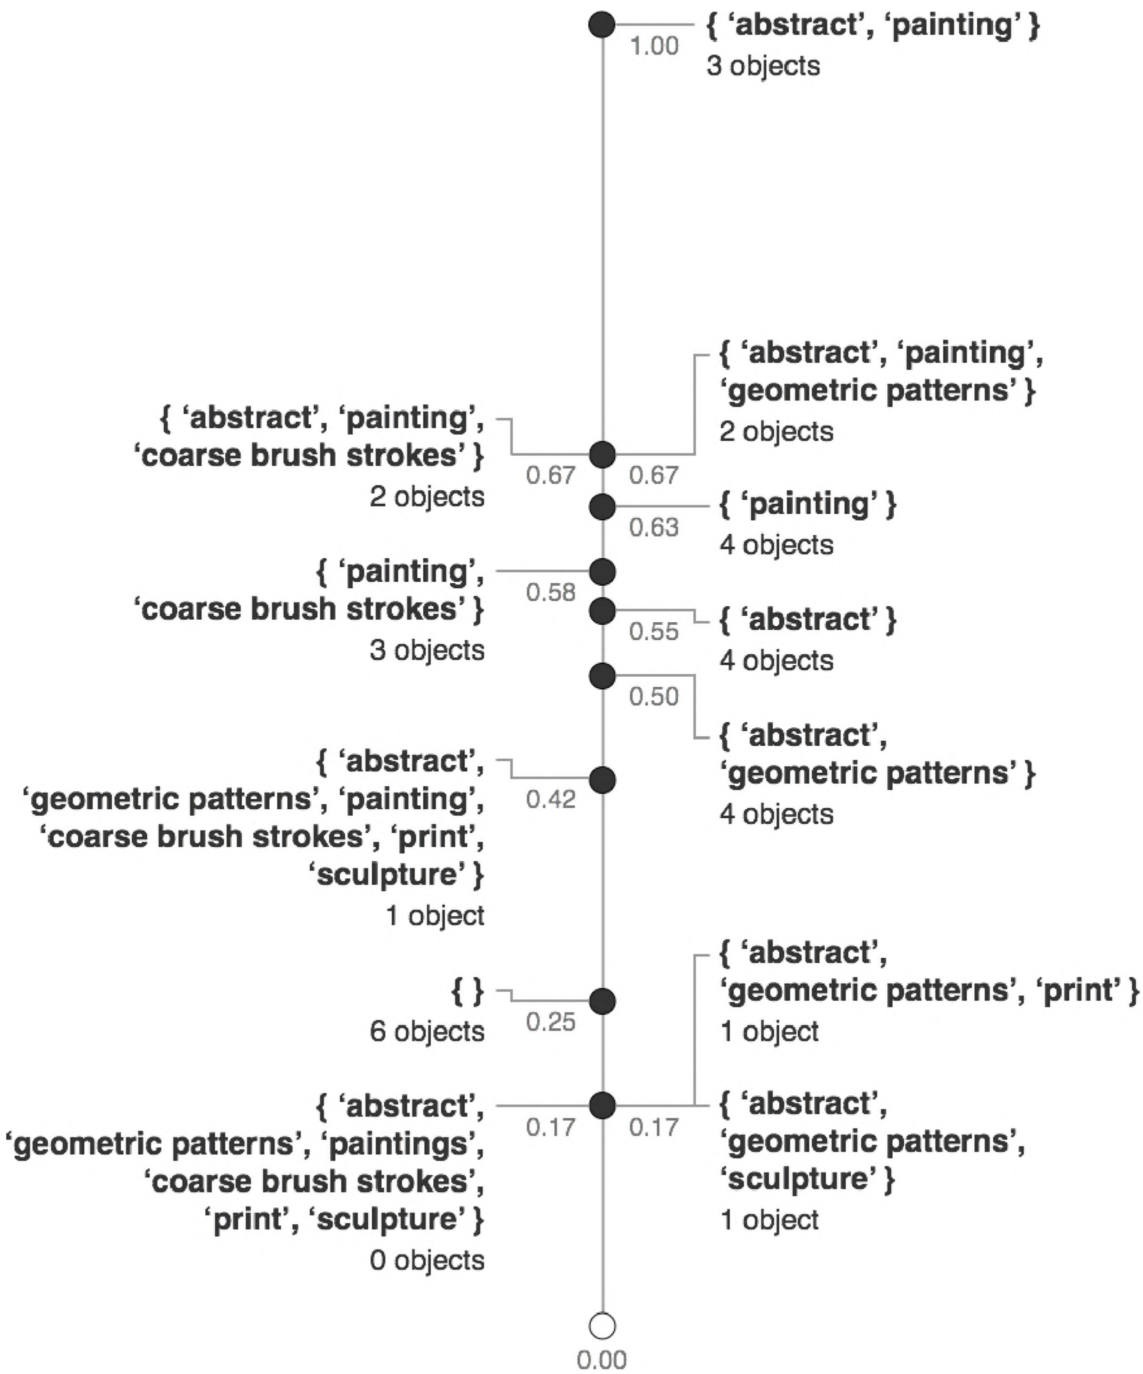

Figure 3.5: Line depicting similarity measures of all formal concepts in Fig. 3.2 relative to formal concept (A, B) where $B = \{$ 'abstract', 'painting' $\}$. More similar formal concepts to (A, B) appear at the top of the line, and less similar formal concepts appear at the bottom of the line.

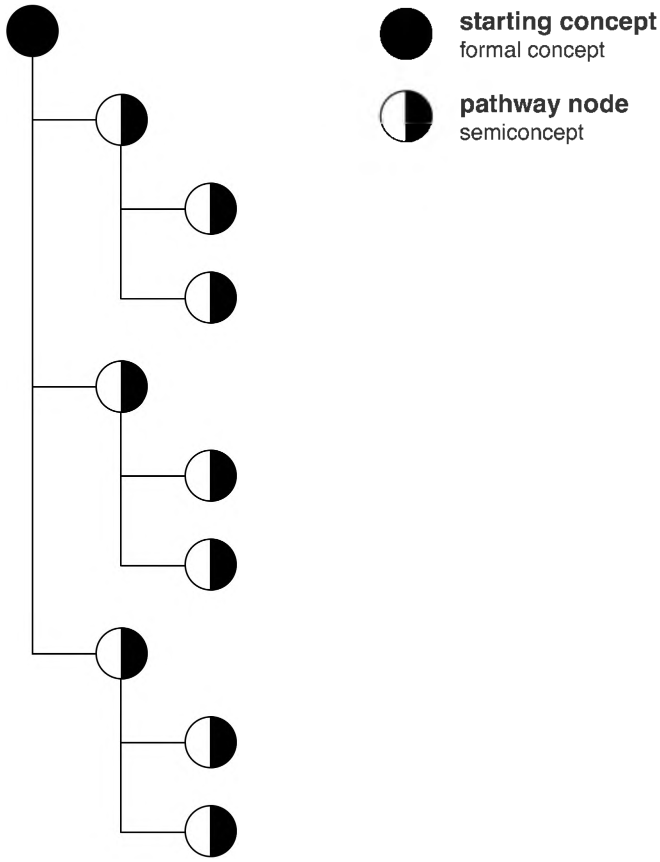

Figure 3.6: A concept pathway depicting a starting concept as its root with pathway nodes as its descendants, each related to their parent by concept similarity.

3.4 Pathways

The *concept similarity* measure, as applied over a set of formal concepts, can be used to generate conceptual structures that can aid in exploration and browsing of an information space. One such example is the *concept pathway*, a conceptual structure based on the *concept similarity* of formal concepts, or objects within formal concepts. This section describes how *concept pathways* are generated, and how they are used to present an associative browsing experience of formal concepts as presented by the case studies in this thesis.

A *pathway* can be described of as a tree of formal concepts, beginning with a root which is referred to as a *starting concept*. The starting concept has a set of child semiconcepts that are conceptually similar to its parent, and each child semiconcept – being a *pathway node* within the tree – also has a set of child semiconcepts that are related to it via concept similarity. Figure. 3.6 shows this representation of a starting concept and its descendant child pathway nodes, each being a semiconcept of its formal context. Although the tree shown in Figure. 3.6 is three levels deep, there is no arbitrary limit to its depth or the number of branches.

As a data structure, the tree consists of a single formal concept from a set of concepts S of its formal context. This concept is referred to as the *starting concept*. This *starting concept* has, associated with it, a generated set of *semiconcepts* as its descendants. A *semiconcept* is a formal concept that contains only a subset of a formal concept's objects or attributes. For example, consider the following pair (X, Y), a semiconcept of the formal context in Table 3.1.

$X = \{$ 'Bush Rocks After the Rain' $\}$
$Y = \{$ 'painting', 'coarse brush strokes' $\}$

Formally, a semiconcept of formal context $K(G, M, I)$ is a pair (V, W) consisting of a set of objects $V \subseteq G$ and a set of attributes $W \subseteq M$ such that $V = W'$ **or** $W = V'$. Given that the definition of a semiconcept also partly defines the definition of a formal concept, all formal concepts are semiconcepts, but not all semiconcepts are formal concepts. Therefore, (X, Y) is a semiconcept of the formal context in Table 3.1 because although 'painting' and 'coarse brush strokes' are attributes of the object 'Bush Rocks after the Rain' they are also attributes of the objects 'Waiting, Port Kembla' and 'Solar Boat'. As a point of comparison, consider the following pair (A, B):

$A = \{$ 'Bush Rocks After the Rain', 'Waiting, Port Kembla', 'Solar Boat' $\}$
$B = \{$ 'painting', 'coarse brush strokes' $\}$

As compared to the semiconcept (X, Y), (A, B) is a *formal concept* (where (X, Y) isn't) as (A, B) satisfies the following: $A' = B$ **and** $B' = A$ where $A \subseteq G$ and $B \subseteq M$ for formal context $K(G, M, I)$ in Table 3.1 – i.e., 'painting' and 'coarse brush strokes' are attributes common to objects 'Bush Rocks After the Rain', 'Waiting, Port Kembla' and 'Solar Boat', which are the objects shared by attributes 'painting' and 'coarse brush strokes'. By contrast, for pair (X, Y), although $X \subseteq G$ and $Y \subseteq M$ for formal context $K(G, M, I)$ in Table 3.1, $X' \neq Y$ and $Y' \neq X$, therefore (X, Y) is a semiconcept, but *not* a formal concept.

In the *conceptual neighbourhood* approach, a user traverses the concept lattice by navigating upper and lower neighbours, as shown in Fig. 3.3. Similarly, in the *concept pathway* approach, a user navigates through a set of formal concepts by traversing between parent and child pathway nodes within the tree structure. Concept pathways are generative and are determined based on current and previous user navigations on the pathway – aside from the starting concept, all subsequent child concepts are determined based on the current and previous concepts that the user

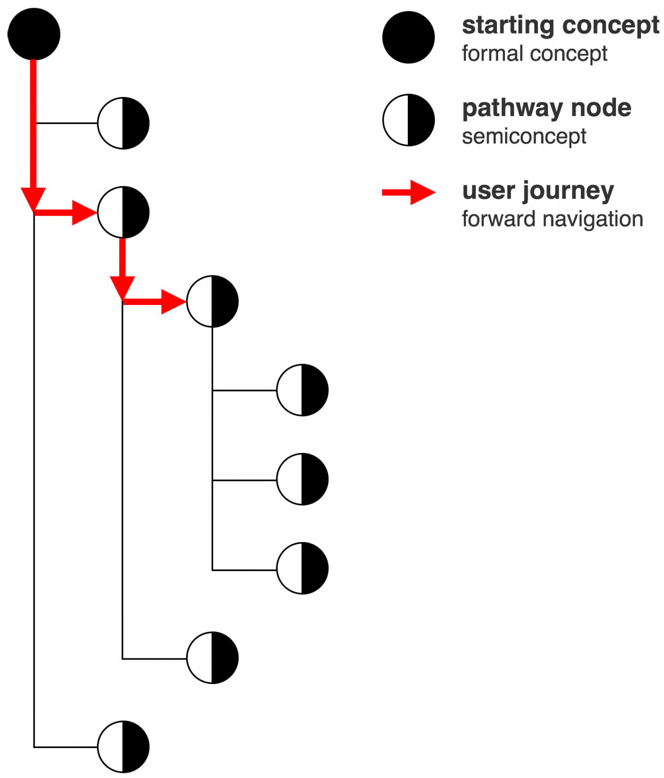

Figure 3.7: A concept pathway depicting a user journey through descendants of a starting concept. The tree structure forms as the user navigates through pathway nodes.

navigated to. In this approach, the user never navigates a fixed, pre-determined pathway structure: the selection of child semiconcepts for any pathway node occurs at query time as the user navigates to it, and the overall tree structure is determined by the user's navigation history, as shown in Figure. 3.7. In this approach, a user is also free to backtrack and traverse through previous pathway nodes and descend to alternate branches while the tree structure itself retains nodes that have been previously traversed, as shown in Fig. 3.8.

There are two kinds of *concept pathways* – each referring to how child concepts are selected on each pathway node: *concept-linked* pathways and *object-linked* pathways. They are described as follows:

3.4.1 Concept-linked pathways

On concept-linked pathways, the child concepts of a pathway node's formal concept are selected according to the similarity measure of itself with other formal concepts from its formal context. The formal concepts that have the highest concept similarity to the pathway node's formal concept are selected as its child concepts, given that those child concepts contain objects that are not already existing in formal concepts on the pathway structure. The number of child concepts selected is arbitrary and can

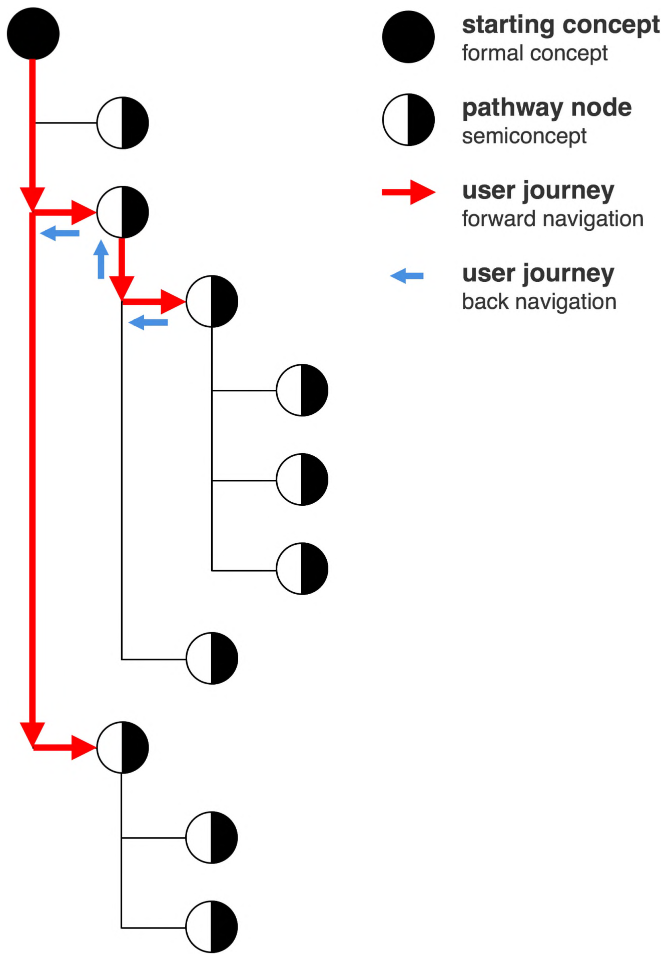

Figure 3.8: A concept pathway depicting a user journey through descendants of a starting concept. Users can backtrack through the structure and create new descendant nodes, while descendant nodes created by their previous navigations are retained.

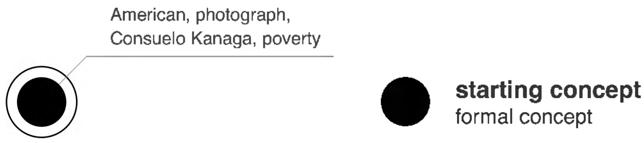

Figure 3.9: The starting concept of the concept pathway showing a list of its attributes { 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }.

be determined by the application or selected by the user. *Concept-linked pathways* provide the navigation and user interface for the *The Brooklyn Museum Canvas* (Section 4.4) and, based on that case study, will be described by example.

The concept-linked pathway begins with a starting concept. For the example demonstrated in Fig. 3.9, the starting concept consists of a set of works from the Brooklyn Museum's collection: 27 photographs by Consuelo Kanaga that depict people in poverty, labelled with its set of attributes, { 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }. The process of pathway formation begins with the starting concept (A, B) and an empty set of *retained objects* R that contains the union of all objects in formal concepts and semiconcepts that are currently on the pathway structure. For now, the only formal concept on the pathway structure is (A, B), so $R = A$.

The next step is to determine all child concepts of (A, B), designated as *children*(A, B). The pseudocode for determining *children*(A, B) is shown in Algorithm 1 and is described as follows: first, the set of all formal concepts S of its formal context are selected. An ordered array, T, is used to store each formal concept of S with its similarity score to (A, B). For each concept (E, F) in S, *similarity*$((A, B), (E, F))$ is computed (see Section 3.3) then combined with (E, F) and then added to T. T is then sorted by each element's *similarity*$((A, B), (E, F))$ value in descending order, so that the most similar concepts with respect to (A, B) appear first in the array of T. Further, for each formal concept (E, F) in T, a semiconcept (X, Y) is formed where $X = E \setminus R$ and $Y = F$. For each semiconcept (X, Y), if $|X| > 0$, then (X, Y) is added to *children*(A, B) until *children*(A, B) contains three child concepts.

This selection process is illustrated by example in Table 3.5. The table shows the 9 'most similar' formal concepts to the starting concept (A, B) where $B = \{$ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' $\}$. Given that (A, B) is the starting concept, the set of retained objects R – or objects already existing on the pathway – is equal to A. The first, third, fourth, fifth, sixth and seventh concepts are excluded from selection, given that for each concept (E, F), $|X| = 0$ in (X, Y) given that $X = E \setminus R$ and $Y = F$. These excluded semiconcepts appear as striked out in Table 3.5. The remaining three semiconcepts – the concepts with attributes { 'American', 'photograph', 'poverty' }, { 'American', 'photograph', 'Consuelo Kanaga',

Input: Concept (A, B); the number of semiconcepts to retreive n, the set of retained objects R; the set of all formal concepts S in formal context $K(G, M, I)$
Output: The set of child semiconcepts for concept (A, B)

$R := R \cup A$;
$children(A, B) := \varnothing$;
$T := []$;

foreach (E, F) *in* S **do**
  $similarity((A, B), (E, F)) := \frac{1}{2}\left((|A \cap E|/|A \cup E|) + (|B \cap F|/|B \cup F|)\right)$;
  $conceptAndSimilarityScore := ((E, F), similarity((A, B), (E, F)))$;
  Add $conceptAndSimilarityScore$ to T;
end

Sort T by $similarity((A, B), (E, F))$ descending;

while $|children(A, B)| < n$ **do**
  $(E, F) :=$ shift first element from T;
  $X := E \setminus R$;
  $Y := F$;
  if $|X| > 0$ **then**
    Add (X, Y) to $children(A, B)$;
  end
end

Algorithm 1: The GetChildConcepts algorithm used for retrieving child semiconcepts of concept (A, B) for the set of all formal concepts S in formal context $K(G, M, I)$.

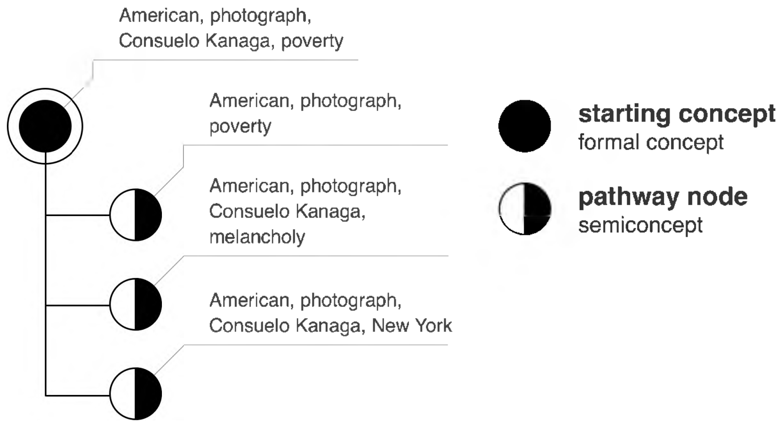

Figure 3.10: The starting concept of the concept pathway and its child concepts, generated via the selection of semiconcepts from Table 3.5

'melancholy' } and { 'American', 'photograph', 'New York', 'Consuelo Kanaga' }, are selected as child concepts and are appended as child concepts to the pathway structure in Fig. 3.9. These selected semiconcepts appear highlighted in Table 3.5. The selection process ensures that for any child concept of a starting concept or pathway node, its set of objects will only contain 'new' objects that do not already exist in other formal concepts on the pathway structure. The process also avoids circularity and repetition of descendant concepts on the pathway, so that $(G, H) \notin children(A, B) | (G = A \wedge H = B)$.

Figure. 3.10 shows the computed child concepts that are now attached to its starting concept. Concept pathways are generated based on user interaction: with the exception of the starting concept, the child concepts for any pathway node are only computed if a user chooses to navigate to it. The running example will now present how the pathway structure would be generated if a user navigated to the second child concept, which is the pathway node with the attributes { 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' }.

For a pathway node, the steps taken to compute its child concepts are almost identical to the way child concepts of the starting concept are computed, except the set of *retained objects* R now contains the objects from the formal concept that the user just navigated from along with the objects of the child concept that the user will be navigating to. Thus, for this example, R is the union of the two sets of objects from the following formal concepts:

{ { 157696, ... }{ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' } }
{ { 157696, ... }{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' } }

Table 3.5: An example of how child concepts are selected for the starting concept (A, B) in Fig. 3.9 where B = { 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }. Selected semiconcepts are highlighted, excluded semiconcepts are striked out.

Starting concept (A, B), **Retained objects** $R = \varnothing$

$|A| = 27$

B = { 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }

$R = R \cup A, |R| = 27$

List of formal concepts T with highest concept similarity score to (A, B).

For each concept (E, F) in T, form a semiconcept (X, Y) where

$X = E \setminus R$ and $Y = F$. Select semiconcepts where $|X| > 0$.

Y	$\|E\|$	$\|X\|$	Sim.
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	27	0	1.00
{ 'American', 'photograph', 'poverty' }	31	4	0.74
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	16	0	0.71
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	15	0	0.69
~~{ 'American', 'photograph', 'New York', 'Consuelo Kanaga', 'poverty' }~~	13	0	0.66
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.56
~~{ 'American', 'woman', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.56
{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' }	25	9	0.56
{ 'American', 'photograph', 'New York', 'Consuelo Kanaga' }	16	3	0.55

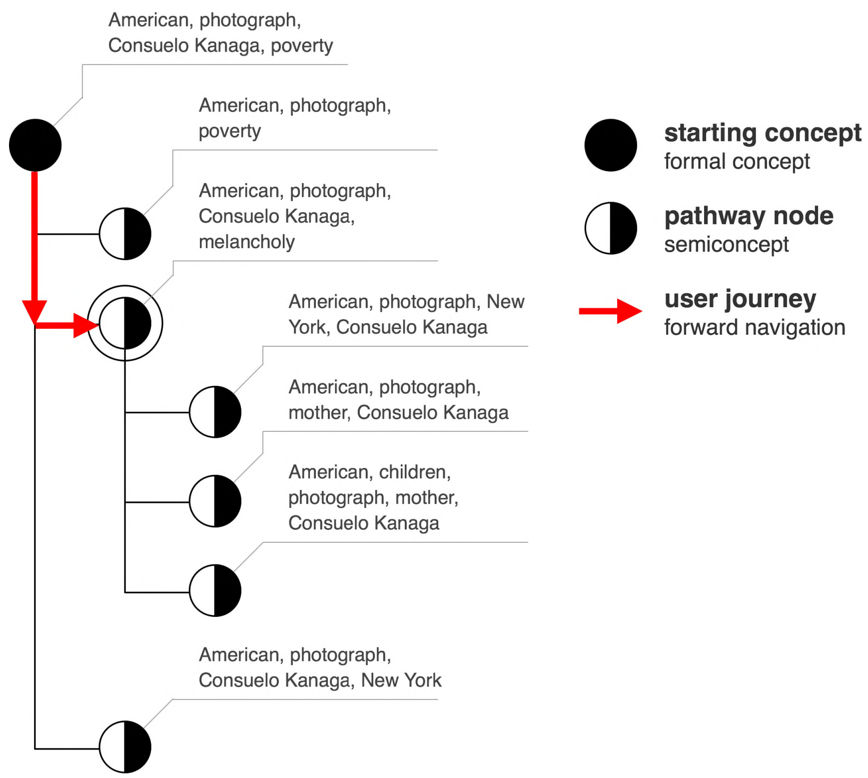

Figure 3.11: The now expanded pathway structure showing a user navigation to the formal concept with attributes { 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' } and its expanded child concepts, generated via the selection of semiconcepts from Table 3.6

Instead of searching for formal concepts that have the highest similarity measure with respect to the pathway's starting concept, the algorithm now searches for formal concepts that are most similar to the formal concept with attributes { 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' }, retrieving only semiconcepts with objects not in R. This selection process is illustrated by example in Table 3.6, with its child concepts now attached to the pathway structure in Figure. 3.11.

Concept pathways permit the expansion of any pathway node on the structure, allowing the user to backtrack and navigate to other nodes further up the tree. The example that follows demonstrates how the pathway structure would look if a user decided to go back one level on the pathway structure and navigate to the third child concept of the starting concept – the formal concept with attributes { 'American', 'photograph', 'New York', 'Consuelo Kanaga' }. Although the pathway node being navigated to is a child of the starting concept, the set of retained objects, R, now contains objects from starting concept, along with the two formal concepts that the user previously navigated to. For this example, R is the union of three sets of objects from the following formal concepts:

Table 3.6: An example of how child concepts are selected for the pathway node concept (A, B) in Fig. 3.10 where $B = \{$ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' $\}$. Selected semiconcepts are highlighted, excluded semiconcepts are striked out.

Pathway concept (A, B), **Retained objects** $|R| = 27$

$|A| = 25$

$B = \{$ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' $\}$

$R = R \cup A, |R| = 36$

List of formal concepts T with highest concept similarity score to (A, B).

For each concept (E, F) in T, form a semiconcept (X, Y) where

$X = E \setminus R$ and $Y = F$. Select semiconcepts where $|X| > 0$.

Y	$\lvert E \rvert$	$\lvert X \rvert$	Sim.
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' }~~	25	0	1.00
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	16	0	0.74
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.58
~~{ 'American', 'woman', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.58
~~{ 'American', 'child', 'children', 'photograph', 'Consuelo Kanaga', 'melancholy' }~~	10	0	0.56
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	27	0	0.56
~~{ 'American', 'portrait', 'photograph', 'Consuelo Kanaga', 'melancholy' }~~	6	0	0.54
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'close-up', 'melancholy' }~~	5	0	0.52
~~{ 'American', 'child', 'photograph', 'New York', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	9	0	0.49
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	15	0	0.48
{ 'American', 'photograph', 'New York', 'Consuelo Kanaga' }	16	3	0.47
{ 'American', 'photograph', 'mother', 'Consuelo Kanaga' }	18	9	0.47
~~{ 'American', 'child', 'children', 'photograph', 'mother', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	9	0	0.46
{ 'American', 'children', 'photograph', 'mother', 'Consuelo Kanaga' }	11	2	0.45

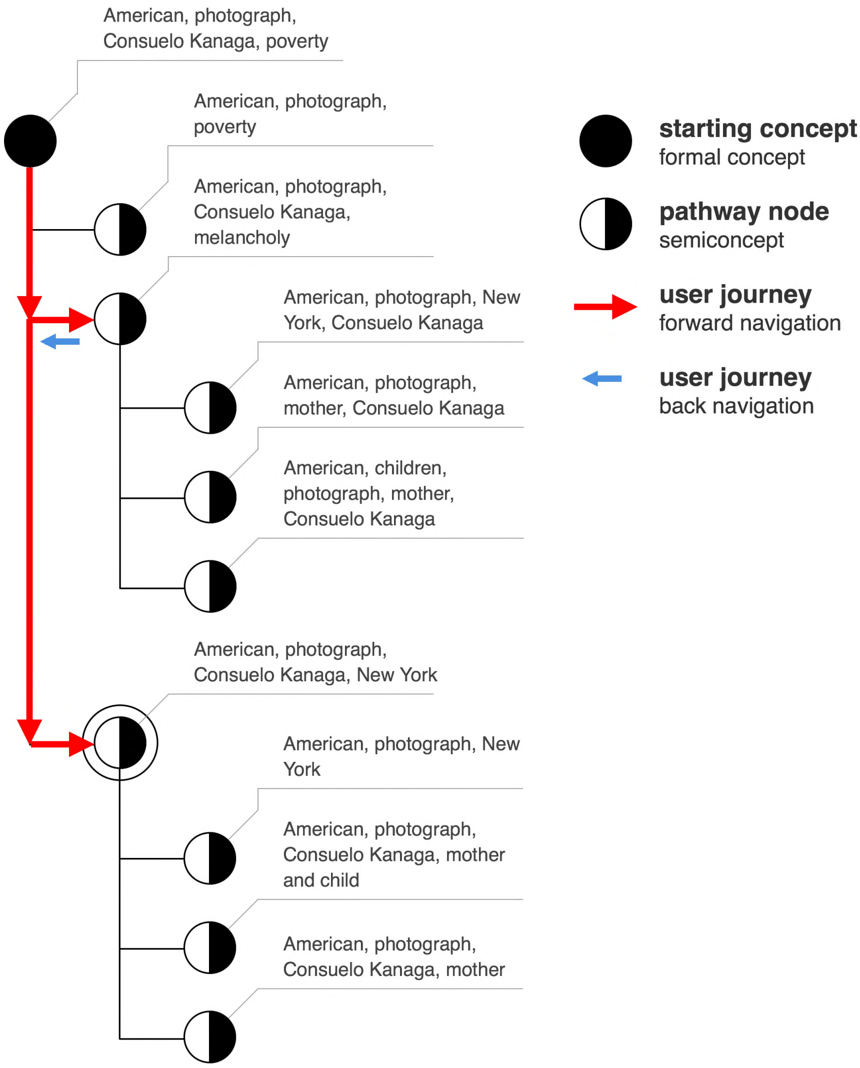

Figure 3.12: The expanded pathway structure showing a user backtracking and navigating to the formal concept with attributes { 'American', 'photograph', 'New York', 'Consuelo Kanaga" } with its expanded child concepts, generated via the selection of semiconcepts from Table 3.7.

{ { 157696, ... }{ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' } }
{ { 157696, ... }{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' } }
{ { 157696, ... }{ 'American', 'photograph', 'New York', 'Consuelo Kanaga' } }

The algorithm again searches for formal concepts that are most similar to the formal concept with attributes { 'American', 'photograph', 'New York', 'Consuelo Kanaga' }, retrieving only semiconcepts with objects not in R. The selection process is illustrated by example in Table 3.7, with its child concepts now attached to the pathway structure in Figure. 3.12.

Based on the example provided, a user navigating concept-linked pathways would begin at a *starting concept* – which can be randomly assigned or selected by the user – and then incrementally *diverge* as they would navigate to conceptually

Table 3.7: An example of how child concepts are selected for the pathway node concept (A, B) in Fig. 3.11 where $B = \{$ 'American', 'photograph', 'New York', 'Consuelo Kanaga" $\}$. Selected semiconcepts are highlighted, excluded semiconcepts are striked out.

Pathway concept (A, B), **Retained objects** $|R| = 36$

$|A| = 25$

$B = \{$ 'American', 'photograph', 'New York', 'Consuelo Kanaga" $\}$

$R = R \cup A, |R| = 39$

List of formal concepts T with highest concept similarity score to (A, B).

For each concept (E, F) in T, form a semiconcept (X, Y) where

$X = E \setminus R$ and $Y = F$. Select semiconcepts where $|X| > 0$.

Y	$\|E\|$	$\|X\|$	Sim.
~~{ 'American', 'photograph', 'New York', 'Consuelo Kanaga' }~~	16	0	1.00
~~{ 'American', 'photograph', 'New York', 'Consuelo Kanaga', 'poverty' }~~	13	0	0.82
{ 'American', 'photograph', 'New York' }	19	3	0.72
~~{ 'American', 'woman', 'photograph', 'New York', 'Consuelo Kanaga' }~~	7	0	0.64
~~{ 'American', 'child', 'photograph', 'New York', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	9	0	0.59
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	27	0	0.55
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'melancholy' }~~	12	0	0.52
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.50
~~{ 'American', 'child', 'photograph', 'Consuelo Kanaga', 'poverty' }~~	15	0	0.49
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	16	0	0.48
{ 'American', 'photograph', 'Consuelo Kanaga', 'mother and child' }	15	8	0.48
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy' }~~	25	0	0.47
{ 'American', 'photograph', 'mother', 'Consuelo Kanaga' }	18	9	0.46

similar pathways. These divergences are driven by the way the algorithm selects conceptually similar pathways that always present new objects to the user. While users are free to navigate to any node within the pathway structure, the concepts generated by new pathway nodes become increasingly distant from their previous points of navigation. In the context of presenting concept-linked pathways over a museum collection, the notion of presenting diverging, yet similar pathways could potentially encourage the user to explore concepts and objects that were previously unknown to them. In effect, a user's exploration of the pathway structure could expose them to new explicit descriptions and interpretations of its data, a trait that's commonly realised in *redescription mining*, *story telling* and *link analysis* applications [105, 106]. Whereas *concept-linked* pathways allow the user to diverge on a trail of similar concepts, a variation of this interaction approach – *object-linked pathways* – allows the user to diverge based on a specific object of interest.

3.4.2 Object-linked pathways

Object-linked pathways are formed based on the concept similarity of a formal concept's set of objects with other formal concepts of its formal context. Like concept-linked pathways, object-linked pathways begin with a *starting concept*, however, its child concepts are formed by creating formal concepts for each object within its object set. Rather than have a user select a child concept that is conceptually related to its parent, the user would select a specific *pivot object* on that concept, which would then display the formal concept that has the highest concept similarity to the pivot object's formal concept, given that the selected formal concept contains objects that are not already existing in formal concepts on the pathway structure. Although object-linked pathways provide the navigation and user interface for *A Place for Art* (Section 4.3), the Brooklyn Museum collection from the *Brooklyn Museum Canvas* case study will be used as an example to describe how object-linked pathways are constructed, and how they differ from concept-linked pathways.

An object-linked pathway begins with the starting concept (A, B) and an empty set of *retained objects* R, along with its list of child concepts generated for each object on that starting concept. Since the starting concept is the only formal concept currently on the pathway, $R = A$. For the starting concept (A, B), for each object $a \in A$, a semiconcept, called the *object concept*, is formed: $\{ \{a\}, a' \}$ [10] [11], which is a semiconcept that consists of the object and its set of its attributes. The resulting pathway structure is shown in Fig. 3.13. For the sake of clarity, a maximum of

[10] The *object concept* is technically a *semiconcept* and may differ from other definitions of an *object concept* used elsewhere within the Formal Concept Analysis literature. For the scope of this thesis, this definition will be retained.

[11] The prime operator a' on a is defined in Section 3.1.

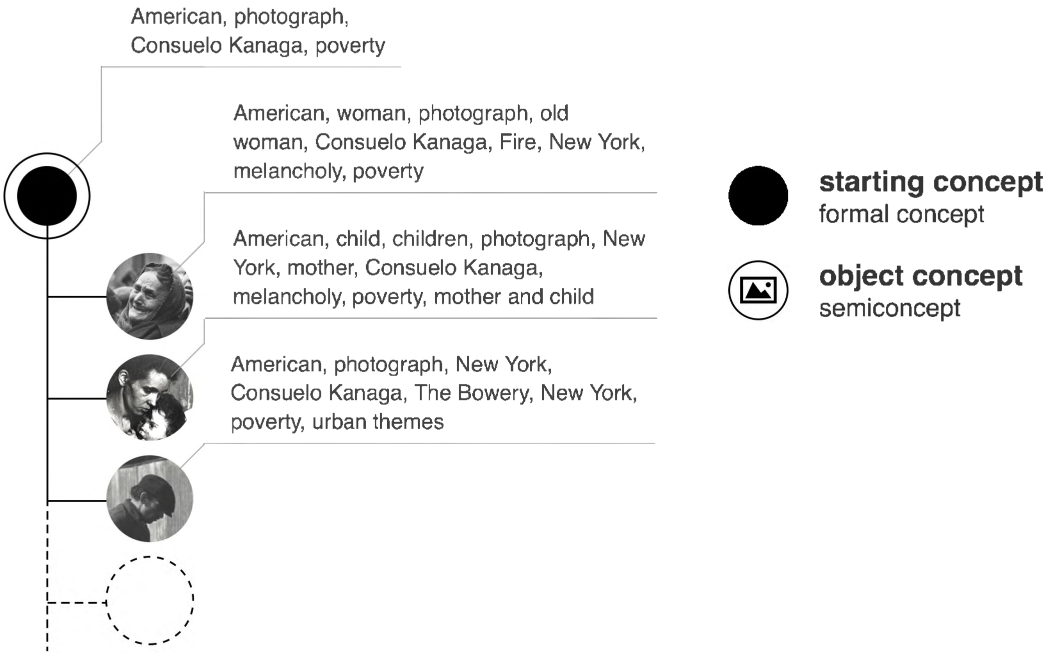

Figure 3.13: The starting concept of the object-linked pathway and its child concepts, each formed from an object within the starting concept.

three child concepts will be shown for each pathway node, although in practice the amount of child concepts per pathway node is usually far greater than three.

At this point, the user would select which object concept they would like to navigate to. From this *pivot object*, the next pathway node is then determined, which is the formal concept that has the highest similarity score to the pivot object that contains objects not in the set of *retained objects* R. The pseudocode for selecting the next pathway node is shown in Algorithm 2 and is almost identical to the GETCHILDCONCEPTS pseudocode in Algorithm 1, except that that the input is the object concept (O, P) and that it returns a single semiconcept (M, N), rather than an array of child concepts.

The selection process is illustrated by example in Table 3.8, with the resulting pathway structure shown in Fig. 3.14. The figure shows the selected pivot object with attributes { 'American', 'woman', 'photograph', 'old woman', 'Consuelo Kanaga', 'Fire', 'New York', 'melancholy', 'poverty' }; its closest matching pathway node – a formal concept with attributes { 'American', 'photograph', 'Consuelo Kanaga', 'woman', 'old woman' } – and its child concepts, a set of semiconcepts formed by the objects of the pathway node's formal concept (A, B), in which its objects are also added to the set of retained objects such that $R = R \cup A$. From this point, the user could traverse deeper into the pathway structure and select another pivot object which would reveal another formal concept closely related to that object, as demonstrated by example in Table 3.9 and Fig. 3.15, or backtrack and explore other concepts within the pathway structure.

Input: Object concept (O, P); the set of retained objects R; the set of all formal concepts S in formal context $K(G, M, I)$
Output: The pathway node, a semiconcept (M, N), for object concept (O, P)

```
R := R ∪ A;
T := [];
(M, N) := null;

foreach (E, F) in S do
    similarity((A, B), (E, F)) := 1/2 ((|A ∩ E|/|A ∪ E|) + (|B ∩ F|/|B ∪ F|));
    conceptAndSimilarityScore := ((E, F), similarity((A, B), (E, F)));
    Add conceptAndSimilarityScore to T;
end

Sort T by similarity((A, B), (E, F)) descending;

while (M, N) = null do
    (E, F) := shift first element from T;
    X := E \ R;
    Y := F;
    if |X| > 0 then
        (M, N) := (X, Y)
    end
end
```

Algorithm 2: The GetPathwayNodeForObjectConcept algorithm used for retrieving the pathway node for a given object concept (O, P) from a set of formal concepts S in formal context $K(G, M, I)$.

Table 3.8: An example of how the pathway node concept is selected for the pivot object (O, P) in Fig. 3.14. The selected semiconcept is highlighted, excluded semiconcepts are striked out.

Pivot object (O, P), **Retained objects** $|R| = 27$

$O = \{157696\}$

$P = \{$ 'American', 'woman', 'photograph', 'old woman', 'Consuelo Kanaga', 'Fire', 'New York', 'melancholy', 'poverty' $\}$

List of formal concepts T with highest concept similarity score to (O, P).

For each concept (E, F) in T, form a semiconcept (X, Y) where

$X = E \setminus R$ and $Y = F$.

Select semiconcept with highest similarity score where $|X| > 0$.

Y	$\|E\|$	$\|X\|$	Sim.
~~{ 'American', 'woman', 'photograph', 'old woman', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	5	0	0.41
~~{ 'American', 'woman', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	11	0	0.31
{ 'American', 'woman', 'photograph', 'old woman', 'Consuelo Kanaga' }	11	6	0.28
{ 'American', 'woman', 'photograph', 'Consuelo Kanaga', 'melancholy' }	15	4	0.26
~~{ 'American', 'photograph', 'Consuelo Kanaga', 'melancholy', 'poverty' }~~	16	0	0.26

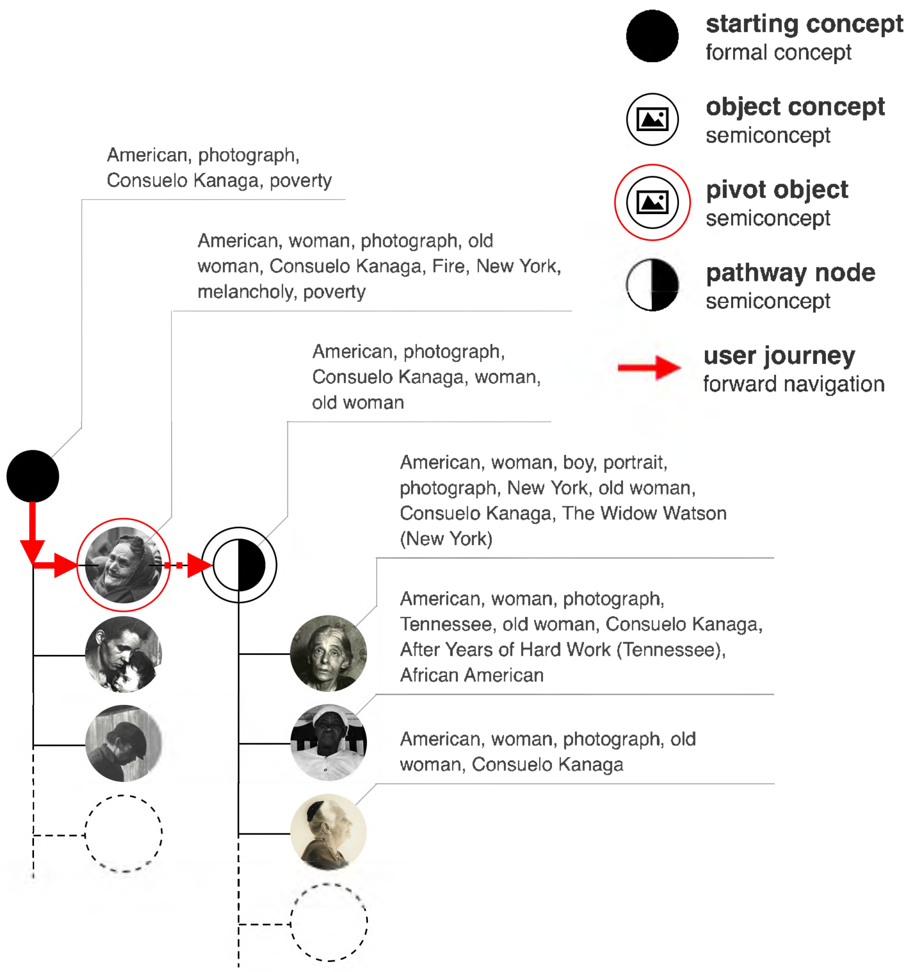

Figure 3.14: An expanded object-linked pathway with a selected pivot object and its associated pathway node from Table 3.8

Table 3.9: An example of how the pathway node concept is selected for the pivot object (O, P) in Fig. 3.15. The selected semiconcept is highlighted, excluded semiconcepts are striked out.

Pivot object (O, P), **Retained objects** $\lvert R\rvert = 33$			
$O = \{157744\}$			
$P = \{$ 'American', 'woman', 'photograph', 'Tennessee', 'African American', 'old woman', 'Consuelo Kanaga', 'After Years of Hard Work (Tennessee)' $\}$			
List of formal concepts T with highest concept similarity score to (O, P). For each concept (E, F) in T, form a semiconcept (X, Y) where $X = E \setminus R$ and $Y = F$. Select semiconcept with highest similarity score where $\lvert X\rvert > 0$.			
Y	$\lvert E\rvert$	$\lvert X\rvert$	Sim.
{ 'American', 'woman', 'photograph', 'Tennessee', 'Consuelo Kanaga' }	4	3	0.34
~~{ 'American', 'woman', 'photograph', 'old woman', 'Consuelo Kanaga' }~~	11	0	0.26
{ 'American', 'woman', 'photograph', 'Consuelo Kanaga', 'African American' }	15	14	0.25
{ 'American', 'photograph', 'Tennessee', 'Consuelo Kanaga', 'African American' }	17	16	0.25
{ 'American', 'photograph', 'old woman', 'Consuelo Kanaga' }	12	1	0.22

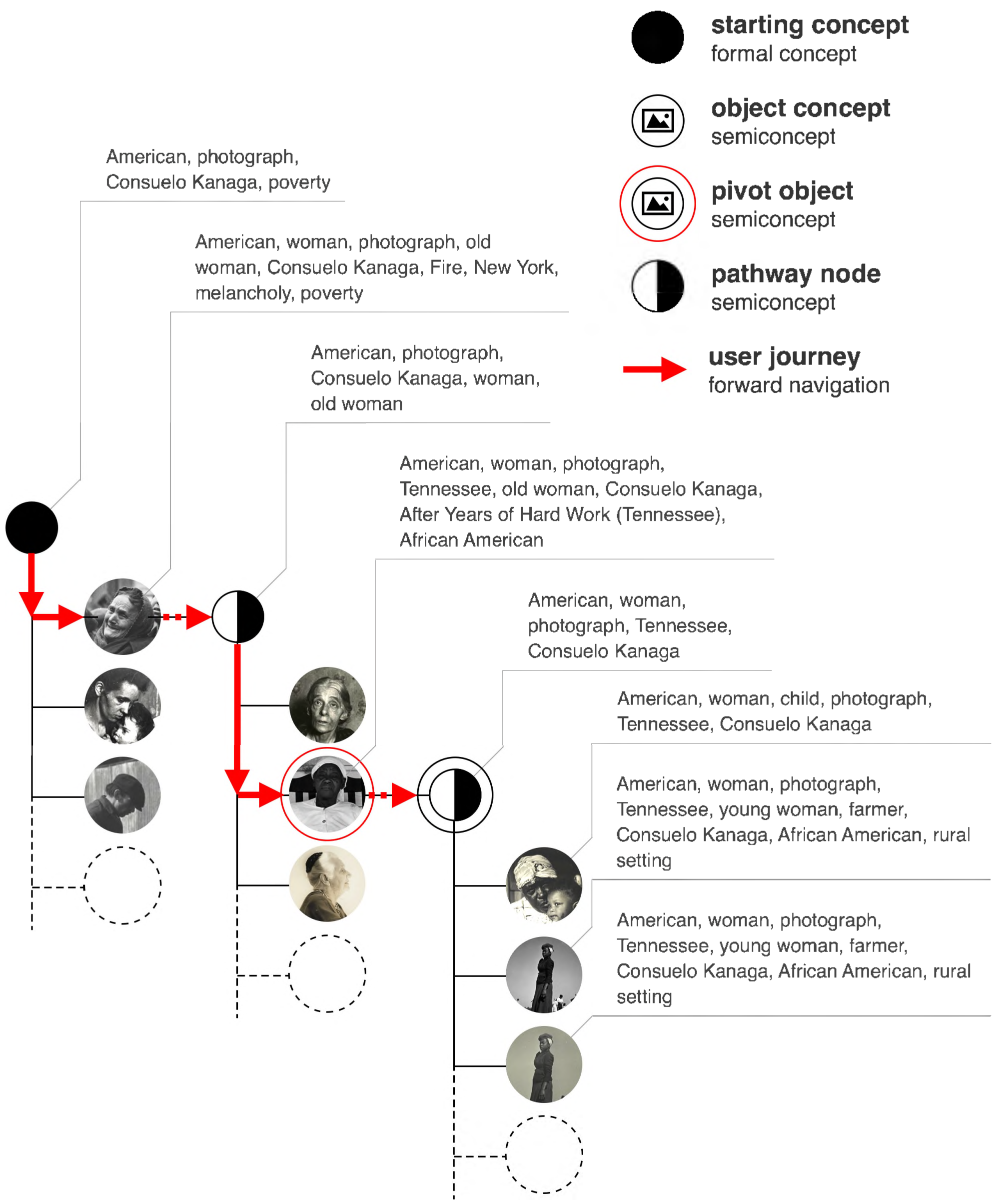

Figure 3.15: An expanded object-linked pathway with a selected pivot object and its associated pathway node from Table 3.9

Chapter 4

Case Studies

This chapter presents the *CollectionWeb* framework and three interactive visualisations: *The Virtual Museum of the Pacific*, *A Place for Art* and the *Brooklyn Museum Canvas*. Given that each section of this chapter describes the design, implementation, evaluation and context of these visualisations, they are presented as case studies.

Section 4.1 presents the *CollectionWeb* framework; its implementation of the *conceptual neighbourhood* and *concept layer* approaches of pathway-based navigation; its natural language processing and expression capabilities and a discussion on how the framework scales with large museum data-sets. The section also presents a performance evaluation of the algorithm that generates and maintains its *concept layer* (Section 4.1.3).

Section 4.2 presents the *Virtual Museum of the Pacific*, a rich Web-based browsing interface developed in collaboration with the Australian Museum that allows users to explore a collection of Pacific objects using the *conceptual neighbourhood* approach of pathway-based navigation. The study reports on the organisational context, fit and feasibility of the navigational approach; issues discussing museum-based information systems and museological perspectives on classifying objects, and user evaluations from museum staff, visitors, researchers and Pacific islanders.

Section 4.3 presents *A Place for Art*, an iPad app that allows users can explore a collection of paintings, works on paper and Australian Aboriginal works from the University of Wollongong Art Collection via the creation of *object-linked pathways*. The study reports on the app's interaction design, followed by results from a user experience evaluation where participants demonstrated exploratory behaviours and their personal connection to the works presented in the app.

Section 4.4 presents the *Brooklyn Museum Canvas* rich, web-based interactive visualisation that allows users to browse *concept-linked pathways* of works linked by artist, theme and visual imagery using objects harvested from the Brooklyn Museum API. The study reports on its design and its relationship to the principles of *land-*

scapes and *parallelism* along with an investigation on how pathways can represent a *narrative* of connected themes and visual imagery.

While all case studies present the *pathway* metaphor of interaction and hence demonstrate the generality of the concept, each case study individually differs in the size, scale and type of meta-data used; the the organisational and cultural context of each collection and for the *Virtual Museum of the Pacific* and *A Place for Art* case studies, the user communities, context, expectations and user experiences associated with each collection. All three interactive visualisations are available as working software and URLs to access them are available in the footnotes on the first page of each case study.

4.1 CollectionWeb

CollectionWeb is a framework that provides interfaces and Web Services for importing, extracting and navigating museum collection metadata using Formal Concept Analysis. The framework implements the necessary data structures required for pathway based navigation for all applications described in this thesis and also provides the algorithms for extracting and expressing them in natural language. In this thesis, *CollectionWeb* specifically refers to the Web Services framework and database layer that imports, extracts and stores museum metadata, whereas a '*CollectionWeb* application' or '*CollectionWeb* app' refers to any website or app that interfaces with it.

CollectionWeb is designed to be vendor and data-format agnostic can be best described as a data-store for museum collection metadata with the following processes and components:

1. An importer for museum metadata that uses a natural language processing pipeline to create a formal context of museum objects.

2. An implementation of FCA algorithms that expose the conceptual relationship among objects by a) creating and maintaining a *concept layer* that can be used to efficiently generate *concept pathways* and/or b) by dynamically computing *conceptual neighbourhoods*.

3. Another natural language processing pipeline that expressively describes formal concepts and the relations between them.

4. A toolkit and set of data structures that facilitates both the extraction of attributes from museum metadata in 1) and concept expression in 3) for any museum data-set.

Section 4.1.1 describes the data structures used by *CollectionWeb* to store museum metadata as a formal context, and Section 4.1.2 describes the natural language processing pipeline and toolkits that are used to extract attributes from museum metadata and express formal concepts in natural language. Section 4.1.3 describes the FCA algorithms used by the framework along with a performance evaluation of those algorithms as applied to museum metadata.

The other case studies in this chapter – *Virtual Museum of the Pacific*, *Brooklyn Museum Canvas* and *A Place for Art* – are all based on the principle of allowing a user to explore a pathway through a museum collection. These applications are all driven by the *CollectionWeb* framework, and this section demonstrates how the framework implements the interaction approaches of *conceptual neighbourhoods* and

concept pathways – both which adhere to the core design principle of *pathways as a mechanism that semantically structures content.*

4.1.1 Data structures

CollectionWeb provides, at its very core, a formal context for museum collection metadata. As mentioned previously in Section 3.1, a formal context is a triple $K := (G, M, I)$ where G describes the set of objects, M, the set of attributes and I the incidence relations between $\langle g, m\rangle \in I$ meaning that object $g \in G$ has attribute $m \in M$. Within *CollectionWeb*, each object g is represented in as an instance of a `Object`, and each attribute m is represented as an instance of `Attribute`. The formal context is stored as a table that references the incidence relations between `Objects` and `Attributes`. *CollectionWeb* also provides additional data structures that add a semantic layer to the `Attributes` that augment the expressiveness of formal concepts to the end user.

An `Object` instance only records a unique identifier that corresponds to an ID in a formal context, along with the museum's identifier for that object. In *CollectionWeb*, an `Object` can be extended to incorporate additional metadata fields that may be used in *CollectionWeb* applications. For instance, *A Place for Art* (Section 4.3) and *The Brooklyn Museum Canvas* (Section 4.4) store museum metadata using instances of `APlaceForArtObject` and a `BrooklynMuseumObject` – both of which inherit from `Object`. In a *CollectionWeb* application, it is possible to create many descendant classes of `Objects` in a single application, allowing the ability for a single app to display and process museum metadata for multiple sources, although the case studies in this thesis only show apps that store and show museum metadata from a single source.

The following is a partial JSON representation of a single `BrooklynMuseumObject` that shows the metadata associated with a work by photographer Consuelo Kanaga shown in Fig. 4.1. [1]


```
{
   "id": "157697",
   "name": "82.65.413",
   "uri": "http://www.brooklynmuseum.org/opencollection/ ... ",
   "title": "[Untitled] (Mother with Children, New York)",
   "objectDate": "1922-1924",
   "objectDateBegin": "1922",
   "objectDateEnd": "1924",
   "medium": "Gelatin silver photograph (from glass plate negative)",
   "description": "Portrait-group Condition: good",
   ...
   "attributeIDs": [8, 10, 107, 333, 974, 2808, ... ]
}
```


Of note are the two identifier fields associated with the `BrooklynMuseumObject`: the `name` field which is a unique identifier that the museum uses to identify the object, and the `id` field, which is an internal identifier that *CollectionWeb* uses to identify the object and reference it within its formal context. Storing these two separate identifiers allows *CollectionWeb* to store a reference to the museum's authority on the object without the possibility of conflicting with any other identifiers used by objects from other museums within a single *CollectionWeb* app.

`Attributes`, like `Objects` have a unique identifier that corresponds to a formal attribute within a formal context, along with a `value` property. The following is a list of attribute identifiers from the above example object record:

8	974
10	2808
107	2845
333	5643
40349	34856
40344	40348

[1]The original record for this object can be accessed at: `http://www.brooklynmuseum.org/opencollection/objects/157697/%7CUntitled%7C_Mother_with_Children_New_York`

Figure 4.1: *Untitled (Mother with Children, New York)* by American photographer Consuelo Kanaga, the image associated with a single `BrooklynMuseumObject` record. Image obtained from `http://www.brooklynmuseum.org/opencollection/objects/157697/%7CUntitled%7C_Mother_with_Children_New_York/`

Dereferencing these attribute identifiers yields the following attribute values:

`20th century`	`children`
`American`	`photograph`
`child`	`New York`
`1920s`	`mother`
`mother and child`	`Consuelo Kanaga`
`melancholy`	`poverty`

These attribute values can be used to describe the object in natural language:

```
An object associated with 20th Century, American, children, 1920s,
photograph, New York, mother, Consuelo Kanaga, melancholy, poverty,
mother and child
```

In *CollectionWeb*, this description can be formed by constructing a `FormalConcept` object based on its `Object` and `Attribute` identifiers:

```
// Usage:       new FormalConcept(objects[], attributes[]);
formalConcept = new FormalConcept([157697], [8, 10, 107, ... ]);
```

And then calling its `getNaturalLanguageExpression()` method:

```
// Retrieves a natural language description of the formal concept
print formalConcept->getNaturalLanguageExpression();

/*
Outputs: An object associated with 20th Century, American, children,
1920s, photograph, New York, mother, Consuelo Kanaga, melancholy,
poverty, mother and child
 */
```

By loading the collection's `FormalContext`, and then instantiating that same formal concept to then create a `FormalConceptInContext`, *CollectionWeb* can compute its upper and lower neighbours:

```
// Loads the entire formal context of the collection
K = FormalContext::load();

// Assign the object ID and its attributes as variables
A = [157697];
B = [8, 10, 107, 333, 974, 2808, 2845, 5643, ... ];

// Creates a new formal concept
formalConcept = new FormalConceptInContext(K, A, B);

// Retrieves its upper neighbours
upperNeighbours = formalConcept->getUpperNeighbours();

// Display its upper neighbours
foreach (upperNeighbour in upperNeighbours) {
   print upperNeighbour->getNaturalLanguageExpression();
}
```

The above code outputs the following:

```
objects associated with 20th century, 1920s, photograph, mother

objects associated with 20th century, American, child, 1920s,
children, photograph, Consuelo Kanaga

objects associated with 20th century, American, child, children,
1920s, photograph, New York, Consuelo Kanaga, melancholy,
poverty

objects associated with 20th century, American, child, 1920s,
photograph, New York, Consuelo Kanaga, melancholy, poverty

objects associated with 20th century, American, child, children,
photograph, Consuelo Kanaga, mother and child
```

These natural language statements only list the attribute values of the formal concept. While accurate in their descriptions, they offer relatively little expression in terms of how the objects are grouped and related to one another.

4.1.2 Natural language processing and expression

One of the design goals of *CollectionWeb* aims to utilise machine learning and human computer interaction to augment understanding of the knowledge embedded within museum collections – this is what Horvitz refers to as *mixed initiative.* [107] *CollectionWeb* applications can express concepts in natural language that allow for easier comprehension of the conceptual structures that represent the relationships among the objects in the collection. While enhancing usability, the design goal of fluid, narrative like expression of formal concepts using natural language also aims to aide the key tasks of *identification* and *recognition* in Wille's key tasks of Conceptual Knowledge Processing [96]. This section describes the data structures and processes used to generate natural language statements for formal concepts, followed by an example of how they can present a user journey as a *narrative* through a pathway of formal concepts generated from the Brooklyn Museum's online collection.

CollectionWeb uses `Predicates` to add an additional semantic layer to attributes and `Parsers` to express these semantics in natural language. A `Predicate` can be thought of as a qualifier that characterises the relationship between the `Attribute` and the `Object`. For example, the work shown in Fig. 4.1 is `byArtist` *Consuelo Kanaga.* It is a `typeOf` *photograph* that `depicts:location` *New York.* The following is a list of all attributes of the *Untitled* work by Consuelo Kanaga from Section 4.1.1, now supplemented with `Predicates`:

```
fromTimePeriod:20th century
artistNationality:American
depicts:person-subject:child
fromTimePeriod:1920s
depicts:person-subject:mother and child
depicts:emotional-subject:melancholy
depicts:person-subject:children
isTypeOf:photograph
depicts:location:New York
depicts:person-subject:mother
byArtist:Consuelo Kanaga
depicts:emotional-subject:poverty
```

Although superficially similar to the `Subject-Predicate-Object` relationship found in the RDF data model, *CollectionWeb* does not intend to aim to replicate, nor surpass, the expressive power of the semantics found in RDF or similar metadata data models. The `Attribute-Predicate-Object` model used in *CollectionWeb* is

deliberately kept simple to avoid the complexity of constructing de-normalised formal contexts from semantic data.

`Predicates` exist independently of `Attributes`: an `Attribute` may be assigned to one `Predicate` but a `Predicate` may be assigned to multiple `Attributes`. *CollectionWeb* does not map explicit relationships between `Predicates`, nor does it allow more than one `Predicate` to be assigned to a single `Attribute` instance. In such cases, if a single `Attribute` is assigned to the same `Object` but with two different `Predicates` – e.g. an `Object` may be `associatedWithAgent:Consuelo Kanaga` and/or created `byArtist:Consuelo Kanaga`, *CollectionWeb* treats both occurrences as two separate `Attributes`, each being potentially assignable to the object.

`Predicates` can be used to determine how an `Attribute`, or a set of `Attributes` sharing a common root `Predicate`, are expressed in natural language. For example, the following attributes:

```
fromTimePeriod:20th century
fromTimePeriod:1920s
fromTimePeriod:1930s
fromTimePeriod:1940s
```

If assigned as `Attributes` to a `FormalConcept`, will produce the following result for its `getNaturalLanguageExpression()` method:

```
works from the 1920s to the 1940s
```

This is because the above attributes were assigned to the `fromTimePeriod` predicate, one of *CollectionWeb's* built-in predicates that can be used to express common metadata fields found within museum collections. Each built-in `Predicate` is linked to a `Parser`. The `Predicate` describes the semantic relationship between an object and an attribute and contains simple natural language processing rules and templates, whereas its `Parser` contains *procedural* functionality on how a certain class of attributes should be mined from its free text field and subsequently expressed in natural language. A `Parser` can be thought of as a pipeline that acts an intermediary between the free-textual data of a meta-field and synthesized natural language statement of any formal concepts that arise from the attributes that are mined from the data. In this case, the `fromTimePeriod` predicate uses `TimePeriodParser` to express a collection of `fromTimePeriod` attributes as an individual decade, a date range or a century depending on the attributes present. Likewise, the parser can also be used to extract such dates from free text metadata from a collection object that expresses a date or date range. *CollectionWeb* contains several built-in

`Predicates` – each with their own `Parser` – that allow natural language extraction and expression of common metadata fields typically found in museum collections – such as date ranges, artwork series, mediums, keywords and titles. The parsers use the Illinois Part of Speech Tagger [2] [108] to recognise nouns and noun phrases, and the Stanford Named-Entity Recogniser [3] [109] to identify people, places and organisations.

The other `Predicates` built into the *CollectionWeb* framework are as follows:

The isTypeOf Predicate The `isTypeOf` predicate uses the `ObjectNounParser` for extracting and displaying textual data that describes actual object types. For example, attribute values that represent category types and classes such as `Contemporary works` or `paintings` would use the `isTypeOf` predicate.

The `ObjectNounParser` ensures that object nouns are represented in their correct form when expressed in singular or plural forms. When parsed from textual data for mining and database storage, it ensures that they are represented as collective nouns (e.g. `warfare objects` instead of just `warfare`) and that they are both expressed correctly in singular and plural form (e.g., an object with the text snippet "painting" will be expressed as `painting` if the formal concept contains one object or `paintings` if the formal concept contains more than one object.)

As an example of how the `isTypeOf` predicate expresses attributes, the following attributes:

```
isTypeOf:photograph
isTypeOf:drawing
isTypeOf:painting
```

If assigned as `Attributes` to a `FormalConcept`, will produce the following result for its `getNaturalLanguageExpression()` method:

```
paintings, drawings and photographs
```

The is Predicate The `is` predicate uses the `ObjectDescriptionParser` to parse descriptive elements and non-iconographical concepts from textual data that succinctly describe a literal depiction of an object by extracting noun-phrases. It also implements a custom method for expressing those entities in a meaningful way.

The `ObjectDescriptionParser` works particularly well if the artwork's title is

[2] http://cogcomp.cs.illinois.edu/page/download_view/POS

[3] http://nlp.stanford.edu/software/CRF-NER.shtml

used as a data source, although very short snippets describing the description of the work can be used as well. The parser itself does not distinguish the level of iconography expressed within the source text, and it assumes that all text fed into the parser describes the literal description of the work, rather than its iconography.

As an example of how the `is` predicate expresses attributes, the following attributes:

```
is:cover
is:box
```

If assigned as `Attributes` to a `FormalConcept`, will produce the following result for its `getNaturalLanguageExpression()` method:

```
objects that consist of boxes and covers
```

The depicts Predicate The `depicts` predicate uses the `ImageSubjectParser` to parse entities and iconographical concepts from short snippets of textual data that succinctly describe an image-based artwork by extracting noun-phrases and named-entities. The parser recognises and distinguishes the difference between iconographical concepts that represent objects, people or places in a work of art, and implements custom methods for expressing those entities in a meaningful way.

The `ImageSubjectParser` works particularly well if the artwork's title is used as a data source, although very short snippets describing the iconography of the work can be used as well. The parser does not distinguish the level of iconography expressed within the source data, and it assumes that all text fed into the parser describes the "aboutness" of the artwork, (i.e., Panofsky's [110] first or second level of image iconography) [4] rather than a literal description of the object.

Although it can theoretically accept any data, this parser is designed to work on artwork titles for two-dimensional visual works, such as paintings, prints, photographs or other 'flat' visual media where the title evokes a depiction, rather than just stating what the work is. Depending on how the artist or creator titles a work, the parser could also accept titles of works in other mediums, such as sculptures or 3-dimensional works.

As an example of how the `depicts` predicate expresses attributes, the following

[4] Panofsky's three levels of image iconography refer to three strata of image subject matter of meaning: the first level refers to the primary subject matter of an image or a work of art, the 'literal' presentation of the work; the second level incorporates iconography and adds cultural context to the imagery in the work; and the third level incorporates iconology, referring to the intrinsic meaning or interpretation of a work: its personal, technical, historical or cultural understanding.

attributes:

```
depicts:person-subject:child
depicts:person-subject:children
depicts:location:New York
depicts:emotional-subject:poverty
```

If assigned as `Attributes` to a `FormalConcept`, will produce the following result for its `getNaturalLanguageExpression()` method:

```
works that depict children in poverty in New York
```

CollectionWeb also contains other general-purpose `Parsers` that are not tied to any specific `Predicate` and can be used to extract data and generate attributes from any metadata field:

- The `StopWordsRemovalParser` can be used to tokenise a string and removes punctuation, numbers and stopwords (such as `and`, `for`, `of`, etc.)
- The `NounPhraseParser` extracts head nouns and noun phrases from free-text fields, useful for parsing text fields that depict a work's medium or title. For instance, the free text that describes a work's medium – `Woodblock color print with silver pigment` – parses as attributes with values { `'woodblock color print'`, `'silver pigment'`, `'print'`, `'pigment'` }. Likewise, the free text that describes a work's title `Sculptor's Model of Male Body Wearing Long Skirt` returns attribute values { `'sculptor's model'`, `'male body'`, `'long skirt'`, `'model'`, `'body'`, `'skirt'` }.

`Predicates` and `Parsers`, when combined, can be used to build complete natural language descriptions of any object or formal concept. *CollectionWeb* provides the ability to specify and configure how the natural language phrasing is expressed via configurable JSON files. These files store what *CollectionWeb* refers to as a *predicate map*, which is a two-dimensional array of `Predicates` where each predicate describes a natural language phrase fragment, such as `made with steel and wood` or `depicts farmhouses and animals`. In *CollectionWeb*, the predicate map can be used to:

- Map the associations between natural language phrase fragments and `Predicates` that in turn map to the metadata fields within the museum collection.
- Describe the order that these phrase fragments occur in, and prioritizes which phrase fragments should be displayed for a given position within the phrase.

- Describe, in a declarative manner, how each phrase fragment is displayed for a given set of attribute values.

Fig. 4.2 shows how a predicate map generates the natural language phrase `photographs by Consuelo Kanaga from the 1920s made with gelatin silver that depict poverty in New York` from the attributes of its formal concept. The predicate map is shown in the centre of the figure, where the horizontal ordering of its columns corresponds with the ordering of the phrase fragments within the natural language statement, and the ordering of each row item corresponds to the priority of predicates displayed. For example, as the first column lists the predicates `is` and `isTypeOf`, *CollectionWeb* will first search for any attributes within the formal concept with the predicate `is`, and if it finds none, it then searches for any attributes with the predicate `isTypeOf`. The algorithm then produces a phrase fragment for the attribute `isTypeOf:photograph` using the `isTypeOf` predicate, which then forms the first part of the phrase.

The algorithm then moves to the second column as it searches for any attributes with the `byArtist` predicate. As there is such an attribute present – `byArtist:Consuelo Kanaga` – it renders its phrase fragment and subsequently ignores any attributes with the `artistNationality` or `fromLocation` predicates. Such ordering could be used to denote an implicit hierarchy of concepts as expressed in natural language without producing excessively long natural language descriptions. As an example, if the attribute `byArtist:Consuelo Kanaga` and `depicts:emotional-subject:poverty` were to be 'removed' from the concept by means of navigating to its upper neighbour, the generated natural language statement would instead read as `photographs by American artists from the 1920s made with gelatin silver that depict New York`.

In the following column, another attribute is omitted from presentation – `fromTimePeriod:20th century` – but this time, for a different reason. As highlighted previously, the `fromTimePeriod` predicate uses the `TimePeriodParser` that omits redundant attributes that describe a century, such as `20th century` if more specific date attributes are present. After the column has been processed, the algorithm continues left-to-right within the predicate map, forming the complete natural language statement of the formal concept.

As mentioned previously, the predicates `is`, `isTypeOf`, `fromTimePeriod` and `depicts` are built in to the *CollectionWeb* framework, each with a `Parser` that dictates how a phrase fragment should be expressed for a given set of attribute values. However, the predicates `byArtist`, `artistNationality`, `fromLocation` and `hasMedium` are not built-in – instead, they are provided in this example as custom predicates. Custom predicates can be defined to express any metadata field on an object within a museum collection. They use an additional field that dic-

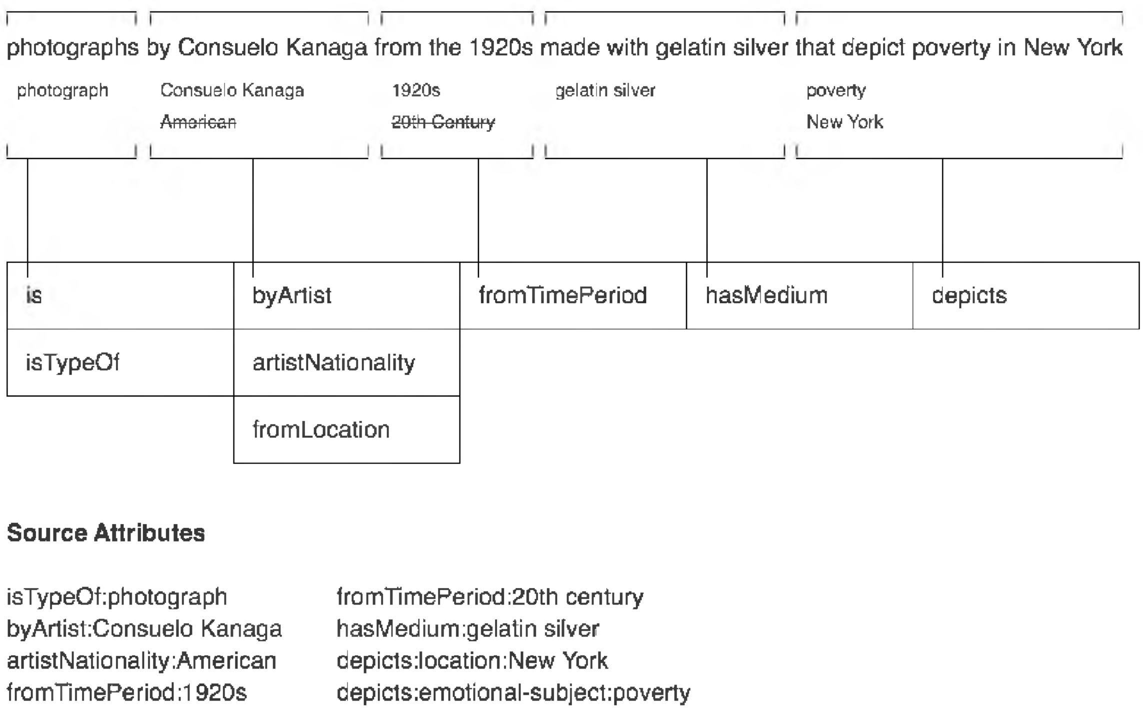

Figure 4.2: An example of how a formal concept can be expressed in natural language from its source attributes using a predicate map.

tates how that predicate is displayed as a phrase fragment, prepending any necessary prepositions, suffixes or prefixes so that its attribute values are inserted into the phrase. For example, the custom predicate `hasMedium` uses `made with {attribute}` so that the attribute `hasMaterial:gelatin silver` appears as `made with gelatin silver`. Likewise, custom predicate `artistNationality` uses `by {attribute} artists` so that attribute `artistNationality:American` appears as `by American artists`. For custom predicates that require more detailed logic or `Parsers` for expressing phrase fragments, it is possible to create custom classes that inherit from `PredicateParser` and link those to the custom predicate.

Using `Predicates` and `Parsers` – including some that are already built into the framework – the following example shows the list of neighbouring concepts shown previously in Section 4.1.1, but with the natural language expression rules applied:

```
/*
List of Upper Neighbours for formal concept:

A photograph by Consuelo Kanaga from the 1920s that
depicts children in poverty and melancholy in New York
*/

other photographs from the 1920s that depict mothers

other photographs by Consuelo Kanaga from the 1920s
that depict children

other photographs by Consuelo Kanaga from the 20th
century that depict mothers and children in poverty
and melancholy

other photographs by Consuelo Kanaga from the 1920s
that depict children in poverty and melancholy in
New York

other photographs by Consuelo Kanaga from the 20th
century that depict children

...
```

The above listing shows concepts that are far more expressive than the ones shown in Section 4.1.1 – in this case, the upper neighbours of the formal concept represented by the *Untitled* work in Fig. 4.1 link to Kanaga's other works that depict mothers, children and poverty from a similar time period.

Formal concepts are inherently spatial and relational: as generated from a formal context, formal concepts can be related to one another in terms of its upper or lower neighbours within the concept lattice (Section 3.2) or as having a certain degree of *concept similarity* to other concepts within the lattice (Section 3.3). In addition to describing concepts individually, *CollectionWeb* also expresses *relations* between two formal concepts in natural language. Expressing the conceptual relationships among formal concepts – as well as their content – allows *CollectionWeb* to describe not only concepts or groups of objects in their isolation, but also describe a chain of concepts as they would appear as part of user's navigation pathway as they would navigate among and across formal concepts.

When exploring neighbouring formal concepts, if a user moves to a formal concept's *subconcept* or *lower neighbour*, they navigate to a more specific formal concept. Likewise, if a user moves to a formal concept's *superconcept* or *upper neighbour*, they navigate to a more general formal concept. The *predicate map* – with its ability to prioritise which attributes are displayed for selected parts of a natural language phrase – can be used to succinctly highlight concepts that are more general or more specific to one another, where attributes that describe more specific concepts (such as the name of an artist) may be displayed or expressed in place of attributes that describe more broader concepts (such as the nationality of an artist). In the example predicate map shown in Fig. 4.2, the display of the more specific artist attribute also implies their nationality – the relevant domain knowledge either being known to the user or specifically expressed once the user navigates to a broader superconcept. As an illustration of this example, a formal concept with the following attributes:

```
fromTimePeriod:20th century
artistNationality:American
fromTimePeriod:1920s
isTypeOf:photograph
byArtist:Consuelo Kanaga
```

Will generate the following natural language statement:

```
photographs by Consuelo Kanaga from the 1920s
```

Note that the the attribute value of `artistNationality:American` is not displayed in the statement due to it being 'overridden' by the `byArtist:Consuelo Kanaga` attribute as shown in the predicate map in Fig. 4.2, and the `fromTimePeriod: 20th century` attribute is not shown due to the way the `TimePeriodParser` omits century attributes if more specific date attributes are present. In this way, the predicate map can be used to imply a hierarchy of attribute values that have different predicates but share a common semantic class. For example, given that `byArtist` attributes are more specific than `artistNationality` attributes, and that `designedIn` attributes are more specific than `associatedWithLocation` attributes, only the more specific `byArtist` and `designedIn` attributes would be displayed – even with the presence of `artistNationality` and `associatedWithLocation` attributes.

If, following the same rules and predicate map as above, the attributes `fromTimePeriod:1920s` and `byArtist:Consuelo Kanaga` were to be 'removed' from the formal concept – for example by means of navigating to a formal concept's upper

neighbour – then the generated natural language statement would appear as follows:

```
photographs by American artists from the 20th century
```

The expression of the broader terms – `American artists` and `20th century`, implies that the concept is broader. Hence, to a user, navigating between the two concepts would appear as follows:

```
photographs by Consuelo Kanaga from the 1920s =>
more photographs by American artists from the 20th century
```

CollectionWeb prepends the word `more` to the concept that the user is navigating to if an attribute with an `is` or `isTypeOf` predicate is present, or `other` if those attributes are not present, as shown in the following example where the user would then navigate to another upper neighbour that does not have the `isTypeOf:photograph` attribute present:

```
photographs by Consuelo Kanaga from the 1920s =>
more photographs by American artists from the 20th century =>
other works by American artists from the 20th century
```

If a user navigates between two concepts that do not have a subconcept – superconcept relationship, such as when navigation occurs on a *pathway* (Section 3.4), *CollectionWeb* only displays the attributes in the concept that the user is navigating to that are not present in the concept that they are navigating from. For example, consider the two attribute sets, each forming formal concepts that are rather similar to one another:

```
artistNationality:American        =>   artistNationality:American
isTypeOf:photograph                    isTypeOf:photograph
byArtist:Consuelo Kanaga               byArtist:Consuelo Kanaga
depicts:emotional-subject:poverty      depicts:melancholy
                                       depicts:person:mother
```

The new attributes introduced by the second formal concept `depicts:emotional-subject:melancholy` and `depicts:person:mother and child` are the only ones displayed when the user navigates to it, as shown in the example below:

```
photographs by Consuelo Kanaga that depict poverty =>
similar photographs that depict mothers in melancholy
```

Note that attributes with `is` or `isTypeOf` predicates are always displayed. In continuing with the running example, the user would then navigate to an upper neighbour of the latter formal concept, removing the attribute `depicts:emotional-subject:melancholy` so that its upper neighbour is displayed.

```
photographs by Consuelo Kanaga that depict poverty =>
similar photographs that depict mothers in melancholy =>
more works by Consuelo Kanaga that also depict mothers
```

These natural language expressions can be used to describe a user journey through a concept pathway. A concept pathway, as described in Section 3.4, is a tree structure that dynamically forms as a user navigates between similar formal concepts. Since *CollectionWeb* expresses the difference between formal concepts that are broader or narrower than one another, along with highlighting the differences between concepts that are closely similar, the framework can use its natural language capabilities to describe a user's experience through a pathway as a *narrative*. To demonstrate this, the example that follows shows the concept-linked pathways from Section 3.4.1 but with generated natural language statements from the *CollectionWeb* framework. The example begins with the starting concept, a selection of photographs by Consuelo Kanaga that depict people in poverty, as shown in Fig. 3.9.

Figure 4.3: The concept pathway from Fig. 3.9 showing a list of its attributes { 'American', 'photograph', 'Consuelo Kanaga', 'poverty' } and a generated natural language expression "photographs by Consuelo Kanaga that depict people in poverty".

Figure. 4.4 shows its child concepts, demonstrating that from this point the user could navigate to similar works that depict people in melancholy or works that depict the city of New York, or explore more photographs by American artists that depict people in poverty. As described in Section 3.4.1, these selections are based on the formal concepts that have been identified as most similar to these group

of photographs, given that they contain new objects that are not on the starting concept.

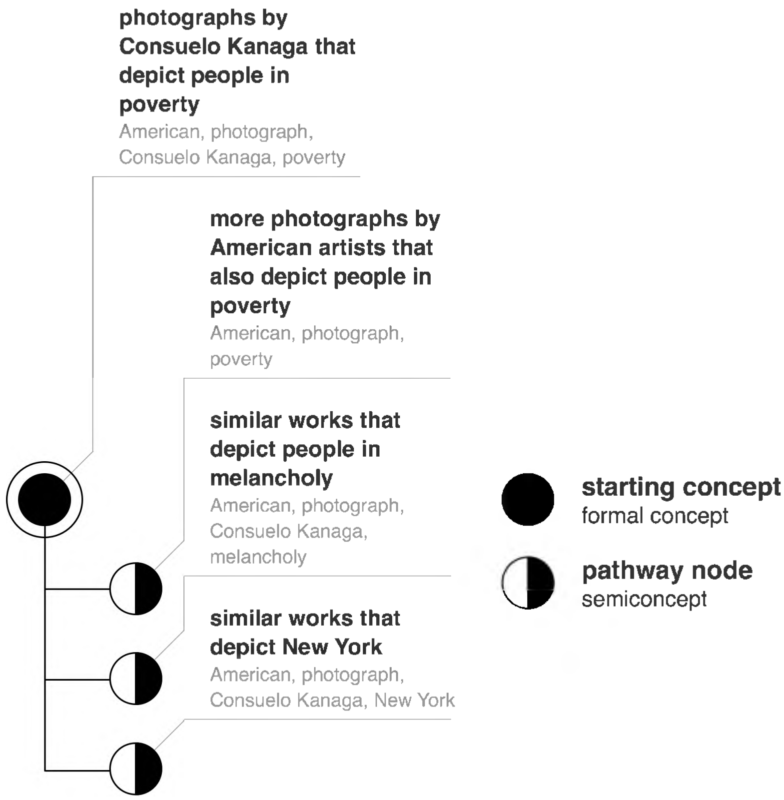

Figure 4.4: The start of a concept pathway from Fig. 3.10 with generated natural language expressions, highlighting the similarities and differences between the starting concept and its child concepts.

In Fig. 4.5, the user then navigates to the second formal concept, where the concept-linked pathway reveals child concepts that are thematically similar to "works that depict people in melancholy", but are further divergent from the original set of photographs by Consuelo Kanaga. Fig. 4.6 shows the user backtracking to the first pathway "photographs by Consuelo Kanaga that depict people in poverty" and selecting its third child concept – "similar works that depict New York" – and from that concept, its first child concept: "more photographs by American artists that also depict New York". The pathway navigation presented in Fig. 4.6 forms a narrative, shown in Fig. 4.7, that describe the user's navigation journey throughout the collection.

The ability for *CollectionWeb* to sufficiently convey these rich conceptual and expressive relationships relies on its ability to apply concept lattices in a scalable manner on sizes typically associated with museum data-sets. The section that follows describes the FCA algorithms used to form conceptual neighbourhoods and concept pathways, along with a user evaluation of those algorithms as applied to real museum data-sets.

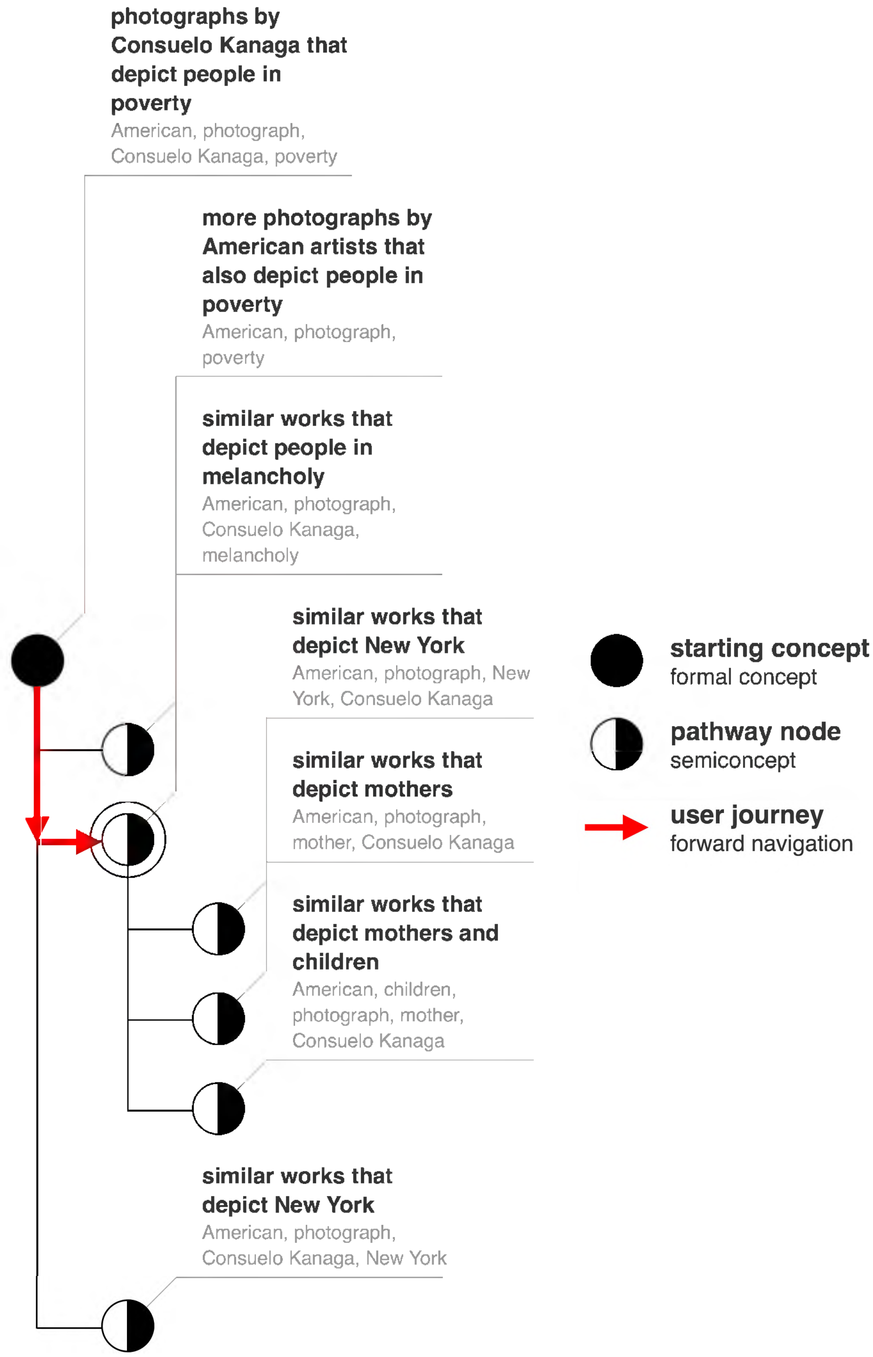

Figure 4.5: The now expanded pathway structure from Fig. 3.11 showing a user navigating from "photographs by Consuelo Kanaga that depict people in poverty" to "similar works that depict people in melancholy", revealing more divergent concepts for navigation.

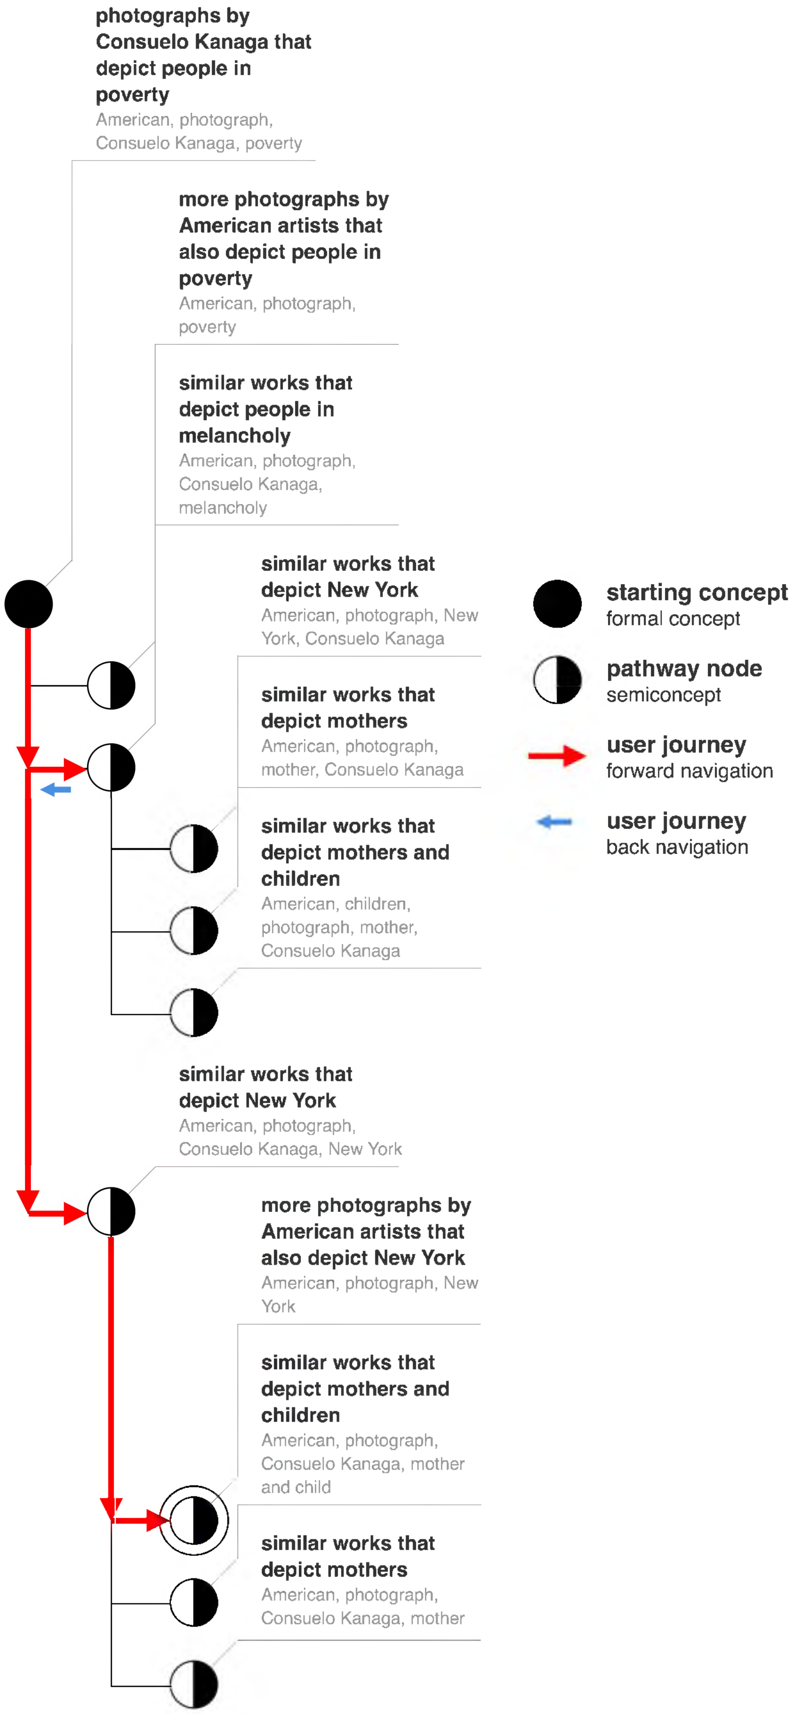

Figure 4.6: The expanded pathway structure from Fig. 3.12 depicting forward and backward navigations through a tree of thematically related concepts.

Figure 4.7: A narrative of pathway navigations from Fig. 4.6.

4.1.3 FCA algorithms: scalability and performance evaluation

CollectionWeb uses Formal Concept Analysis to create and maintain the associations between its objects. Rather than generate a concept lattice to visualise an entire collection, the framework generates and maintains relationships between formal concepts in two ways. The first approach – the *conceptual neighbourhood* approach – dynamically computes the conceptual neighbourhood of a formal concept. This a user to view a formal concept and navigate incrementally to its upper and lower neighbours as demonstrated previously in *ImageSleuth* (Section 3.2) and subsequently the *Virtual Museum of the Pacific* (Section 4.2) . The second approach – *the concept layer* approach – computes, stores and maintains a set of formal concepts for a dynamically changing formal context in which concept similarity and distance metrics can then be applied to generate pathways of related content, as it does in the *The Brooklyn Museum Canvas* (Section 4.4), and the *A Place for Art* (Section 4.3) case studies. This section reports on the FCA algorithms used for each approach, along with the scalability limitations and performance evaluations of each.

The conceptual neighbourhood approach

In all case studies reported in this thesis, Formal Concept Analysis is used to provide conceptual structures that can be navigated by the user. The conceptual neighbourhood approach, as reported previously in Section 3.2, offers the ability for a user to view individual concepts and move between neighbouring concepts within a concept lattice. A logical implementation of this approach would be to compute and store a complete concept lattice that can then be traversed by the user. However, complete concept lattices – while adequate for visualising relatively small data-sets – are both computationally prohibitive to produce and visually complex to interpret on larger data-sets typically associated with museum collections [93].

The time and space complexities of pre-computing and storing a complete con-

cept lattice are best understood through a discussion of how the approach scales with respect to the size of a formal context. Following an analysis of algorithms that build complete concept lattices, Carpineto and Romano [93] have identified their time complexities: the best result being the ConceptsCover algorithm which has a worst-case time complexity of $O(|C||M|(|G|+|M|))$, [5] which is dependent, in part, on the number of formal concepts generated from a formal context. The number of formal concepts $|C|$ generated from a formal context $K := \langle G, M, I \rangle$, can be linear (in the best case) or quadratic (in the worst case) with respect to $|G|$ (the number of objects) or $|M|$ (the number of attributes) within the formal context, depending on the number of attributes per object. However, even withstanding the time and space complexities for initially computing and storing the complete concept lattices from a large formal context (which, if the system employed incremental algorithms to update the concept lattice, would only need to be run once), the worst case time complexity for updating a pre-computed concept lattice – i.e., only computing a portion of a concept lattice given changes to a formal context – is quadratic with respect to the number of formal concepts $|C|$, although experimental results [111, 112] cited in [93] suggest that in practice, the growth may be linear, rather than quadratic. Despite this, updating and storing a complete concept lattice for conceptual navigation poses major scalability and space concerns for large formal contexts.

CollectionWeb implements an alternate approach that does not require computation of the complete concept lattice and therefore negates the above scalability issues, but still allows the user to navigate between neighbouring formal concepts – via the reduction and inclusion of query attributes. This method, called the *conceptual neighbourhood approach* was first shown in *ImageSleuth* [98, 100] and is again used in the *Virtual Museum of the Pacific* case study in Section 4.2. In both cases, the user is never presented with the entire concept lattice: instead, they only view partial concept lattices in the form of a formal concept and its immediate neighbours.

The algorithm used by *CollectionWeb* for generating conceptual neighbourhoods is the NearestNeighbours algorithm [93], presented in Algorithm 3. The conceptual neighbourhood of a formal concept can be formed by finding both the upper and lower neighbours of a formal concept which can be computed separately. In the description of the algorithm that follows, a formal context is denoted by the triplet $\langle G, M, I \rangle$ with the finite non-empty sets of objects $G = \{0, 1, \ldots, g\}$ and attributes $M = \{0, 1, \ldots, m\}$ and $I \subseteq G \times M$ being an incidence relation with $\langle g, m \rangle \in I$ meaning that object $g \in G$ has attribute $m \in M$.

[5] In this description of worst-case time complexity and for all others that follow in this section, for a formal context $K := \langle G, M, I \rangle$, $|G|$ refers to the total number of objects in a formal context, $|M|$ refers to the total number of attributes in a formal context, and $|C|$ refers to the total number of formal concepts generated from the formal context.

Input: Context $\langle G, M, I\rangle$, concept (X, Y) of context $\langle G, M, I\rangle$
Output: The set of lower and upper neighbours of (X, Y) in the concept lattice of $\langle G, M, I\rangle$

```
// Returns the lower neighbours of (X,Y)
lowerNeighbours := ∅;
lNCandidates := ∅;
foreach m ∈ M \ Y do
    X1 := X ∩ {m}';
    Y1 := X1';
    if (X1, Y1) ∉ lNCandidates then
        Add (X1, Y1) to lNCandidates;
        count(X1, Y1) := 1;
    else
        count(X1, Y1) := count(X1, Y1) + 1;
    end
    if (|Y1| − |Y|) = count(X1, Y1) then
        Add (X1, Y1) to lowerNeighbours;
    end
end

// Returns the upper neighbours of (X,Y)
upperNeighbours := ∅;
uNCandidates := ∅;
foreach g ∈ G \ X do
    Y2 := Y ∩ {g}';
    X2 := Y2';
    if (X2, Y2) ∉ uNCandidates then
        Add (X2, Y2) to uNCandidates;
        count(X2, Y2) := 1;
    else
        count(X2, Y2) := count(X2, Y2) + 1;
    end
    if (|X2| − |X|) = count(X2, Y2) then
        Add (X2, Y2) to upperNeighbours;
    end
end
```

Algorithm 3: The NearestNeighbours algorithm used for generating a conceptual neighbourhood for formal concept (X, Y) in formal context $\langle G, M, I\rangle$, cf. [93]

The theoretical time complexity of the algorithm is $O(|G||M|(|G|+|M|))$ which is obtained by summing the time taken to find its lower neighbours, $O(|G||M|^2)$, and the time taken to find its upper neighbours $O(|G|^2|M|)$. Hence, the maximum running time of the algorithm is quadratic with respect to the number of objects or the number of attributes within the formal context – whichever is higher. As implemented in *CollectionWeb*, the NEARESTNEIGHBOURS algorithm runs dynamically at query time – i.e., every time a user views a formal concept or moves to an upper or lower neighbour, the new concept and its neighbouring concepts are computed dynamically. For the *ImageSleuth* (Section 3.2) and the *Virtual Museum of the Pacific* (Section 4.2) case studies, this means that any changes to the formal context – such as when new tags are added or removed from the collection – are immediately reflected in its underlying concept lattice, allowing the collection and the relationships among the objects to dynamically respond to user interaction, annotation and tagging.

However, the advantage offered by dynamically computing the conceptual neighbourhood – namely in that it negates the need to compute or store a potentially large concept lattice while still offering the ability to dynamically expose sections of it for user interaction – also presents another scalability limitation as the size of the collection grows. Given the dynamic nature of the query and the quadratic time complexity with respect to the number of objects in a collection, the *conceptual neighbourhood* approach becomes less suited for use in larger collections, as the response time for user interaction (in the worst case scenario) grows quadratically with respect to the number of objects in the collection. While the approach is well suited for dynamically presenting relatively smaller-sized collections at a specialist or 'exhibition' sized scale, such as the 427 objects present in the *Virtual Museum of the Pacific* or the 80 objects present in *A Place for Art*, the approach remains unsuited for larger collections, such as the one presented in the *Brooklyn Museum Canvas* case study.

The concept layer approach

For all other case studies in this thesis, *CollectionWeb* constructs and maintains a set of formal concepts from a formal context of collection objects. The set of all formal concepts for the formal context in *CollectionWeb* is called the *concept layer*. The framework relies on a concept layer in order to efficiently create the pathway structures as described in Sections 3.3 and 3.4 so that users can associatively browse and navigate the pathways within the collection. Although the pathway structures are dynamically generated from these pre-computed concepts, all formal concepts are stored and indexed by both their objects and attributes, allowing the framework to efficiently retrieve single or entire sets of formal concepts at query time.

To create and maintain the concept layer, *CollectionWeb* relies on an algorithm with a low running time for computing formal concepts from a formal context, and for incrementally recomputing formal concepts if any objects or attributes in the formal context changes. Specifically, the algorithm would be required to accommodate changes to a formal context in large museum data-sets if a single object (or a relatively small batch of objects) changes, ensuring that it can maintain pathway structures on a dynamically changing concept layer for a large museum data-set.

There are a number of high performance algorithms that compute formal concepts from formal contexts [113–115] along with associated evaluation studies of those algorithms as applied to data from the Web [116]. As these algorithms offer high performance batch computation of an entire set of concepts from a formal context, they work sufficiently well for large museum collections that are not likely to change over a period of time. However, it is common for museums to add or modify objects to their online collections, and some require the ability for the data to be kept up-to-date as it changes. For instance, the Brooklyn Museum data-set used for the *Brooklyn Museum Canvas* case study (Section 4.4), along with other large public facing data-sets such as the one provided by the Rijksmuseum used for this performance evaluation, require that as part of their terms of use, that all front-facing applications or any other representation of their content must show data that is up-to-date. [6] Further, large-scale collaborative tagging efforts such as the *steve.museum* project [39] and the Flickr Commons recognise museum collections as dynamic rather than static data-sets. This calls for a need for an efficient FCA algorithm that can accommodate incremental changes to a formal context, rather than perform the potentially time consuming task of recomputing the entire set of formal concepts.

CollectionWeb employs the FCbO algorithm to initially compute all concepts of a formal context [114], and a modification of that algorithm called *incremental FCbO* [117] to update formal concepts as objects in the formal context are updated, modified or deleted. The FCbO algorithm is an improved version of Kuznetsov's Close-by-One algorithm [118, 119] and *incremental FCbO* – although presented by Outrata [120], is again presented here for the purposes of self-containment. The example described here presents a scenario where new objects are added to the formal context, which results in the algorithm producing new and updated formal concepts.

In the description of the algorithm that follows, a formal context is denoted by the triplet $\langle X, Y, I \rangle$ with the finite non-empty sets of objects $X = \{0, 1, \ldots, m\}$ and attributes $Y = \{0, 1, \ldots, n\}$ and $I \subseteq X \times Y$ being an incidence relation with

[6] `http://www.brooklynmuseum.org/opencollection/api/docs/terms` `https://www.rijksmuseum.nl/en/api/terms-and-conditions-of-use`

$\langle x, y\rangle \in I$ meaning that object $x \in X$ has attribute $y \in Y$. Concept-forming operators defined on I are denoted by ${}^{\uparrow_I} : 2^X \mapsto 2^Y$ and ${}^{\downarrow_I} : 2^Y \mapsto 2^X$ [97]. New objects to be added to $\langle X, Y, I\rangle$ and not present in X are denoted by X_N, (i.e. $X_N \cap X = \emptyset$); Y_N is the set of attributes shared by at least one of the objects X_N and either present or not present in Y (but usually $Y_N \subseteq Y$), and $N \subseteq X_N \times Y_N$ is an incidence relation between X_N and Y_N. By the triplet $\langle X', Y', I'\rangle$ we denote the formal context which results as a union of $\langle X, Y, I\rangle$ and $\langle X_N, Y_N, N\rangle$, both extended to X' and Y', i.e. $X' = X \cup X_N$, $Y' = Y \cup Y_N$ and $I' \subseteq X' \times Y', I' \cap (X \times Y) = I$, $I' \cap (X_N \times Y_N) = N$ and $I' \cap (X \times (Y_N \setminus Y)) = I' \cap (X_N \times (Y \setminus Y_N)) = \emptyset$.

The algorithm is represented by the recursive procedure UpdateFastGenerateFrom (Algorithm 4) which is a modified form of recursive procedure FastGenerateFrom – the core of the FCbO algorithm as described in [114] (Algorithm 2). The procedure accepts as its arguments a formal concept $\langle A, B\rangle$ (an initial formal concept), an attribute $y \in Y_N$ (first attribute to be processed) and a set $\{N_y \subseteq Y \mid y \in Y\}$ of sub-sets of attributes Y (see [114] for the meaning of the set), and uses a local variable *queue* as a temporary storage for the computed formal concepts and M_y ($y \in Y$) as sets of attributes which are used in place of the third argument for further invocations of the procedure. When the procedure is invoked, it recursively descends, in a combined depth-first and breadth-first search, the space of new and modified formal concepts of $\langle X', Y', I'\rangle$ resulted by adding new objects X_N described by attributes Y_N to $\langle X, Y, I\rangle$, beginning with $\langle A, B\rangle$. For a full description of the procedure, see [117] and [114], recalling that the set $Y_j \subseteq Y'$ is defined by: $Y_j = \{y \in Y' \mid y < j\}$. In order to list all new and modified formal concepts of $\langle X', Y', I'\rangle$ which are not formal concepts of $\langle X, Y, I\rangle$, each of them exactly once, UpdateFastGenerateFrom is invoked with $\langle \emptyset^{\downarrow_{I'}}, \emptyset^{\downarrow_{I'}\uparrow_{I'}}\rangle$, y being the first attribute in Y_N and $\{N_y = \emptyset \mid y \in Y\}$ as its initial arguments.

There are also a number of other algorithms that allow for incremental object-by-object updates to a set of formal concepts, or concept lattice, given changes in a formal context, such as [121, 122] along with the algorithms in [93]. AddIntent [121] is considered to be one of the more recent and efficient of these algorithms, however, along with the other algorithms, it requires prior computation of the complete concept lattice. The incremental FCbO algorithm [117] differentiates itself from other incremental algorithms in that it does not require the concept lattice as its input. However, the number of concepts computed from these data-sets – even without the complexities of storing a lattice – can range in the hundreds of thousands (see Figs. 4.8 and 4.9). In light of this, incremental FCbO not only computes changes only based on a set of objects marked for additional, modification or deletion, but it also outputs only the new, modified or deleted formal concepts, rather than the entire set of formal concepts. This allows for the both quick execution of the algorithm

```
// list ⟨A, B⟩, e.g., print it on screen or store it
if (A ∩ X)^{↑I'} = B then
    if (A ∩ X) ⊂ A then
        list ⟨A, B⟩ as modified;
    end
else
    list ⟨A, B⟩ as new;
end
if B = Y' or y > n' then
    return
end
for j from y upto n' do
    set M_j to N_j;
    // go through attributes from Y_N only
    if j ∉ B and j ∈ Y_N and N_j ∩ Y_j ⊆ B ∩ Y_j then
        set C to A ∩ {j}^{↓I'};
        set D to C^{↑I'};
        if B ∩ Y_j = D ∩ Y_j then
            put ⟨⟨C, D⟩, j⟩ to queue;
        else
            set M_j to D;
        end
    end
end
while get ⟨⟨C, D⟩, j⟩ from queue do
    UpdateFastGenerateFrom(⟨C, D⟩, j + 1, {M_y | y ∈ Y});
end
return
```

Algorithm 4: The UpdateFastGenerateFrom($\langle A, B\rangle, y, \{N_y \mid y \in Y\}$) algorithm used in incremental FCbO, cf. [117]

Table 4.1: Running time of computing all formal concepts from a formal context using the incremental FCbO algorithm, average of 10 iterations

Data-set	No. of attributes	No. of objects	No. of concepts	Avg. running time (ms)
Brooklyn Museum	8,952	10,000	98,547	36,218
Rijksmuseum	1,716	100,000	994,967	68,792

and ingestion of its results where changes to its formal context are relatively minor: strengthening the algorithm's utility in applications where data-sets are large but updated frequently and in small increments.

The algorithm was evaluated on two museum data-sets: the first being the Brooklyn Museum collection consisting of 10,000 objects and 8,952 attributes and the second being the Rijksmuseum collection consisting of 100,000 objects and 1,716 attributes. The formal contexts of these two data-sets were extracted by identifying key words and phrases from their data-sources using parsers from the *CollectionWeb* framework. [7] Unlike the data-sets that are traditionally used in the literature to evaluate FCA algorithms, namely the University of California, Irvine machine learning data-sets [8], the formal contexts of these two data-sets have very low FCA matrix densities [116] of 0.18% and 0.89%, respectively.[9] As reported in Kirchberg et al. [116], these characteristics typically describe Web data sources.

All tests were performed on a machine with a single-core Intel processor at 2.8 GHz with 8GB of memory and 8GB of swap space running Ubuntu 10.04 LTS. The presented algorithm tested was implemented in ANSI C in a manner similar to its non-incremental counterpart [114]. All execution times reflect the actual time it took the algorithm to read, pre-process and compute the formal concepts, and does not account for the time it took for the algorithm's output to be processed, stored or displayed.

To compare the time differences, and hence utility, of incremental vs. non-incremental formal concept computation, Table 4.1 shows the time it took to compute the entire set of concepts from each formal context. For the sake of clarity, a *batch* or *non-incremental* computation – such as the one demonstrated in Table 4.1

[7] The objects from these online museum collections were extracted from their respective data sources at `http://www.brooklynmuseum.org/opencollection/api/` and `https://www.rijksmuseum.nl/en/api/`

[8] The University of California, Irvine (UCI) machine learning data-sets are commonly used for performance evaluations of machine learning algorithms: `https://archive.ics.uci.edu/ml/datasets.html`

[9] The FCA matrix density, as reported in Kirchberg et al. [116], describes the number of incidence relations within a formal context relative to its size. For a given formal context as a triplet $\langle X, Y, I \rangle$ where $|X|$ refers to the number of objects, $|Y|$ refers to the number of attributes and $|I|$ refers to the number of incidence relations between each object and attribute, the matrix density is defined as $|I|/(|X| + |Y|)$. As a comparison, the UCI data-set `mushrooms.dat` has a matrix density of 19.167%.

– is defined as a computation that computes the entire set of formal concepts from formal context, whereas an *incremental* formal concept computation is defined as a computation that outputs only the formal concepts marked for addition, deletion or modification based on a set of objects marked for addition, deletion.

A performance evaluation was conducted on the incremental FCbO algorithm for how quickly it could update formal concepts given changes to a formal context. Given that objects can be *added*, *removed* or *updated* within a museum data-set, these three operations are defined and evaluated separately with respect to the running time of the incremental FCbO algorithm. Assuming that a full set of formal concepts have already been computed, each operation produces a number of *changed concepts* that refer to the formal concepts that are added, modified or removed as a result of each operation. In addition to the time it takes to perform each operation, the number of changed concepts serves as an important indicator of complexity due to the fact that a system implementing this algorithm would need to process a potentially large set of formal concepts from its output, even if the time to compute those concepts by the algorithm is relatively small. The three operations, as implemented in the performance evaluation, are defined as followed:

- *Add*: Defined when a new object, or set of objects, are added to the data-set. For each test iteration, the new object(s) selected for addition are randomised. The *changed concepts* are the set of new formal concepts induced by the addition of new object(s) along with the set of existing concepts in which their extents incorporate the addition of the new object(s).

- *Remove*: Defined when an existing object, or set of objects, are removed from a data-set. For each test iteration, the new object(s) selected for removal are randomised. The *changed concepts* are the set of formal concepts induced by the removal of the object(s) along with the set of existing concepts in which their extents require the removal of such object(s).

- *Update*: Defined when an existing object, or set of objects, are updated, modified or replaced within the data-set. Within the practice of museum collection cataloguing, it is common to update the metadata of objects already within the collection. To realistically test this scenario, each test iteration randomly selects an existing set of objects for modification, and for each of these objects, 20% of their attributes are randomly replaced with other attributes from the collection. Each iteration of the operation causes the algorithm to run twice – first to remove the modified objects, and second to reinsert them with their new attributes. As a result, the number of concept changes may be greater than the number of formal concepts within the context for very large batch

updates. The *changed concepts* are the concepts that were modified and removed as a result of its first run, along with new and modified set of concepts as a result of its second run.

One of the benefits of an incremental concept update algorithm is the ability to apply small changes on an individual object and quickly update those changes to the concept layer. For example, visitors can add their own interpretations to the objects by adding their own keywords or 'tags'. This has been known to introduce new perspectives on the works within the collection [39] and could be potentially used as a way of re-framing the way objects are related to one another [5]. By allowing the formal concepts to incrementally update as users add new tags without needing to recompute the entire concept lattice, these equivalencies can be realised in real-time. In many other cases, updates to museum collection data are provided as a batch – i.e., whole groups of objects may be added or modified to the collection as a result of changes to objects within a museum data-set. For example, the Smithsonian Cooper-Hewitt National Design Museum uses GitHub [10] to host their collection data [11] – allowing anyone to easily access, update and provide updates to the collection. Many other museums provide a timestamp in their object records to indicate when it was last updated, so that data harvesters and aggregators can collect updated records as a single batch. In other situations it may be more feasible to implement updates to the data-set as a batch rather than as a set of frequently occurring single object updates.

The results of a performance evaluation demonstrating *add*, *remove* and *update* operations for the incremental FCbO algorithm are shown in Fig. 4.8 for the Brooklyn Museum data-set, and Fig. 4.9 for the Rijksmuseum data-set. For each operation, the results demonstrate the average running time of the algorithm over the two collections, along with the number of modified concepts [12] from its output. The figures also demonstrate how the algorithm scales, with each operation adding, removing or modifying 1, 5, 50 or 500 objects to their respective data-sets. As a way of comparing the the running time of the incremental FCbO to its non-incremental counterpart, the performance metrics of the incremental FCbO algorithm – its running time and number of modified concepts – are shown along with the total running time and number of formal concepts produced by the non-incremental FCbO algorithm, as indicated by the dashed line in Figures 4.8 and 4.9.

For the smaller Brooklyn Museum collection, the number of modified concepts and time taken to compute them is within reasonable ranges when adding 5 or

[10] GitHub is a popular source code management system traditionally used for making available, committing and providing updates to, program source code.

[11] See: `http://www.cooperhewitt.org/collections/data`

[12] The term 'modified concepts' in this instance refers to any formal concepts added, removed or modified as a results objects being added, removed or modified to its formal context.

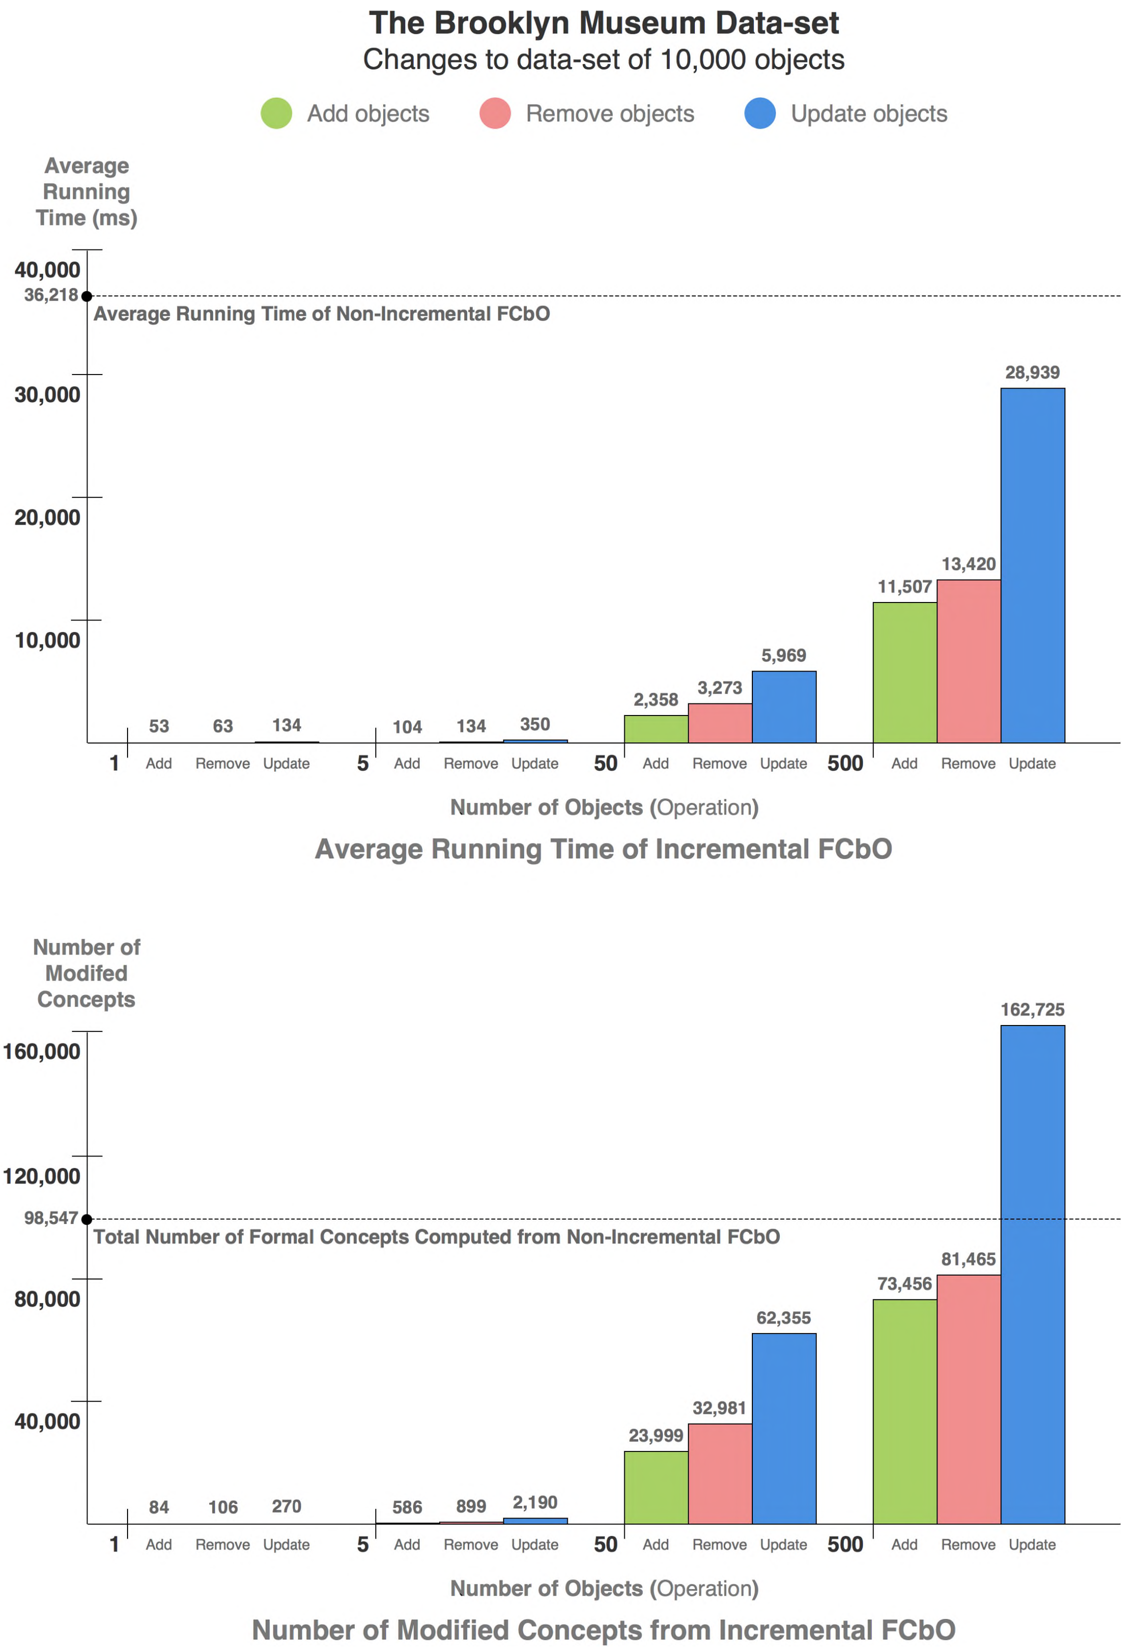

Figure 4.8: Average running time and number of modified concepts for adding, removing or updating objects to a formal context and incrementally recomputing the set of formal concepts using the incremental FCbO algorithm on the Brooklyn Museum data-set. The horizontal axis demonstrates how the algorithm scales with respect to the number of objects added, removed or modified.

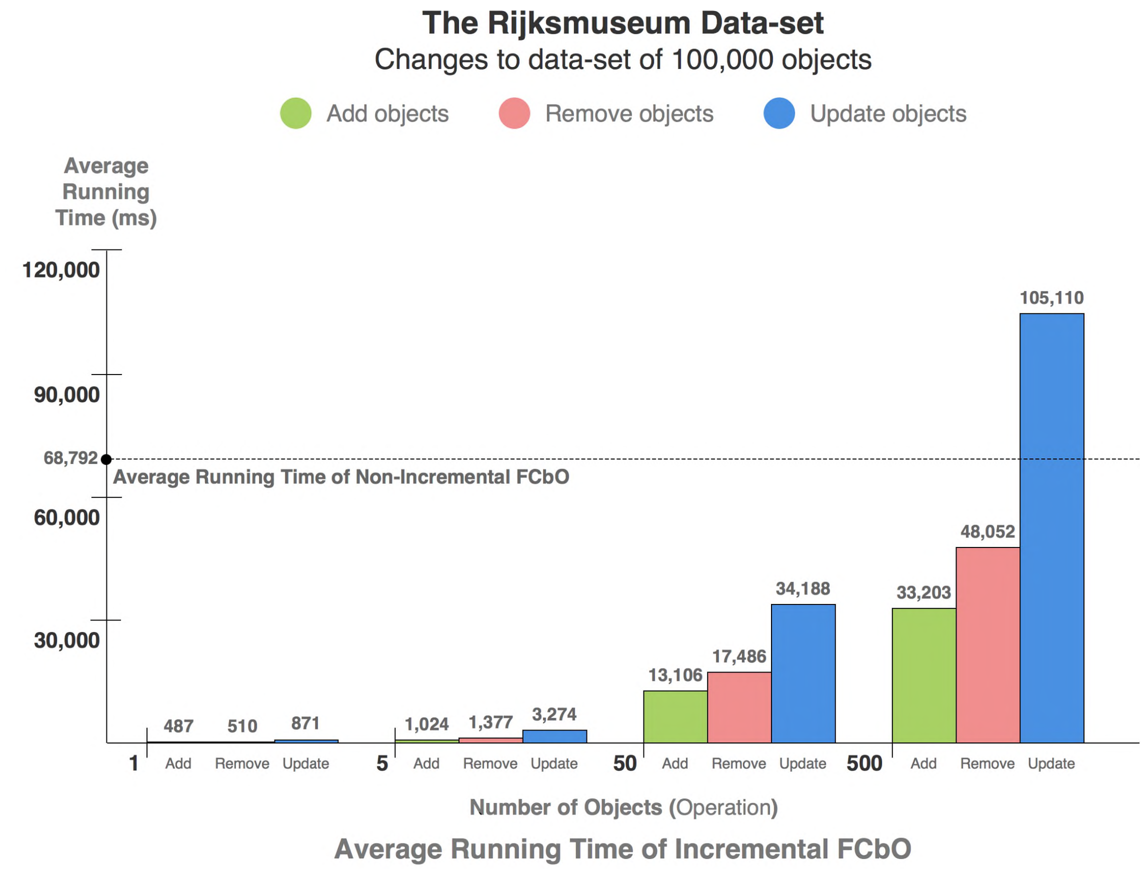

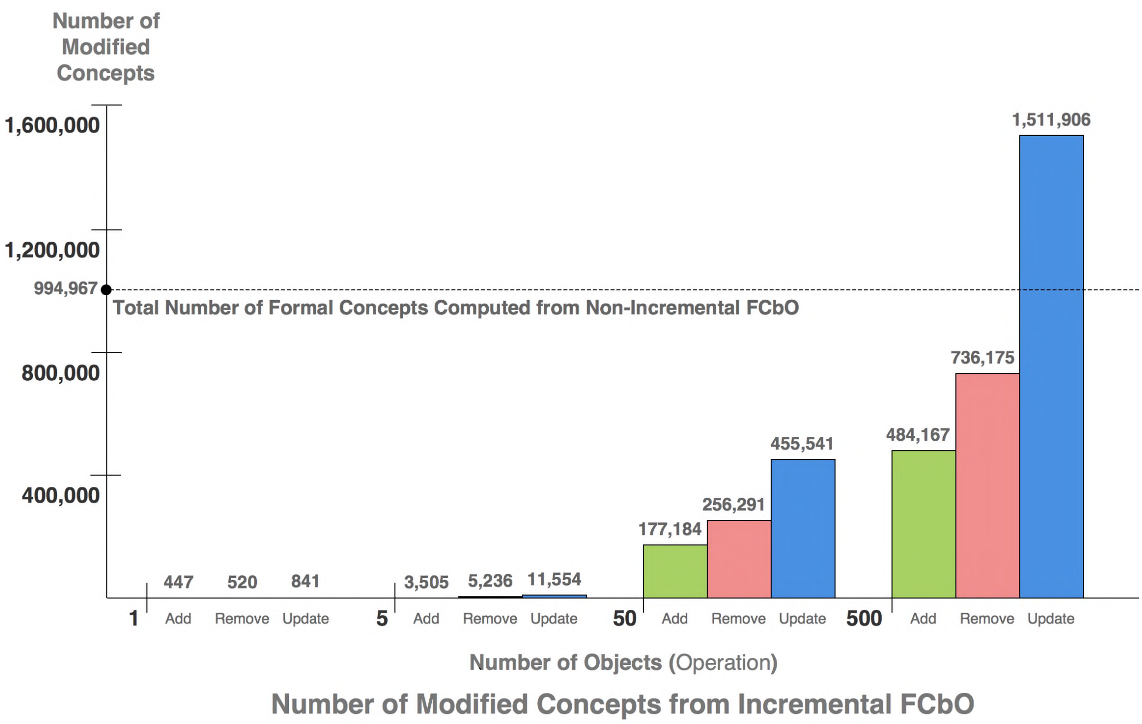

Figure 4.9: Average running time and number of modified concepts for adding, removing or updating objects to a formal context and incrementally recomputing the set of formal concepts using the incremental FCbO algorithm on the Rijksmuseum data-set. The horizontal axis demonstrates how the algorithm scales with respect to the number of objects added, removed or modified.

50 objects, with running times far less than the time it takes for the algorithm to non-incrementally recompute the entire set of formal concepts. However, in the larger Rijksmuseum collection – due to the smaller number of attributes and higher context density – removing and updating a larger batch of objects requires the re-computation of a very high number of concepts where in some cases, as indicated in Figures 4.8 and 4.9, the time taken to update the set of formal concepts incrementally is greater than the time to recompute the entire set as a non-incremental batch operation.

Overall, the incremental FCbO algorithm – as implemented by *CollectionWeb* to construct and maintain its concept layer – provides a fast way to update formal concepts from large and dynamically changing museum data-sets, given that the changes within those data-sets are relatively small relative to the size of the formal context. The algorithm provides a scalable way to construct and maintain a concept layer once the initial and potentially time costly computation of the entire set of formal concepts from a formal context is complete. The algorithm is less efficient at adding, removing or updating large changes to the collection where, in such cases, it may be preferential to recompute the entire set of formal concepts non-incrementally.

4.2 The Virtual Museum of the Pacific

The *Virtual Museum of the Pacific* is a rich Web-based browsing interface that provides access to 427 ethnographic objects from the Australian Museum's Pacific cultural collections. [13] The site contains high resolution imagery, interpretive descriptions of its objects and offers the ability to explore the collection using the *conceptual neighbourhood* paradigm described in Section 3.2. It allows the user to associatively navigate the collection through the inclusion and exclusion of query terms – shifting their focus on the information space as they move to different, but tangentially related aspects of material Pacific culture. It also allows creator and Pacific islander communities, curators and exhibitors to upload rich media and add their own tags and attributes to the collection, that in turn, update its underlying conceptual structures and influence the relationships among the objects.

The app embraces the design principles of *pathways*, *pliability*, *visual imagery* and *similarity-based suggestions* and Section 4.2.1 describes their implementation. It also presents a case study in the challenges that arise in constructing and presenting a rich and associative browsing experience within the institutional and organisational context of a museum, and Section 4.2.2 examines how the Pacific cultural collection was presented with respect to existing organisational processes and information classification structures of the Australian Museum, along the selection process behind the objects in ensuring that meaningful browsing experience can be conveyed. In Section 4.2.3 two independent user experience evaluations of the *Virtual Museum of the Pacific* are presented: one from internal staff and stakeholders of the Australian Museum, and another from external users and stakeholders from three key demographic groups: social media savvy users, academics and Pacific Islanders. Overall, the results from the evaluations conclude that the *Virtual Museum of the Pacific* makes a useful addition to the museum's representation and knowledge practices, particularly with respect to making parts of the collection accessible to the general public, although they also pointed to a number of issues with respect to the terminology used within the app, the way the objects were classified, and in some cases, the complexity of the interface.

As a case study, the *Virtual Museum of the Pacific* represents a highly specialised domain of a museum's cultural collection that also examines the organisational context of presenting a well curated selection of content with high quality metadata. As an anthropological collection, participants who identified themselves as Pacific islander also displayed a strong connection the geographic context of the objects. Although it illustrates a best case scenario in terms of metadata qual-

[13] The *Virtual Museum of the Pacific* can be accessed at `http://epoc.cs.uow.edu.au/vmp/`. Log in with username `vmpuser` and password `uowvmp`.

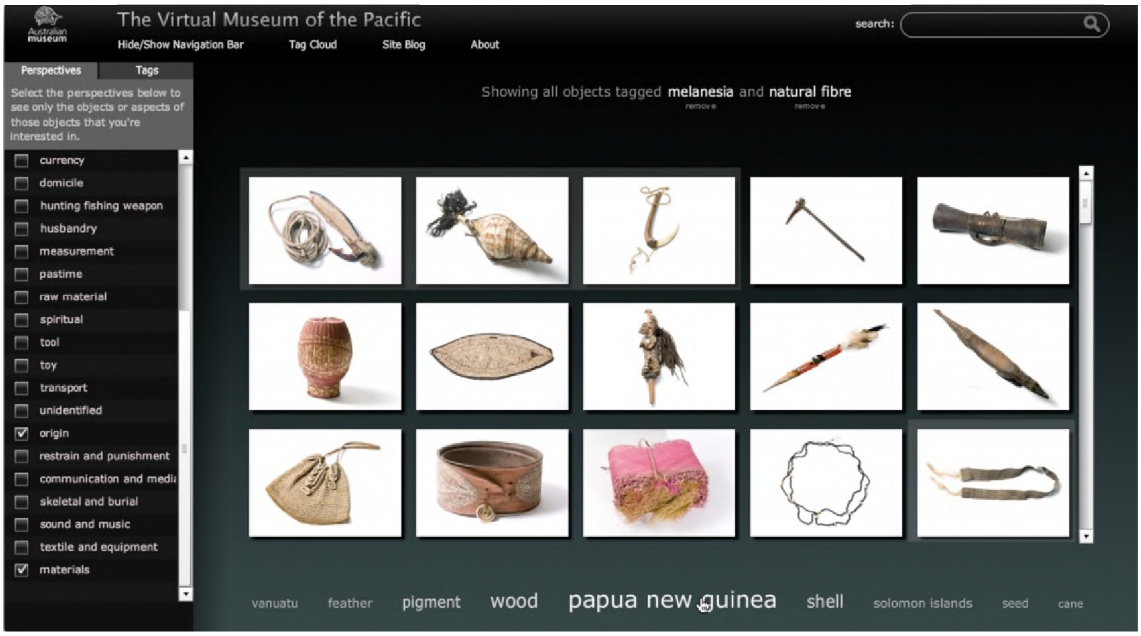

Figure 4.10: Screenshot of The *Virtual Museum of the Pacific*.

ity and consistency, the museum professionals from the internal evaluation noted that the terminology and meta-data models used to describe and classify the objects needs to reflect a wider range of perspective than the current 19th century vocabulary. The study also considers the significant time and organisational costs of manually curating and normalising data to a standard that can be readily applied for concept-lattice based browsing.

4.2.1 Design and implementation

The *Virtual Museum of the Pacific*, shown in Fig. 4.10, was originally conceived as a project to experiment with the use of Formal Concept Analysis to organise and browse digital library content. Rather than align objects to a fixed taxonomy, objects are presented as an associative network based on semantic themes according to a museological view of cultural heritage. The *Virtual Museum of the Pacific* also allows the creation of extensible views, called *perspectives*, and provides a social tagging interface that encourages the museum's stakeholders to annotate content and build custom views of the collection according to their communities of interest.

While its interface affords traditional keyword search-based information retrieval, its primary mode of interaction uses the *conceptual neighbourhood* paradigm of navigating pathways across its collection described in Section 3.2. As implemented within the *Virtual Museum of the Pacific*, users currently view a single formal concept at a time and can navigate to upper and lower neighbours by clicking on tags above and below the thumbnail list of objects. The thumbnail list represents the objects of the current formal concept, whereas the attributes of the formal concept

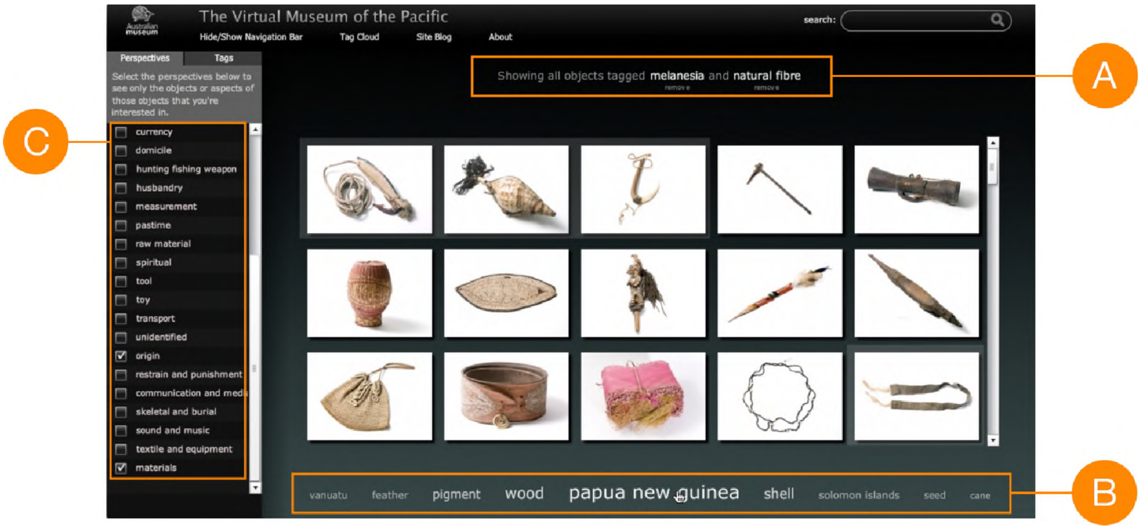

Figure 4.11: The main user interface components of the *Virtual Museum of the Pacific*, consisting of a) upper neighbor navigation links, b) lower neighbour navigation links and c) perspectives.

are represented as the tag links above the thumbnail list, which can be 'removed' as the user moves to an upper neighbour. Clicking on a tag below the grid of image thumbnails moves the user to a more specific formal concept and reduces or 'filters' the objects on the current view with that attribute. Dually, clicking on an attribute above the grid of thumbnails removes the filter criterion, causing more objects to be displayed. The tags that appear above or below the image thumbnail grid represent the corresponding attributes are that are removed or excluded when a user navigates between upper and lower navigations. These navigation links are respectively labelled *A* and *B* in Fig. 4.11. As the user navigates these neighbouring concepts, they create a pathway through the collection. The lower neighbour navigation links are also weighted in size according to the frequency of objects: links that navigate to relatively larger sets of objects appear larger and hence convey a visual overview of the prominence of each attribute within the current group of objects. For example, in Fig. 4.11, the larger weighting of the tag 'Papua New Guinea' would suggest that it contains a large number of objects relative to the number of objects in the current view, and hence, it also represents the most frequently occurring attribute of the current objects in view.

Embracing the design principles of *pliability* [14], the collection appears to 'contract' and 'grow' as a direct response to attributes filtering and unfiltering the objects in view. These incremental changes, along with the use of animations and visual momentum to convey them, are intend to represent discreet, yet fluid transformations as users navigate their pathway through the concept lattice.

An example of conceptual neighbourhood navigation is presented in Fig. 4.13. The first screenshot, labelled *A*, presents all objects tagged *melanesia*, *papua new*

Table 4.2: The formal context $K := (G, M, I)$ containing information about the objects (G) and their attributes (M) for the sub-context resulting from applying *body accessory* and *origin* scales.

K	body accessories	fan	head ornament	ankle ornament	fly whisk	neck ornament	melanesia	polynesia	fiji	papua new guinea	samoa
e002509	×	×					×		×		
e090525	×	×					×		×		
e058551-004	×				×		×		×		
e091567	×					×		×	×		
e091570	×					×	×		×		
e002415	×	×						×			×
e002416	×	×						×			×
e058169	×		×					×			×
e058169	×				×			×			×
e011543	×			×			×			×	

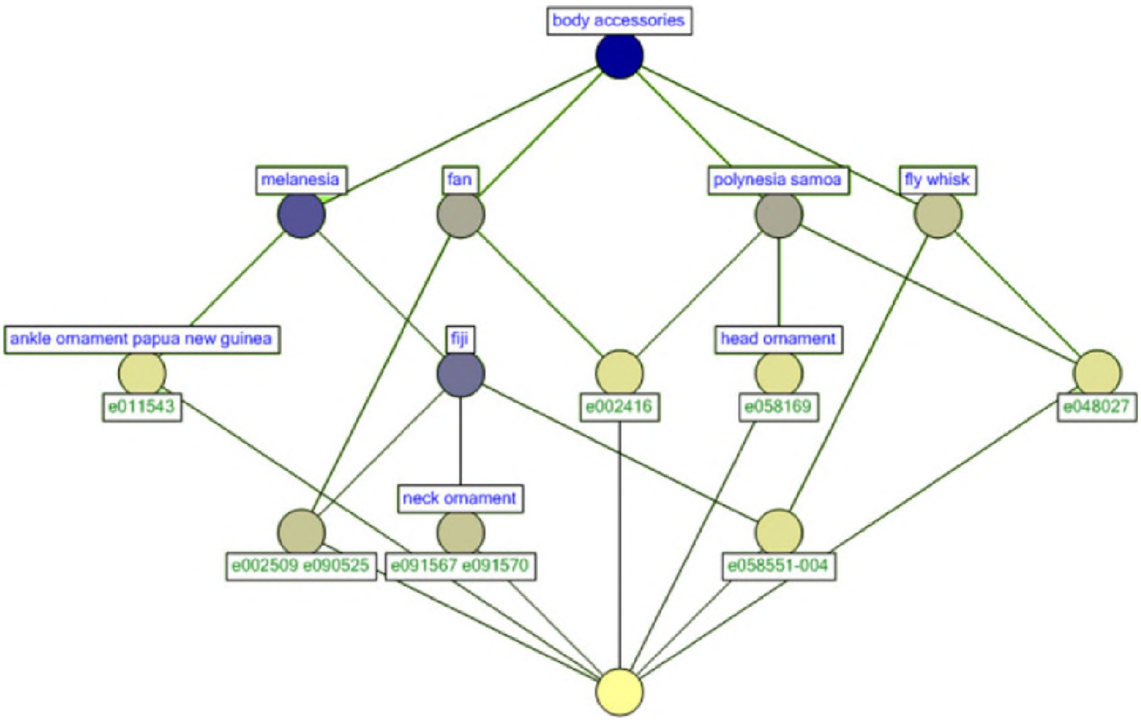

Figure 4.12: The line diagram of the concept lattice of the *Virtual Museum of the Pacific*, derived from the context in Table 4.2.

guinea, *natural fibre*, and *east sepik province* with the user's mouse cursor hovered over the lower navigation link *pigment*. Clicking on this link 'moves' the view towards the concept's lower neighbour, filtering the collection and revealing a subset of those objects tagged *pigment*, resulting in the set of objects shown in screenshot *B*. Navigating to a concept's lower neighbour restricts the object set by adding another attribute to the view, so that the collection of objects appear to shrink, while an additional attribute is added to the list of upper neighbour navigations at the top of the screen.

The *Virtual Museum of the Pacific* provides another feature, called *perspectives*, that provide faceted browsing of its collection. A perspective is an affordance that allows the user to apply or remove a *conceptual scale*, that, when selected or combined, can be used to provide focused views on the collection. These views are *derived contexts* – parts of the collection in which objects are delimited by the attributes that they possess. As an example of how perspectives are implemented in the *Virtual Museum of the Pacific*, consider a user who is currently browsing the entire set of 427 objects within the collection, navigating across neighbouring concepts that are tagged using the entire set of attributes that are used to tag the objects. In this example, the user would then select one of the 29 pre-defined perspectives: the *body accessory* perspective. Selecting one or more perspectives creates a derived context that consists of only the attributes in that perspective, and the objects that are tagged with those attributes. In this case, the *body accessory* perspective contains only attributes that represent object names that are categorised under 'Body Accessories' from the Australian Museum Thesaurus, such as *face ornament*, *kap kap*, etc. Selecting this perspective 'filters' the view so that the user is only presented with these objects, and that the lower and upper navigation options consist only of various sub-categories of body accessory objects.

The user would then apply the *origin* perspective that contains attributes that describe where the object originated from. This generates the derived context shown in Table 4.2. Within this context the attributes are hierarchically related, so that for example, all objects tagged *fiji* are also tagged *melanesia*, along with all other objects tagged *papua new guinea*. This hierarchy is evident in the concept lattice presented in Fig. 4.12. [14] The user would then navigate conceptually within this lattice: moving to upper and lower neighbours that are linked via the type of body accessory object and its origin. Further perspectives could be applied so that, for example, the user could further navigate by what the object was made of by applying the *materials and medium* perspective, or navigate within and across other categories,

[14]For the purposes of clarity and brevity, the concept lattice presented in this figure contains only a sub-set of objects from the *body accessory* and *origin* scales, as the actual lattice used in the application would be too complex to to be legible on a single page.

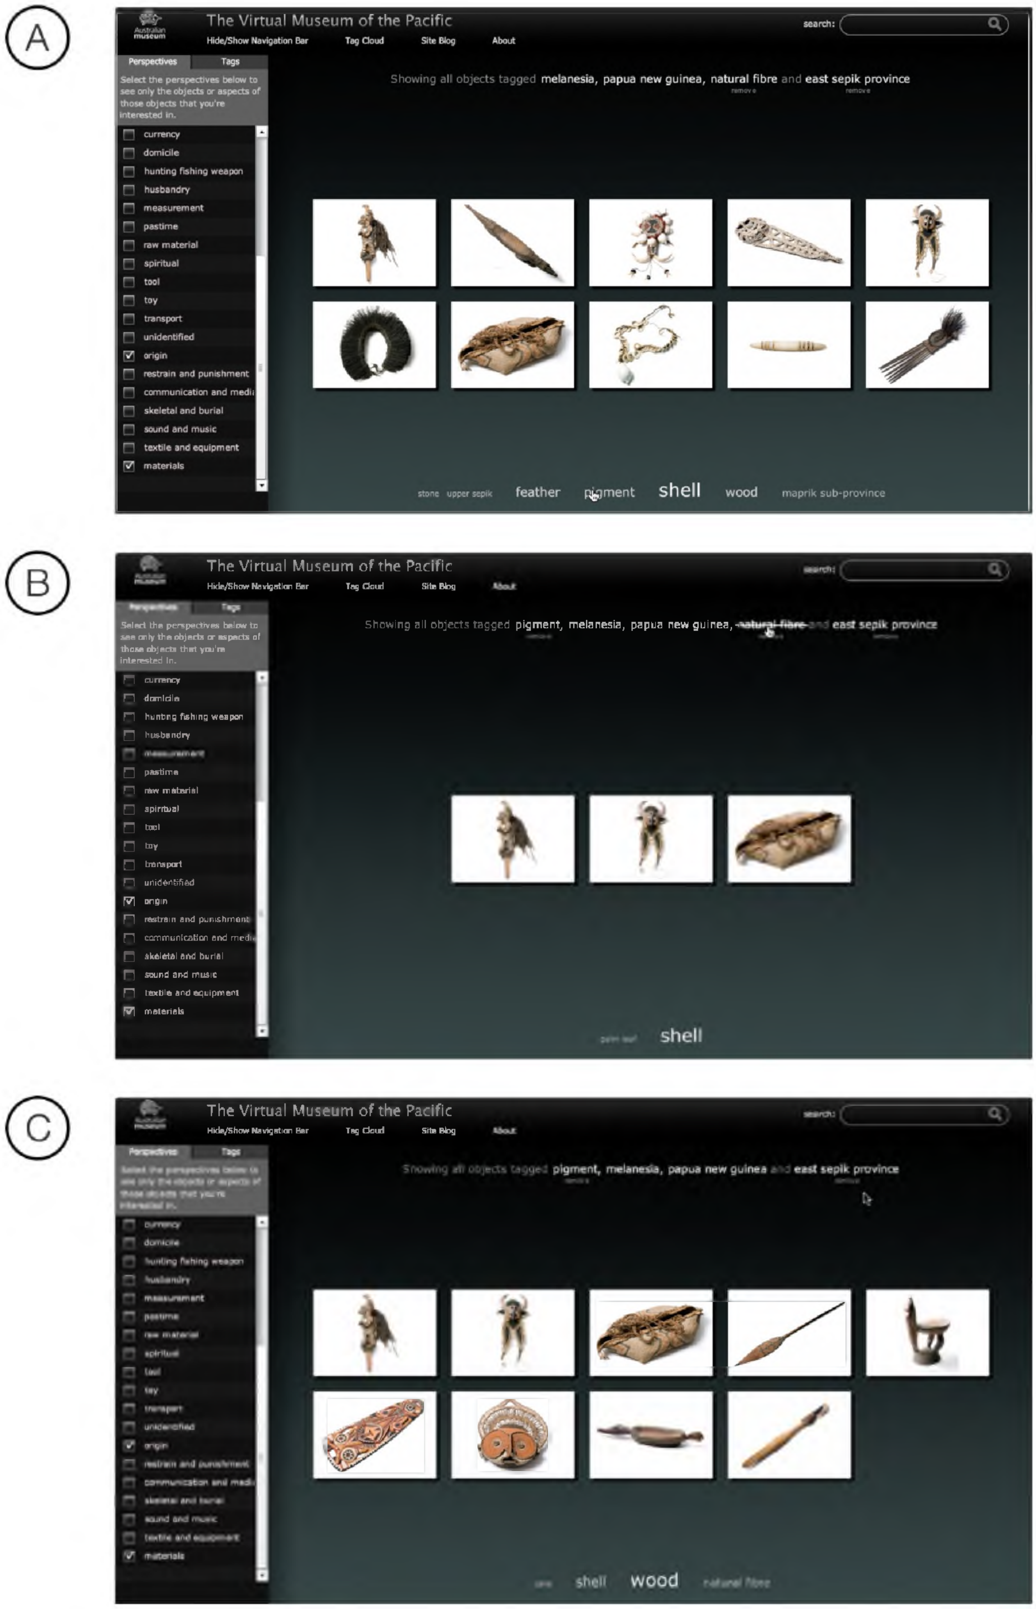

Figure 4.13: An example of conceptual neighbourhood navigation within the *Virtual Museum of the Pacific.*

such as *clothing*, *currency* and *hunting* objects.

The *Virtual Museum of the Pacific* also allows creator communities to add their own tags to the objects that can be selected from the museum's control vocabulary. Since the concept lattice – the very structure that determines the navigation pathways within the collection – is built from a formal context of objects and tags, tagging within the *Virtual Museum of the Pacific* not only affects the meta-data displayed but the structure and presentation of the collection itself. This has implications in that the structure and organisation of the collection can be entirely data-driven from the tagging efforts of museum stakeholders and creator communities. Perspectives – given that they delimit attribute sets – provide the ability to separate tags categorically or according to different stakeholder or institutional lenses on the collection. For example, within the *Virtual Museum of the Pacific*, perspectives are used to represent the various facets provided by the Australian Museum's thesaurus, but also provide an alternate lens that incorporates Pacific Islander terminology and descriptions of objects. Using perspectives, it is possible to view a representation of the collection using any of these lenses. Collaborative tagging allows a diverse range of communities to apply their own views and interpretations of the content, whereas perspectives provide granular control over what views are represented on the collection, and the way metadata and folksonomy tags can both shape the content and interpretation of an object, along with the way objects are contextually linked and represented within the collection.

The *Virtual Museum of the Pacific* provides also provides a content-based retrieval feature, so that, for any given object, a user could search for other, similar objects. Fig. 4.14 shows the 'Objects Like This' feature that presents an order-ranked list of formal concepts that are similar to the formal concept formed by an object of interest, using the *concept similarity* metric described in Section 3.3. This metric traverses the concept lattice and produces a list of formal concepts that are ranked according to their *concept similarity*. From this list of results, a user can then select any formal concept that is represented in the list, in which they are then taken to its conceptual neighbourhood view (Fig. 4.10) and can subsequently browse the collection within its conceptual neighbours.

4.2.2 Organisational context: artefact selection, classification and data normalisation

In this case study, 427 objects were selected and prepared for the specific purpose of exhibition within the *Virtual Museum of the Pacific*. Examining this process provides insights to the organisational fit, feasibility and costs of the research and digitisation efforts required to produce high quality imagery and metadata. Working

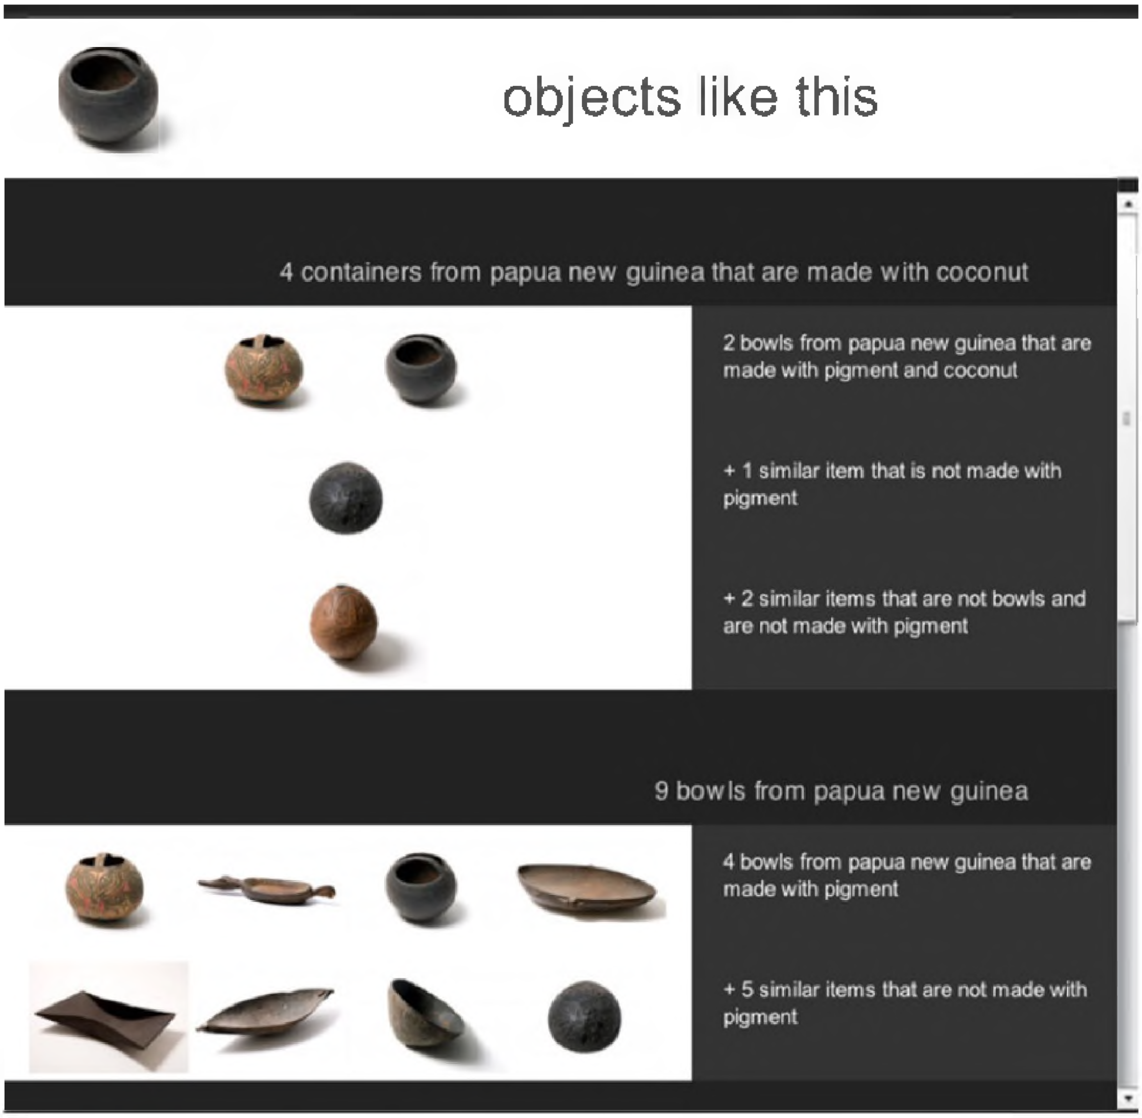

Figure 4.14: The 'Objects Like This' feature of the *Virtual Museum of the Pacific*, showing a list of formal concepts order-ranked according to their *similarity*, and then clustered according to matching superconcepts.

with the museum, a project anthropologist was commissioned to select 427 objects from the museum's collection of approximately 60,000 Pacific objects. The objects were selected based on the following criteria:

- Their knowledge of Pacific culture.
- Their knowledge of Pacific collection.
- The aim to create a representative sample of Pacific objects that are relevant and have wide appeal for community members, researchers and the general public.

Critical to the object selection process was the need to include objects with similar and overlapping qualities, so that the *Virtual Museum of the Pacific* can convey new and serendipitous connections via its conceptually-based browsing paradigm. This was done by selecting multiple types of objects from different categories (body accessories, ceremonial objects, ornaments etc.) without restricting each category to objects from a particular provenance, or to objects made with a particular material. This ensured that the *Virtual Museum of the Pacific* could represent a diverse cross-section of Pacific material culture while also providing opportunities for exploration by linking objects across multiple facets and dimensions.

Overall, the criteria for the selection of objects from The Australian Museum's Pacific Collection is based on the following:

- The object's cultural relevance.
- A diverse, yet overlapping selection of objects that may share commonalities in object type, provenance, materials, etc.
- Objects that include a variety of materials and production techniques.
- Objects that could be seen as interesting from stakeholder groups.
- The relative importance of the object, based on historical value.

The selection of objects for the *Virtual Museum of the Pacific* also considered the quality, consistency and provenance of the object metadata. The process required that all objects be linked to a set of attributes or tags that are applied in a consistent fashion, or preferably, linked to a controlled vocabulary. The metadata for the objects was initially imported from the Museum's Content Management System: the one currently used being the third effort by the museum to digitally store its records of the Pacific Collection.

The Australian Museum acquired the objects in its Pacific Collection from many sources over 150 years. The process of adding an object to the collection, called accessioning, involves creating a registry entry ''to associate provenance information and collection metadata with the object. Accessioning instantiates a registration number – for example, the object names in Figure. 4.12 and Table 4.2 'e002509', 'e090525' etc. are registration numbers for some of the objects in The Australian Museum's Pacific Collection. In the 1950s an index card system was created which summarized the object's provenance, added more detailed descriptive text, and on its reverse, the object's physical measurements. Later again, as objects were added to one of the three generations of content management systems used by the museum, objects were further described. During the transition to the most recent commercial Museum Content Management System, a simple, practical corporate taxonomy was developed: the museum's *Archeology and Anthropology* taxonomy is two-level, with 27 categories and 709 item types distributed across those categories. This taxonomy was used as the control vocabulary to provide attributes for the 427 objects represented within the *Virtual Museum of the Pacific*. Over time, The Australian Museum has developed a rigorous process for metadata quality assurance to ensure that terms entered into the Museum Content Management System were compatible with the attribute tags for the *Virtual Museum of the Pacific*.

Metadata for the selected objects are exported from the Museum Content Management System, then a strict metadata ingestion process is applied to extract 'tags' from exported fields which contain either keywords or more informal narrative-like text. The data is carefully curated by The Australian Museum Cultural Collections sections to conform to a formal taxonomy using rules which are based on inferences derived from the existing Museum's thesaurus, along with contextually placed keywords and terms relating to an object's *materials* or *indigenous name*. Fig. 4.15 demonstrates this process. Within the Museum Content Management System, both the object's Western name (e.g. 'ornament') and indigenous name (e.g. 'kap kap') are recorded, and both fields are implemented as separate *perspectives* within the VMP so that a visitor browsing the site can view the objects and their related concepts from a Western or indigenous perspective.

From statistics collected during preparation of the 427 objects for the *Virtual Museum of the Pacific*, about 70% of the objects in the Pacific Collection have at least a basic entry in the Museum Content Management System. Nearly all objects required metadata preparing and cleaning to bring them to a uniform exhibition standard. This involves normalizing spelling and thesaurus checking, for instance testing whether the attributes *mother of pearl* or *pearl shell* should be used or whether the synonymous terms *dagger* or *knife* should apply to a particular object. Based on observing the average effort required for the project's anthropologist

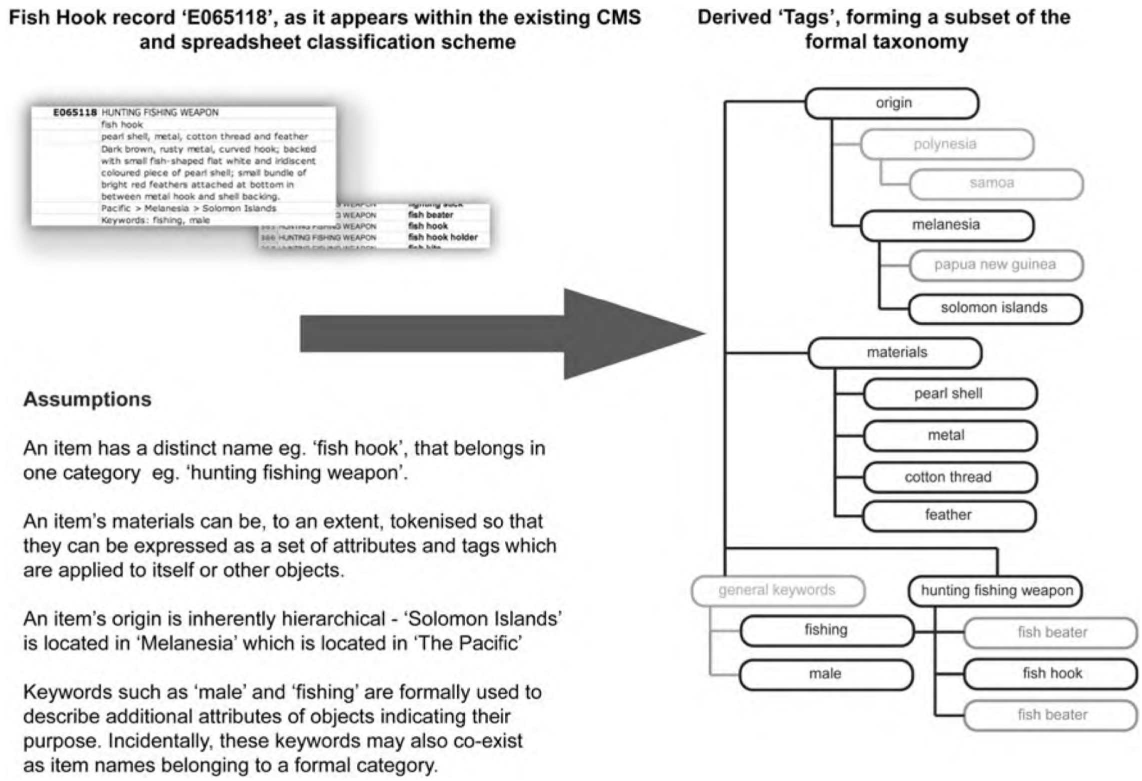

Figure 4.15: The process of extracting attributes from the Australian Museum's metadata using its corporate *Archeology and Anthropology* taxonomy.

to research, clean and prepare the metadata, it took approximately one hour per object to normalise the object's metadata, and another hour to write its descriptive and interpretive label. While the metadata adds enormous value to the value of the object within the *Virtual Museum of the Pacific*, it also carries significant monetary and time costs.

As a result of the effort required to bring objects to a presentable exhibition standard, the *Virtual Museum of the Pacific* is considered suitable for exhibiting a featured collection or exhibition of curatorially selected and well-researched works that – within the context of museum resourcing – would not be feasible beyond hundreds of objects. The manual effort in researching and selecting the metadata ensures that the use of concept lattices for navigation and browsing can be exploited in the best possible way over a relatively small and specialised data-set. Manually constructing the attribute sets from high-quality metadata provides a best-case scenario for metadata quality and consistency, although this approach is not feasible for larger data-sets that could benefit from interactive pathway-based browsing, or in scenarios where museum metadata is rich and readily available but it's intractable to manually index and construct a formal context. In this thesis, the *The Brooklyn Museum Canvas* (Section 4.4) case study circumvents this restriction, where metadata for their respective collections were automatically extracted from publicly available datasets.

4.2.3 User experience evaluations

Two separate and independent user-experience evaluations were conducted for the *Virtual Museum of the Pacific*. The first evaluation was an internal stakeholder evaluation where museum professionals were asked to assess the *Virtual Museum of of the Pacific* as a potential e-learning tool for collection access and annotation. The second evaluation was an independent, external stakeholder study consisting of key demographic groups that identified usability issues, user requirements and future design suggestions.

Both evaluations were conducted as semi-structured interviews setting where participants were required to perform scripted tasks within the *Virtual Museum of the Pacific*. As they were completing these tasks, participants were asked to report on their experiences using the 'think-aloud' usability protocol. Although both evaluations examined the content of the *Virtual Museum of the Pacific*: i.e., the objects and the way that they were classified, the internal evaluation reported on issues of collection access, representation and how the *Virtual Museum of the Pacific* represented a platform for its Pacific collection, whereas the external evaluation focused on usability, design, aesthetics and content.

The results of the user experience evaluations are summarised in Table 4.3. Both evaluations identified common themes: one of the more prominent issues concerned the accessibility of the site and the relatively steep learning curve of its interaction paradigm – namely its implementation of the *conceptual neighbourhood* metaphor, and in some cases, the way results were presented in its 'Objects Like This' feature. From the internal and external user evaluations, both cohorts recommended that the interface implemented a form of contextual help in order to improve its learnability. Other recurring themes included site aesthetics and colour schemes, the Anglo-centric design and representation of the Pacific collection and the learnability of the concept lattice metaphor for browsing and collection access. Overall, both evaluations conclude that, given that the learnability issues are addressed in the form of contextual help or providing more conventional forms of navigation, the *Virtual Museum of the Pacific* provides the capability to complement physical exhibitions, or provide access to objects that are otherwise not on display to key stakeholder communities.

Results of internal evaluation from key staff members within the Australian Museum

16 staff members from the Australian Museum, with backgrounds of collection management, education anthropology and senior management were asked to perform scripted tasks within the *Virtual Museum of the Pacific* and then answer a series

Table 4.3: A summary of key findings from the user experience evaluations for the *Virtual Museum of the Pacific*.

Internal evaluation: museum staff and stakeholders
Positive Aspects
Use of visuals, and photographs makes the collection more accessible for audiences. Creator communities can participate in knowledge creation via social tagging.
Negative Aspects
Perspectives are too closely aligned with the museum's view on the collection.
Recommendations
Implement non-hierarchical ways to keep the visitor interested and engaged. An introductory tutorial to familiarise users with the interaction approach. Perspectives should be aligned to the provenance and language groups of tags.
External evaluation: Pacific islanders and researchers
Positive Aspects
Overall design and layout. Detailed object descriptions and clear photographs. The 'Objects Like This' feature. Noted that the application was very good at 'grouping' objects.
Negative Aspects
Difficult to learn, especially given that it's unfamiliar interaction approach – too many navigation options. Difficulty using the search and understanding 'clustered' search results. A sense of feeling 'lost' within the collection.
Recommendations
The use of a 'back' button to easily backtrack through navigation sequences. A way to easily show culture and provenance among the objects. The ability to pan and zoom on object photographs.

of open-ended interview questions [123]. The interviews, each lasting from 60 - 75 minutes, began with a scripted demonstration of key features of the *Virtual Museum of the Pacific*. Each participant would then need to repeat the demonstration of features, but on a different data-set so that participants would learn and gain a 'hands-on' experience with the app. In both cases, the demonstrations were scripted and replicated for all participants in order to ensure validity. After the demonstration, participants were then asked the following questions:

- What are/were your initial impressions of the *Virtual Museum of the Pacific*?
- Do you expect to make regular use of the *Virtual Museum of the Pacific*?
- If so, how would it be useful to you?
- Are there any specific comments you would like to make with respect to the way you navigate within the *Virtual Museum of the Pacific*?
- How do you see the *Virtual Museum of the Pacific* developing as a museum collection management tool?
- Can you identify some key issues that may need addressing in the further development of this as a digitised social media platform?
- How do you see the relationship between the *Virtual Museum of the Pacific* and the traditional exhibition, i.e., the real and the digital object, unfold?
- What are the strengths and weaknesses in having objects digitised as opposed to displaying objects within its physical context?
- How does the *Virtual Museum of the Pacific* impact on ideas about learning, computer interaction, museum visitation and objects?
- How would you describe the *Virtual Museum of the Pacific* in terms of audience access?

These questions targeted usability issues and the potential for the *Virtual Museum of the Pacific* to be adopted as a platform for collection access and research. The objective of the evaluation was to determine whether the *Virtual Museum of the Pacific* could effectively represent a virtual domain of the Australian Museum's Pacific collections. Some of the major challenges for museums using virtual platforms to create meaningful learning experiences, as identified by museum staff, are connected with audience expectations, accessibility and the textual and visual presentation or layout of the content. Others relate to addressing cultural specific protocols for the public display of objects.

Participants stated that the system needs to implement non-hierarchical ways to keep the visitor interested and to acknowledge the short attention spans of its users. The implementation of the *conceptual neighbourhood* paradigm was designed specifically to present the collection in an associative, non-hierarchical way, and the use of animations and visuals within the interface, described in Section 4.2.1, were intended to aid users in forming a mental model of conceptual navigation. However, despite this, participants noted the relatively high learnability of the navigation paradigm, compounded by its unfamiliarity and lack of contextual help. They recommended a short introductory video, demonstration or some other form of contextual help, such as floating help labels, to help acquaint users with the navigation style.

In this case study, the metadata from the *Virtual Museum of the Pacific* was sourced from the object labels and descriptions provided by the projects anthropologist, which were in turn sourced from the Australian Museum's thesaurus. Within the user interface, participants noted the visual cues and relationships towards the Australian Museum's data model, and emphasised the importance of embodying multiple languages and representations. They stated that the current terminology does not reflect the social requirements of 21st century audiences, and the pre-defined categories or 'perspectives' were too closely aligned with the museum's view on how objects should be represented and contextualized. While Formal Concept Analysis can provide dynamic, conceptual structures of objects based on their attributes, this result may suggest that it has failed to properly embody a wider cross- cultural representation of the collection. However, it further highlights the importance of embracing alternative ways of annotating objects, where social tagging has proven to proven to contribute new knowledge to collections [39, 124]. The conceptual navigation induced by social metadata could provide stakeholder generated object descriptions, which in turn are supported by a conceptual structure generated from the synthesis of social meta-data. The results suggest that perspectives could be redefined to filter attributes according to the provenance or language group of the social tags, rather than reflect the predefined categories of the Australian Museum's thesaurus.

Participants stated that the *Virtual Museum of the Pacific* would work best in an exhibition context as an interactive tool and would be suitable for hosting blogs and discussions concerning the objects and the site itself. Furthermore they stated that the use of visuals and animation to convey navigation rather than the use of hierarchical text labels would make the *Virtual Museum of the Pacific* more appropriate and accessible to its audiences, and that the use of Formal Concept Analysis and its non-hierarchical means of exploring an otherwise flat (or near-flat) taxonomy of objects partially achieves this objective.

Overall, participants felt the *Virtual Museum of the Pacific* could make a useful addition to the museum's representation practices and knowledge generation, once the issues of terminology and accessibility are resolved. Furthermore, the need to make parts of the collection visible to the public that otherwise would remain in storage was almost unanimously pointed out as a progressive museum strategy and in accordance to international developments. Here, the possibility for creator communities to actively participate in the knowledge creation opens not only up to the preservation of traditional knowledge for their respective communities but provides the museum with the opportunity to update mostly 19th century annotations and descriptions.

Results of external evaluation from key demographic groups

An external usability evaluation conducted by an independent usability consultancy gathered 11 participants to conduct a user experience evaluation of the *Virtual Museum of the Pacific* [125]. The 11 participants were selected from key demographic groups as identified by a screening questionnaire: 4 participants identified themselves to be of Pacific Islander heritage, 3 participants identified themselves as researchers, librarians or academics, and another 4 participants were identified, via a screening questionnaire, as regular users of Facebook, Twitter and or other social media platforms. The participants were guided through a series of semi-scripted interviews lasting up to 75 minutes each where they were given hands on interaction and experience with the *Virtual Museum of the Pacific*. The evaluation entailed usability, user acceptance and requirements gathering components. The objective of the evaluation was to identify fine-grained usability issues and to seek requirements, feedback and opinion from representatives of key stakeholder groups. Qualitative comments and observations from the transcripts were independently analysed under the following categories: *overall impressions*, *instructions and orientation*, *navigation and user flows*, *content* and *aesthetics and visual treatment*. The following summarises the results of all evaluations under each category, each proceeded by selected quotes from participants.

Overall impressions

> "I think it's really good – it gives specialised information."
>
> "There should be more of this kind of thing – but there's room for improvement."
>
> "It's a great idea – coming from a Fijian background, some of it looks familiar to me, but I'm learning things I didn't even know."

Overall, the design and concept of the *Virtual Museum of the Pacific* was well received by the majority of participants, especially those with a Pacific Islander background. Along with the researchers and librarians, these users often praised the detailed object descriptions and clear photographs. Being able to search for objects based on similarity, such as the 'Objects Like This' feature described in Section 4.2.1, was particularly well received by the participants. However, despite these positive comments, there were a number of common issues, mostly relating to the unfamiliarity of the interface and the lack of contextual help.

Instructions and orientation

> "I wouldn't need help – it's not hard to use – you can learn by clicking."
>
> "I find it hard to navigate ... you have to use trial and error."
>
> "People could learn. You have to learn by clicking."

Although some participants found the *Virtual Museum of the Pacific* easy to learn thanks to the visual cues employed in conceptual navigation, some had not fully grasped or misinterpreted its operation. For example, many participants simply opted to use the keyword search, rather then the attribute clouds and perspectives. The issue, which was mirrored in the first evaluation, was that the navigation methods were unorthodox and required instruction, orientation or demonstration, and it was again recommended that contextual labels on the navigation elements should be put in place so that users can easily familiarise themselves with the controls. Some participants also demonstrated familiarity with the navigation options once they were given time to experiment with the controls.

Navigation

> "It's good to start ticking boxes to find what I'm interested in then narrow my search."
>
> "Browsing the objects is not intuitive and normally you'd want to find out how to browse the collection – not just search for one thing."
>
> "I like this ... it blocks things into meaningful categories."

With respect to navigation architecture and user flows, some participants had understood the concept of perspectives as facets, with many recognising that the combination of selecting perspectives, and then using the tag cloud to refine their view on the objects as a natural searching behaviour. Many participants noted that the VMP was very good at 'grouping' objects and some noted that the upper and lower tag cloud elements (highlighted 'A' and 'B' in Fig. 4.11) allowed them to find

objects even if they didn't know the correct terms. Overall, through the use of perspectives and tag clouds, the design allows for great flexibility by allowing users to search and browse, but it also provides many opportunities for users to become confused or to misinterpret their function. A recommendation was considered to implement a more constrained approach by reducing the number of navigation options provided within the application, with the option for users to switch between them.

Users also had difficulty understanding the notion of conceptual navigation as a way of navigating an information space, rather than a fixed hierarchy with a well defined 'home' state. For example, Brooklyn Museum's Collections Online site presents users with a well-defined home page with the option to navigate into specific categories of pre-defined collections. The *Virtual Museum of the Pacific*, however, presents a different approach where concept lattices place objects into multiple inheritance hierarchies. When were subject to a conceptual, rather than categorical representation of an information space, many participants reported that they felt 'lost' when navigating the collection. A recommendation was put forward so that users could at least back track through the navigation sequences (in the form of a 'Back' button) or that users could easily go back to a 'home' or 'reset' state.

Participants who identified as Pacific Islander highlighted the importance of being able to browse by region or province. For them, this seemed like a natural starting point and it was common to see participants enter keywords such as 'Fiji', 'Papua New Guinea' and 'Maori' as a way to begin their search. While it was possible to search for objects using these terms, feedback from participants suggested that more emphasis should be placed on an object's region, and as such it should be represented in a different way at least at the interface level. Participants wanted a way to easily and graphically compare both culture and provenance, and it was recommended that a map widget be placed within the user interface so that users can graphically browse across locations and cultures as a complement to the conceptual navigation provided by the interface.

The search and query extension feature of the *Virtual Museum of the Pacific*, described in Section 4.2.1, was intended to provide an entry point into the collection. When a user would search for an object, the interface would display search results for matching formal concepts in addition to results for individual matching objects. By matching search terms to a formal concept, it would allow the user to navigate to that concept and explore other concepts using the attribute clouds. These were presented within the interface as 'clustered' search results, however, many users had difficulty understanding what they meant within the interface: particularly for clusters that weren't directly related to the search string, but were somehow conceptually related. For example, the search "objects made of bone" also returned

objects that were made of tooth, and many users couldn't understand why or thought it was not related to their search. A recommendation was made to provide a more traditional layout for the display of search results with an optional feature to display matching object clusters.

Content

> "I think it's really good – it gives specialised information."
>
> "I want to learn about the objects – not the philanthropists."
>
> "Nice for people who may not be quite on the right one." [15].

The imagery, photographs and description of the content within the *Virtual Museum of the Pacific* were positively received by the participants, particularly for the Pacific Islander and researcher cohorts. Users positively commented that both physical and interpretive descriptions of the objects were provided, although some wanted more detail in the object description fields. Many users also found the 'Objects Like This' content-based retrieval features (Sec. 4.2.1) to be quite useful, although some users preferred to view results presented as individual objects, rather than as 'object clusters.' A highly requested feature was the ability to pan and zoom in on the object images, given the importance of patterns and carving details within Pacific material culture.

Aesthetics and visual treatment

> "The aesthetics are okay for student research – not enticing, but intelligent."
>
> "The Pacific is very colourful – not black and grey ... it's sun, sand and sea. Make it brighter and more colourful."
>
> "It's a bit hard to read the text it's a bit small."

Given its importance and appreciation in information systems design [126, 127], issues of aesthetics, look and feel were considered in the design of the *Virtual Museum of the Pacific* and addressed in its evaluation. Its interface was given a simple treatment for its design to emphasise the objects rather than the UI. Participants were critical of the dark colour scheme, the placement of blackspace, and the small size of text. Participants felt that the design of the site did not reflect the colourful diversity of the Pacific, and recommended to alter the visual treatment of the site to increase the size of the text font and change the color scheme of the interface.

[15] This quote was spoken in response to the user clicking on the 'Objects Like This' feature, described in Section 4.2.1

Summary

Overall, the study concludes that conceptual based browsing offers significant merit for browsing Pacific collections, and that by resolving the identified issues, the core functionality of searching and browsing via concept lattices can be considerably easier for users to learn and engage with. The study also recommends a simplification of the navigation options and an ability for users to browse by region or culture via an interactive map.

4.3 A Place for Art

A Place for Art is an iPad app that publicly displays 80 paintings, works on paper and Australian Aboriginal works from the University of Wollongong Art Collection. [16] [17] The app allows users to navigate the collection by using gesture-based interactions to navigate pathways of conceptually related content via *object-linked pathways* (Section 3.4.2). It embraces the design principles of *pathways*, *pliability*, *visual imagery*, *visual momentum* and *similarity-based suggestions* and Section 4.3.1 describes their implementation. Shortly after the app was released, a user experience evaluation was conducted that correlated users' perceptions of the app with affect: i.e., whether certain features or design elements were perceived to be positive or negative. The findings of the study, described in Section 4.3.4, report the participants' experiences in terms of information seeking behaviour and their personal engagement with the UOW Art Collection but also identified a number of features and usability improvements. In terms of affect, the majority of positive comments centered on the aesthetics and visual design of the app, the participants' personal engagement with the UOW Art Collection, and the way they perceived the app to offer free-form and open-ended exploration of the artworks. Likewise, the majority of negative comments centered on the lack of more conventional options for exploration and browsing, perceived lack of fine-grained control over its navigation and the sense that the app does not provide an overview of the collection.

As a case study, *A Place for Art* examines how an exploratory visualisation with a highly tactile, tightly-coupled and malleable pathway-based interface can facilitate discovery and serendipity in a collection of artworks that, for most participants and target users of this app, surround them in the built environment of the university campus environment. The app is also a digital counterpart to an exhibition book – also called *A Place for Art* – that uses the same artworks but exhibits them in a completely non-linear way. Key elements that were observed from the user experience evaluation of *A Place for Art* centered on this aspect of non-linearity vs. a curated experience, and how its design fosters serendipity, knowledge discovery and meaningful engagement with a publicly accessible art collection.

4.3.1 Design and implementation

A Place for Art was conceived as a digital companion piece to a print publication

[16] The app was released to the general public in 2013 and is freely downloadable from the app store at `https://itunes.apple.com/au/app/a-place-for-art/id638054832?mt=8`

[17] A video that demonstrates some of the gesture interactions of *A Place for Art* is available at `https://vimeo.com/55439706`.

of selected works from the University of Wollongong Art Collection. [18] The original publication was produced to showcase key pieces of the collection along with the personal stories of staff and students of the university as they interacted with the public art on campus. The app presents high resolution images of these works, along with their textual labels and, for the works that are publicly available on campus, their location. The app was developed for the iPad: the choice of screen size and form factor of the tablet device allowed the user to view its high resolution imagery and textual labels simultaneously while providing a large enough surface for them to visibly interact with the app using broad touch and swipe gestures.

The presentation of the works as a visible, sliding pathway is a key design feature of the app, and is presented as its primary means of interaction, shown in Figure 4.16. The pathway is depicted as a horizontal lateral gallery of images that represent thematic concepts derived from the metadata of the artworks. Users can move back and forth along the pathway by swiping left and right with their finger, or they can tap on an individual work to view a larger image and read its label, as shown in Figure. 4.17. When a user first starts the app, they are presented with a randomly selected pathway of works. As the user swipes the pathway, it's 'title' – a natural language representation of the attributes of the formal concept that represents the group of works – is fixed and always remains in view. The prominence of imagery within the app follows key design principles of serendipitous [74, 91] and 'generous' [50] interfaces through its emphasis of large visual imagery.

Each pathway within the app represents a single formal concept, with its objects represented by the artworks, and its attributes represented by its title. The app employs the use of *object-linked pathways* where a user can perform a two-fingered tap on a work to 'displace' that work in a new concept which – according to the mechanisms of *object-linked pathways* – is the most similar formal concept that contains objects that the user has not yet seen. (see Section 3.4.2) The interaction that follows is highly fluid and animated: in order to convey to the user that they are 'displacing' the work in a new and renewed context, the app animates the pathway in such a way that the work appears to be physically 'popped out' while the previous path rotates 90 degrees behind it, forming a 'T' shape shown in Figure 4.18. It is then joined to a new path that slides in from the right. The work that was tapped on now appears at the beginning of a new pathway, in which at that point the user can explore further or 'displace' other works. Figure 4.19 shows an example of a displacement: on the left is the previous state of the pathway before displacement, and on the right is the new pathway that's formed perpendicular to the previous, noting that the 'displaced' work now appears to be partially overlapping the previous

[18] Details on the print publication of *A Place for Art* can be found at: `http://lha.uow.edu.au/taem/uowac/UOW132043.html`

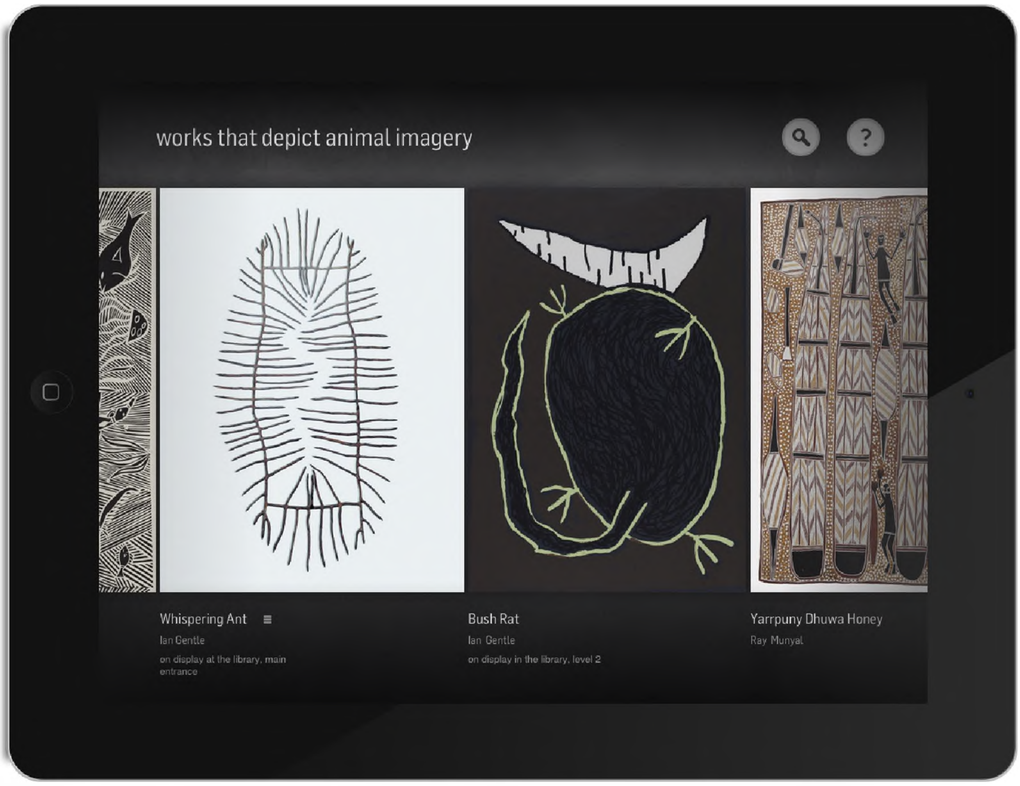

Figure 4.16: *A Place for Art* presents the collection in the form of pathways: galleries of laterally arranged images.

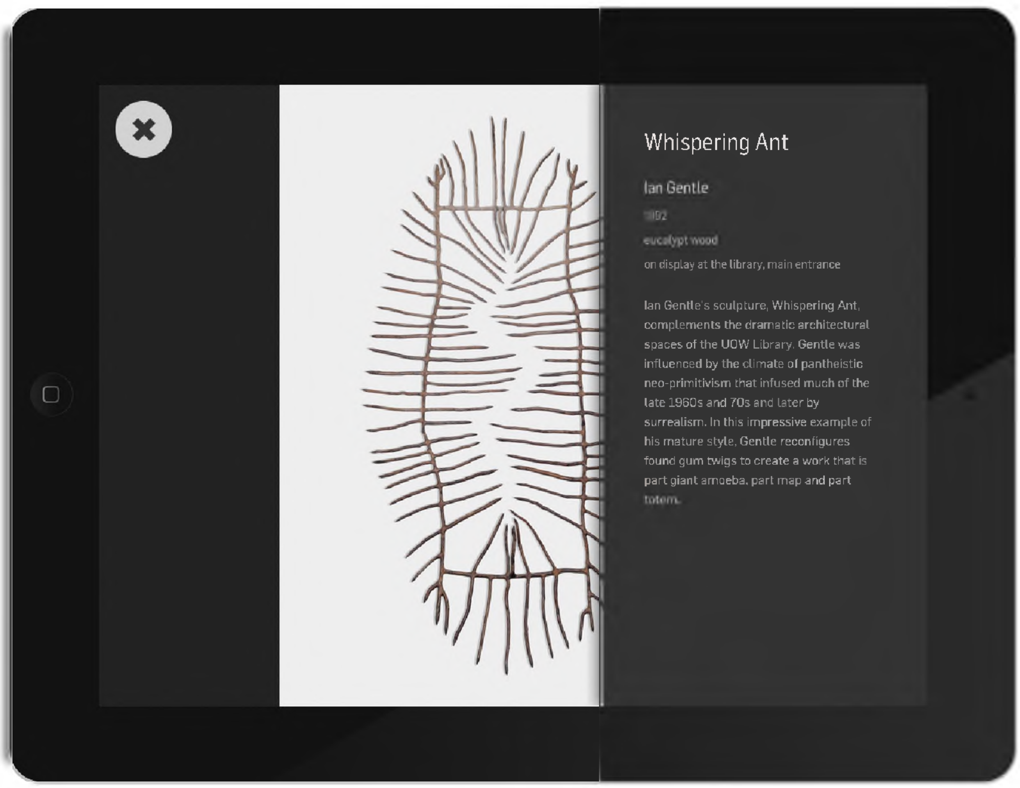

Figure 4.17: A full screen view of an object in *A Place for Art*.

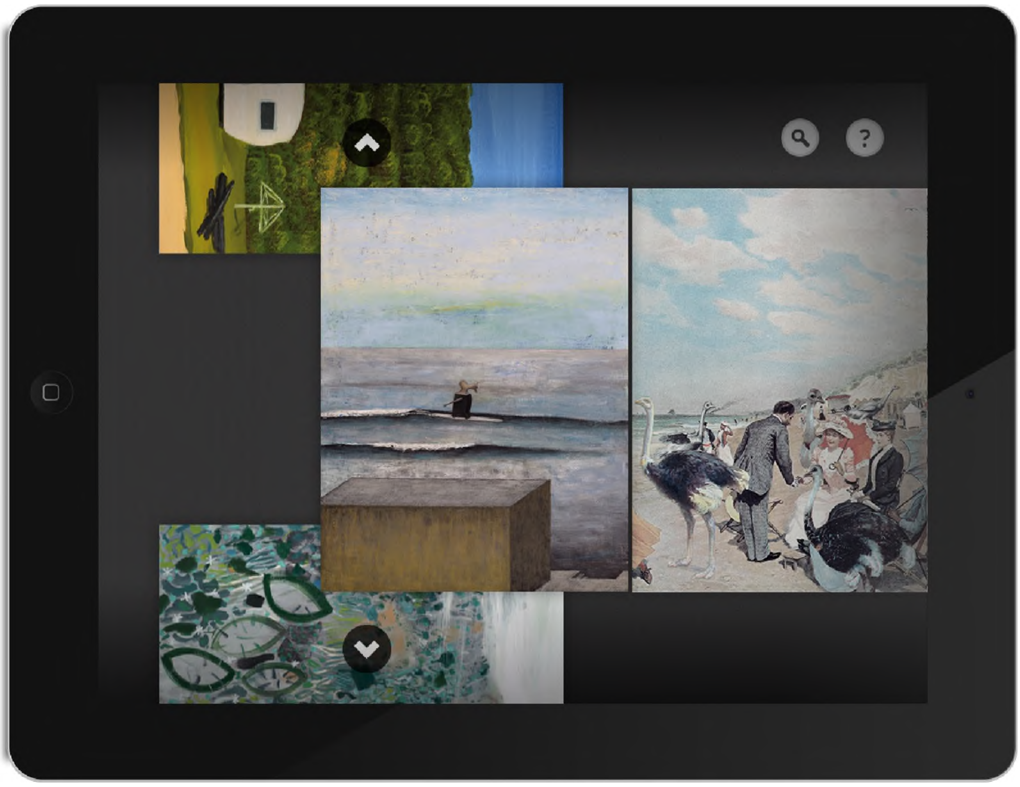

Figure 4.18: A 'displacement' in *A Place for Art* that shows the intersection of two pathways.

pathway. Also of note are the differences between the titles on the left and the right – before displacement it was part of a collection of "works by indigenous artists" and after displacement it is repositioned into a new context of "intricate works that depict nature." The displacement process not only acts as a surrogate for similarity-based content retrieval, but was a way to elicit the multiple contexts that a single work can exist in. For example, if the user were to displace the striking sculpture in Figure 4.16 – Ian Gentle's *Whispering Ant* – it would displace it from its current context of "works that depict animal imagery" and juxtapose it alongside "other abstract sculptures." A displacement produces a new pathway of works that, just like the one previous to it, can also have its works displaced as well. This can result in multiple, cumulative displacements that branch from one another.

These multiple, connecting, displaced pathways represent a user's 'pathway history.' A user, if they wished to do so, can go back to a previous pathway by swiping back to the point of displacement, and tap on the previous path that was rotated outwards by 90 degrees. When the user scrolls all the way back to a previously displaced pathway, small arrows appear above and below the current pathway, as shown in Figure 4.18. These arrows indicate to the user that they can 'go back'. Tapping on the previous pathway causes the interface to animate a displacement in reverse, where the current pathway disappears and the previous pathway rotates

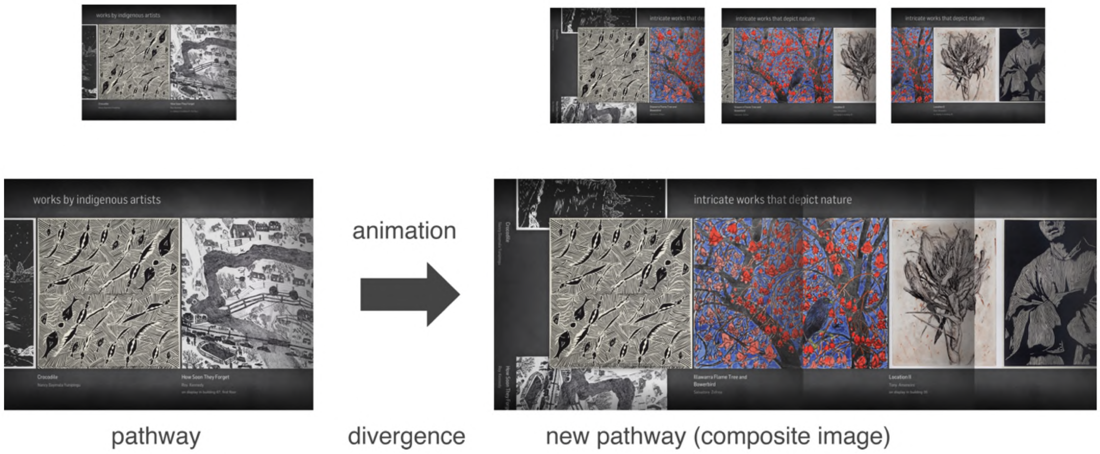

Figure 4.19: A displacement in *A Place for Art*, showing the before and after state. The use of animation and contiguous forms are intended to help users form mental models of state changes.

back to its original position, causing the previous work to 'fold in.' The intention of this process was to engender a sense of reversibility, given that one of the main negative points raised by participants from the *Virtual Museum of the Pacific* evaluation was the presence of multi-directional navigation options with no clear way to 'go back'. The design of *A Place for Art* opted a more minimalist approach, where directionality was reduced to 'forward at this object' or 'back to the previous pathway.'

Multiple displacements, over time, can also reflect emerging themes in the collection, such as the series of displacements in Figure 4.20. As the user creates more displacements that unfold different parts of the structure, they would thematically traverse the collection in a non-linear, self-determined manner. Over time, users unravel implicit phenomena that connect the objects and gain a sense of the thematic concepts within the collection.

4.3.2 User experience evaluation

A user experience evaluation was conducted that targeted participants who demonstrated an interest in visiting museums and galleries and had experience in using a smartphone or tablet device. Participants were selected from a diverse range of backgrounds, with the majority of them having expertise or experience in graphic design, visual arts, psychology, curatorial practice and digital media. In all, 24 participants – grouped into pairs of 12 – were recruited to conduct the study. Most participants were under the ages of 35, and more than half of them were under the age of 25. Approximately half of the participants reported visiting museums and galleries less than three times a year, although a quarter reported visiting more than

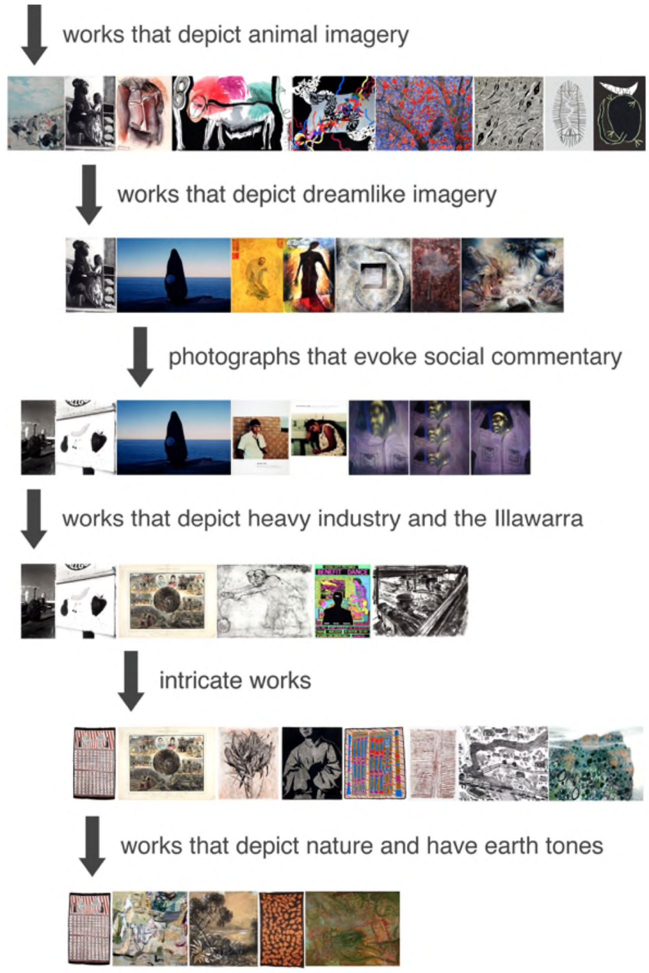

Figure 4.20: A series of six displacements depicts the movement through the conceptual space from "works that depict animal imagery" to "works that depict nature and have earth tones."

once a month.

The evaluations were grouped into two three parts: first, each participant completed a survey that collected data such as their age, professional background and experience with museum and collections apps. [19] In pairs, the participants would then freely use the app and and describe their experiences, in which the interviews were then transcribed and coded. Finally, participants would respond to a post-evaluation questionnaire that consisted of the following questions:

- If you were to think of one word that describes your overall impression of the app, what would that be?
- What was your impression of the interface and the navigation concepts, what seemed natural and/or unnatural with the interaction?
- How would you describe your overall user experience? Was there a sense of flow or serendipity in your interaction?
- What did you learn by using the app? What difference did the app make in terms of your awareness or knowledge of the UOW Art Collection?
- What comments can you make about the way information was presented?
- Can you provide any more comments on how this app shapes your awareness of or experiences with the University of Wollongong's Art Collection? What does this app mean for you in terms of your engagement with the art on campus?

The results, in Section 4.3.4, describe the key themes that are derived from the coded analysis of the evaluation and questionnaire, along with selected quotes to describe these themes in detail or to provide individual insights by the participants. In line with the app's design principles, the following factors informed the evaluation approach:

- The app does not have specific functional goals for the user to achieve: it is entirely concerned with exploration and experience. This introduces new factors into the analysis.
- The app employs non-traditional techniques for user interface engagement; this becomes an issue as we seek to see how users attempt to interact with the system.
- The outcomes of the analysis are not primarily concerned with the number of mistakes made by the user or with interface issues. Rather, the focus is gaining insight into the user experience and flow through the collection.

[19] The pre-evaluation survey and results are available in Appendix A.

The chosen evaluation approach of using pairs to evaluate *A Place for Art* was based on research on computer-supported cooperative work [61–64, 128], as reviewed in Section 2.3. When applied to human computer interaction evaluation methods, Downey [63] discusses a formative usability testing method in which, in some cases, pairs work through a process of co-discovery that can result in a greater number of system issues being identified. Similarly, Miyake [62] describes a paired-evaluation approach called *constructive interaction* as an attempt to understand how people interact with complex physical devices. Miyake argues that having two or more people interacting with each other during a problem solving exercise removes the unnatural aspects usually associated with the think-aloud method, as in normal circumstances a participant would be silent as they interact with systems. Having more than one person using the system means that processes that are usually invisible become visible, as the pair not only explains what they think but also why they think certain things. The research identified the need for participants to interact with a single iPad and pass it amongst one another in an authentic use environment, and that participants would need to know one another so that they shared enough *common ground* to build a constructive critique of the interface.

4.3.3 Results of pre-evaluation survey

One of the key objectives of the study is to determine how *A Place for Art* facilitates the awareness and diversity of the UOW Art Collection for those who demonstrate a strong connection and familiarity to the campus and the collection as compared to those who are not from the area and otherwise unfamiliar with the collection. Based on the questions described in the survey, 16 out of the 24 participants identified themselves as students, staff members or faculty from the University of Wollongong and all reported demonstrating some sort of engagement – or at least awareness – with the public artwork that surrounded them on campus. Out of these 16 participants, 12 reported "noticing the artwork around campus and wanting to know more about the work", and 10 stated "that at least one point in time they felt curious by a particular work that caught their attention, and wanted to read the label." 11 of these participants also expressed a desire to "know more about the artists behind the works" – a finding that was further reinforced by a prominent suggestion from the majority of participants that *A Place for Art* should show more information about the artists. Although the remaining 8 participants were not from the University of Wollongong, 6 out of those 8 participants reported that "the app could be used to highlight the awareness and value of the University of Wollongong Art Collection."

The majority of participants reported that they have accessed an online museum collection before with the majority of them reporting that they did so using a Web

browser on a desktop, smartphone or tablet device. Only a relatively small proportion of participants reported downloading an app to browse or explore a museum's collection. Furthermore, most participants accessed – or found out about – an online museum collection via a Google search of an individual object (13 participants), via the museum's website (10 participants), or as shared on social media, such as Twitter or Facebook (8 participants).

Participants were also asked of their reasons and preferred technology platform for accessing online museum collections. The majority of participants reported that they would "use an online museum collection to learn about artists they are already familiar with" or "as a way of discovering new works". Participants also reported that they would "serendipitously explore the online collection", "access more information about works seen in an online exhibition", or "access artists' perspectives on the works". The majority of participants stated that they would prefer to access an online collection via the Web from a desktop or laptop computer, and more than half stated that they would also prefer to do so via an app downloaded to a tablet device. Some participants reported that a larger screen is better for viewing the works, with one noting that technology platform of choice depends if it is used for research or pleasure: they would use a desktop Web browser if they were conducting research so that they could take notes as they went, but would prefer "an iPad with a comfy chair" if they were browsing for pleasure and leisure.

The full results of the pre-evaluation survey are presented in Appendix A. Due to the small sample size, these findings were not intended to explain how audiences at large would interact with online collections. However, they do provide insights to their responses and feature suggestions for *A Place for Art*. Throughout the course of the study, the participants were asked to freely experience and interact with the app with no further instruction or set task goals. They were encouraged to report on usability issues, provide constructive feedback, or describe how the app facilitates their engagement with the UOW Art Collection.

4.3.4 Results of user experience evaluation

Overall, the findings describe the information seeking behaviours of participants as afforded by the app, and how its presentation of the collection engages participants in a meaningful away. Rather than focus on traditional usability issues such as efficiency and task completion, the results are presented in terms of affect, emotion and motivation: i.e., the feelings felt by participants as they were experiencing the app. This is based on the idea that a greater emphasis on pleasure and leisure should be considered in designing user experiences for digital products and devices: [13] a factor worthy of consideration when crafting user experiences for online mu-

seum collections. The results of the user experience evaluations are summarised in Table 4.4.

Emotion and motivation

Participants experienced a range of positive experiences with the app, such as appraisal or particular features: "I really like that it is got where you can view the works."; expressions surprise and delight: "It's good that you can whirl and twirl. It is quite fun."; meaningful engagement with the collection: "I enjoyed that one how there's a story around it" and expressions that highlight intuitive or natural interactions: "It's good that there are many gestures that are quite natural to people who use an iPad." Participants also reported a number of negative experiences, such as missing features, functionality or content: "It's upsetting me that sizes are not represented"; feelings of obstruction and disorientation: "There's no navigational tools, you can't see where you are"; feelings that the app presents the collection in an obscure or unorthodox way: "How's that similar to that? They are so different!" and feelings of boredom: "I think we exhausted the app's functionality and we're bored." Throughout the evaluation, the majority of positive comments were centered on the way the app affords exploration of the collection, its meaning making qualities, its aesthetics and visual design. Likewise, the majority of negative comments were based on a number of usability issues identified by the participants along with other aspects of the information seeking experience that were considered limited, obstructive or unorthodox.

As *A Place for Art* affords exploration of the collection primarily via the means of pathway creation and exploration, we expected participants to embrace this as their primary means of interaction. Participants in 10 out of the 12 evaluations demonstrated exploratory behaviour using the pathway structures, with two noting that the presence of a 'search function' within the app "felt out of place." In 6 of the evaluations, participants spent a large amount of time interacting with the works on the first pathway that was presented. In 3 out of the 12 evaluations, participants spent much of their effort in using the app's search box as a way of exploring the collection by repeatedly entering search terms that they are interested in, or thought they might know are in the collection. This indicates that for certain participants, the more familiar search box – despite its apparent limitations [3, 50] – may still be employed by participants as a way of accessing or exploring the collection.

Information seeking, exploration and meaning making

Throughout the evaluation, a number of dimensions were noted that relate to information seeking when participants were engaging with *A Place for Art*: exploration,

Table 4.4: A summary of key findings from the user experience evaluation for *A Place for Art*.

Positive Aspects
The ability to view the physical location of the works as they appear on campus.
The ability to view the story and history behind some of the featured works in the app.
The aesthetics, visual design and 'flow' of app.
The ability to discover connections among the works and the 'categoryless' nature of their presentation.
Fosters a greater awareness of the art collection that surrounds students and staff on campus.
Negative Aspects
Dimensions of the works are not shown / there is no sense of scale with the works.
For some participants, feelings of obstruction and disorientation.
For some participants, a lack of overview or 'starting point'.
For some participants, highly subjective or unusual representations of pathways and their connections.
Unable to zoom in on the artworks using pinch gestures.
Recommendations
Display the sizes and scale of the works.
Provide an overview of the collection.
Show more information about the artists and 'connect' to them via social media; view other works by the same artist, etc.
Provide more conventional navigation options, such as viewing a list of artists, browse by medium, etc.
Show a visible map of the campus with the locations of the artworks.

meaning making and search. These dimensions – derived from this analysis – are modeled on the constructs identified by Skov [37] in their study on information seeking behaviours in online museum collections.

Exploration

Most of the comments made by participants in relation to exploratory behaviour described how the pathways in *A Place for Art* afforded creative, non-linear exploration of its collection, and how they relate to the notion of the information flaneur – the notion of the creative, curious and critical information seeker [15].

In relation to how the pathway structures afford creative exploration, the majority of positive comments centered on the way they convey a sense of 'movement' through the collection, and in particular, how they empowered visitors to discover connections among the works. Some participants reported that they could get a sense of the collection's content and structure entirely through the progression of 'unraveling' new pathways. One participant reported that in order to gain an overview of the collection, it was necessary to explore: "I wanted to use it more and figure out where the connections are." Others hinted that pathways convey a sense of transition or connectedness through the collection: "It actually does feel connected ... it gives a sense of transition It is not what I expected to see but it works really well."

In some instances, participants briefly discussed how much freedom the visitor should have in determining their own path through the collection – whether they should be guided, or follow a more self-determined approach. In one case, the participant was expecting the app to be much more linear in the way the collection unfolded: "I was expected to be taken on a journey through the app." Another remarked that the app empowers the visitor with a higher level of self-determination, noting that while the visitor should be somewhat guided by the experience, they should still have a certain level of control in where they were going: "I suppose there are different ways – like there is a self directed journey, or like a 20 minute video that starts." In another discussion, the two participants highlighted the dichotomy of immersion within the collection vs. 'getting lost' and not having enough control:

> PARTICIPANT A: "There's a whole lot of ... pathways that I want to ... explore now that I did not realise where at the beginning. So there's a whole lot more depth that I did not realise ... "
>
> PARTICIPANT B: "I do like how everything does ... how everything links to something else. It is so thoroughly done. It is really clever."
>
> INTERVIEWER: "So regarding the actual concept of navigating by collection pathways, can you comment on how you feel about that?

> Whether it is positive or negative?"
>
> PARTICIPANT B: "I definitely think it is positive. I think it allows you to look at things in a different way rather than have it set out in a collection where someone's got a room and they've got a wall. You go into a gallery and they got this, this, this, this, this, this - it is the order of the collection - so, by being in charge of that, you can sort of take control of your own experience, and I think it is pretty good."
>
> PARTICIPANT A: "Going around tangents and, people do like YouTube related videos for example ... you choose your own story ... chose your own ... where you're going ... rather than being confined to a physical space."

Participants often held differing opinions on the way the app presented the collection in such a deliberately loosely structured manner. In one case, one of the participants was particularly responsive to the idea: "I think that's the appeal of it, being able to get lost in it and explore your way around." The other also responded positively to the idea, but noted that the pathways could induce a sense of disorientation: "Part of me wants to see it all, I'm a bit overwhelmed, and my mental map is lost, but part of me is also curious." This comment mirrored one of the negative aspects of the user experience: the fact that some visitors felt a sense of loss or disorientation when navigating pathways – particularly in relation to a loss of context, or from not being able to 'go back'. This was exacerbated by the fact that the app did not appear to sufficiently convey an overview of the collection: in 8 out of the 12 evaluation sessions conducted, participants indicated that the app did not sufficiently provide an overview. One of the participant pairs even questioned the necessity of the pathway concept in the first place, and opted for the works to be instead presented as a single, continuous stream.

The persona of the *information flaneur* was brought to life through a discussion between a pair of participants who described their experiences in using the app, noting that there was no mention to any of the participants of the design theory and principles for *A Place for Art*, nor there were any mentions of triggering words such as 'flaneur':

> PARTICIPANT 1: "You navigate by meaning ... It is a nice way to navigate the works in a collection. It is more flaneur-like ... "
>
> PARTICIPANT 2: "It would be nice to be able to explore more than one pathway from a given point, at a given time ... Because if you find the same thing in different ways you could end up on pathways."

> PARTICIPANT 2: "So you would rather have an option to have it go on, rather than have it just take you down one way ... and then ... "
>
> PARTICIPANT 1: "... or have it give you one way but then perhaps give you a way to go back and try a different one, if there is more than one available."
>
> PARTICIPANT 2: "I actually prefer it this way I think."
>
> PARTICIPANT 1: "You prefer the option of not having to go back and change your mind?"
>
> PARTICIPANT 2: "Yeah. I actually prefer that it is in one direction. That there isn't ... That I don't have to make that decision. I like that it actually meanders ..."

Many other participants alluded to the *information flaneur* in a similar way, particularly for those who demonstrated a strong connection to the University of Wollongong Art Collection. Some participants claimed that the pathway metaphor and the aesthetics of its interaction played a large role in fostering curiousity and exploration:

> "It's a lot easier to peruse and go through ... it encourages the passion in me."
>
> "You can explore so easily. It is really fun. Just taking a peruse."
>
> "You definitely get a sense of serendipity, and just being encouraged to meander ... moving through a spiral, as we were unraveling our journey."
>
> "It really does remind me of a museum. You just walk around and look for things."

Participants observations and experiences reflected the key design principles of *serendipity* and *visual momentum*, describing the pathway animations and transitions as "moving through a fluid" and being "very pleasing on the eye":

> "I really like the movements in the whole app ... I really like the way that it moves just seamlessly"
>
> "It's a really nice animation going from one pathway to another. It is kinda like a Rubiks cube."
>
> "I like the movements for navigation, how it goes to a different orientation, how it rotates 90 degrees."

Participants also noted other serendipitous aspects of the user experience, such as the way the app induces curiousity about the works and the connections among them: "I definitely felt intrigued by it"; the way it introduces the unexpected: "There was a cool unexpectedness – I did not expect that coming."; and the way users can diverge or create new paths: "the option to jump feels natural ... something that made it obvious that you wanted to explore." Overall, participants responded positively to the way *A Place for Art* affords creative exploration of its collection, although participants in 5 out of the 12 evaluations did indicate that they also felt a certain sense of loss or disorientation while exploring the pathway structures. Participants stated that the ability to navigate by more traditional means, such as browsing through a list of artists, works sorted by genre or medium or by navigating an 'overview' of the collection are likely to resolve these issues of loss and disorientation.

Meaning making and engagement with the UOW Art Collection

Participants noted how pathways were used to structure content and show the relationships among the works within the collection, and most demonstrated strong engagement with the works presented in the app – particularly in relation to their place on campus and their own experiences of these works throughout their day-to-day life. Participants noted that the pathways within the app grouped objects into multiple, overlapping contexts, and many of them made comments about the way these pathways described the objects. Their meaningful engagement and interpretation of the UOW Art Collection was consistently denoted as one of the more positive aspects of the *A Place for Art* user experience, and many participants – particularly those who work and study at the University of Wollongong – demonstrated personal engagement with the works that describe the history or the natural geography of the local region.

Participants from 10 out of the 12 evaluations commented on how pathways within the app not only afforded creative, non-linear exploration of its collection, but also how they are used to highlight the differences and similarities among groups of objects within the collection. Some participants praised the 'category-less' nature of the app and in particular, how objects were often represented in different categories or themes:

> "It was interesting to see which paintings came around again and again, as part of different themes. It gave a sense of how big the collection might be."
>
> "It's interesting now that I've started to use it after a while to see how they are tagged."

> "It's good how you can create multiple paths from any work."
>
> "I love it! It is so interesting! It is so interesting to see what relates to one another."

Other participants, however, noted the unorthodox and unconventional representation of metadata within the pathway titles, with some noting that the depiction of certain metadata terms was against convention, too subjective, or even inaccurate:

> "That's a weird theme, I would be looking at more of the styles."
>
> "Are they random?"
>
> "Some of these works really aren't that vibrant."

In 6 out of the 12 evaluations, the participants explicitly acknowledged that the pathways were grouped "by theme", although the participants in 3 of the evaluations thought that they were grouped by visual and iconographical similarity. This was the cause of contention and confusion in one of the evaluation sessions, where one of the participants thought that the pathways presented the relationships among the objects in a natural and intuitive manner as they were grouped by theme, whereas the other participant relied purely on visual similarity and iconography. When the participant attempted to create a new pathway from the work *Port Kembla, NSW, 23rd August 1997* – a photograph depicting a wall of drawn fruits with the furnaces and industrial landscapes of Port Kembla in the background – the participant expected to see more works showing fruit: "I want more works about fruit" – and were surprised to see that instead, the app presented more works about the Illawarra: [20] "why does it say more works about the Illawarra?" In another evaluation, one of the participants interpreted similarity visually and often appeared perplexed at the fact that the connections among the objects were not what they expected: "I pick something that's got clean lines, and then I get things that are abstract and messy." In 4 out of the 12 evaluations conducted, there was at least one mention by the participants that the connection and relationship between adjoining pathways were either obscure, unintuitive or inaccurate. Such observations suggest the need to employ less subjective metadata to describe and link the pathways, or match and link artworks based on their visual similarity.

The selection of works presented within the app – all from the University of Wollongong's Art Collection – received universal praise by the participants. Almost all participants agreed that the app showcases the diversity and value of the collection. Those who work or study at the University of Wollongong demonstrated a strong connection to the works that surround them on campus, and many of them wanted to locate and visit after browsing the app:

[20] The Illawarra is the name given the to local region that the university is located in.

> "I know where two of these works are - and I instantly want to go find them."
>
> "You did not even know that some of these works are on display ... Have you seen this stuff in real life?"
>
> "Even now I'm quite intrigued to go looking at some of the works."
>
> "I cannot believe all of these are at the uni? I did not know it was so art focused."

The collection also appealed to participants who are not from the University of Wollongong, or otherwise did not know anything about the University of Wollongong's Art Collection:

> "Seems like the University of Wollongong would be a cool space."
>
> "If I had some time to kill in Wollongong I would definitely check out the campus."

In 7 out of the 12 evaluations, the participants stated that they would felt more comfortable if the app conveyed an overview of its collection, and many of them were not sure at first if the app contained all works from the University of Wollongong's Art Collection, or just a small sub-selection of works. Participants felt that an app should convey the size and scale of the collection, along with a list of its main themes, mediums or prominent artists.

Participants from the University of Wollongong often demonstrated a personal connection to the works, sometimes seeing works that they were familiar with from around campus. There was a strong emphasis from participants that the app should better showcase the artists behind the work, and many wanted to access more information about the artist, or in the very least, be able to easily view more works by the same artist. Participants in 11 out of the 12 evaluations wanted to access more information about the artists, and one of the participants made the explicit observation that the app seems very 'object-oriented.'

In one of the evaluations, there was an extensive discussion on social media integration. The participants stressed the importance of being able to deep-link to artist Facebook pages, Wikipedia or external portfolio sites: "As an artist, you would want to be represented, and you would want your career to be represented as well." The participants noted that the artworks within the app piqued their interested and wanted a lot more information, and thought it would be helpful to be able to provide Web or social links to external artist profile pages. They also suggested the idea of having a Facebook 'Like' or 'Share' button associated with the works, so that, if a user 'Like's' a particular work within the app, it would appear on their news-feed, and that the user could also share the work with other users.

Instagram was also suggested as another avenue for sharing and deep-linking: "For visual people, Instagram is a really good way to go." The participants noted that the ability to deep-linking works on Instagram could allow them to visually browse more works by the artist, or similar works on that platform. In particular, one noted that the ability for artists to share and link their works on these social networks could allow them to build a profile and bring greater awareness of their works. One of the participants hinted at the idea of being able to 'tag' works, potentially as a way to add additional metadata and meaning to the artworks, although the other countered that tagging is not an as an effective form of classification as it may seem: "One thing about social media is that hashtags are not really that good – people just tag things."

Overall, the discussion on social media integration was based on the idea of connecting artists and their works to students, scholars and art enthusiasts, and for these people to learn more and forge connections with the artists: "The best way for this to work is to have it connected: you want artists connected, you want teachers connected; you want students connected – it just keeps going." The participants recognized the novelty of showing connections to the works within the app, but also underscored the importance for the potential for a user to connect elsewhere by engaging or learning more about artists on other forums. To this end, 'connecting' was an underlying theme: "Everything is so dependent on linking things and connecting things ... everything is about connection."

Search

In the evaluation sessions, the search function was used to observe to determine if participants would feel more comfortable typing in keywords than exploring pathways, and was also used as a means for participants to discuss the differences between the two interaction approaches. In some cases, the participants would resort to the app's search function if they felt uncomfortable, frustrated or otherwise obstructed by the pathway-based approach for navigation.

In almost all of the evaluation sessions, the participants interacted with the app's search functionality. In 3 out of the 12 evaluation sessions, participants used search as their primary way of exploring the collection as they would enter search terms for works, artists or genres that they thought they might be interested in. In 8 out of the 12 evaluations, participants engaged in exploratory search [79] – entering keywords and terms that span their interest or at least as a way for them to make sense of the collection. Participants would most commonly search for a number of different mediums and campus building numbers as a way of getting an overview of the collection.

In 7 out of the 12 evaluations, the participants engaged in known item searching

– searching for specific objects, works or artists that they knew previously or wanted to view. In some cases, participants wanted to be able to 'go back' to an item that they previously saw in the pathway, and if they were unable to do so, they would attempt to search for it instead by entering a description. In other cases, they would remember the name of the pathway that the work was categorized under, and use that to search for the item.

User experience issues and feature suggestions

Participants noted a number of usability issues that impeded or otherwise negatively impacted on their user experience. Most of these issues centered on the unconventional gestures employed in the app, deficiencies in the way its help system was presented and missing features and functionality. Likewise, the participants made a number of feature suggestions that could be used to improve the user experience of *A Place for Art*. This section describes usability issues and feature suggestions that have been consistently noted in at least 4 out of the 12 user experience evaluations.

All participants in all 12 evaluations reported a usability issue when browsing pathways or when interacting with the works in the full screen view (shown in Figure 4.16). In 8 out of the 12 evaluations, the participants noted that they were unable to zoom in on the works by pinching outwards with their fingers, although it was explained to some participants that they were unable to zoom in on the works due to copyright restrictions. [21] Another major usability issue was that participants were unable to create a new path when viewing the object within this view, and were unable to swipe left-to-right to view adjacent works on the pathway: another gesture that they thought felt natural and familiar to them, but one that was not supported in *A Place for Art.*

Another major usability issue, as identified in 8 of the 12 evaluations, was that the participants felt the app did not sufficiently provide an overview of the collection. Similarly, in 7 out of the 12 evaluations, the participants felt that the app did not sufficiently convey the context of the collection. Some participants also noted issues in relation the way certain works were represented within the app, with some claiming that the pathway titles used to describe them were incorrect.

Another issue, as identified from 7 out of the 12 evaluations, was that the app employed gestures that appeared to be unconventional or unintuitive – particularly the gesture that requires the user to tap with two fingers on a pathway in order to create and branch a new path. Although the app provides a tutorial system that describes the gestures within the app, participants in 4 of the evaluations actually misinterpreted or did not fully comprehend the app's help system, and as a result,

[21] At the time the user experience study was conducted, the app was unable to provide the ability to pinch to zoom on the images due to copyright restrictions.

were not certain on where they would need to perform the two-fingered tap gesture. Many of these participants incorrectly assumed that the gesture needed to be invoked in the full screen view of the object, rather than on its thumbnail in the pathway.

Due to the unfamiliarity of the pathway interaction paradigm and the gestures required to use, participants at one point in 6 out of the 12 evaluations reported that pathways actually hindered rather than promoted exploration. Participants in 4 out of the 12 evaluations reported that the collection appeared inaccessible within the app, and also noted that the app should employ more conventional navigation options for the user.

Participants also made a number of feature suggestions. Although some of these feature suggestions address the above usability issues, others have been put forward as ideal features participants would wish to see in any app that showcases a museum's collection.

All participants stated that they would like the app to offer more control over the way they navigate the collection. In particular, participants expressed a desire to browse by artist or medium: whether that would be being able to choose a selection of artists and mediums from a predefined list, or the ability to create pathways based on specific facets, such as being able to browse a list of similar works by the same artists, or works that share the same medium and materials. In this manner, participants want to be able to select the facets of their browsing experience that represent specific domains of interest. Participants in 4 out of the 12 evaluations preferred a more traditional way of exploring the collection: to simply be able to select the works, or works grouped by artist, medium or category, from a list.

One of the most consistent requests by participants was for the app to highlight greater exposure to the artists. Participants in 11 out of the 12 evaluations wanted the app to showcase additional information about the artists, and those in 10 out of the 12 evaluations wanted the app to showcase more works by the same artist. A few participants stated that art apps should, in the very least, showcase artist biographies and their own perspectives on the works.

A large proportion of participants – those from 10 out of the 12 evaluations – requested that the app should better convey an overview of the collection – ideally in the form of a list of works – as visual previews or as a map showing the works distributed across campus. On a related note, 5 out of the 12 participants requested more information to be displayed about the works, and for each work to have a visible map depicting its location on campus, accessible from its full screen view. Participants in 5 of the evaluations also requested the app to have a 'shuffle' button: a button that shows a random pathway or selection of works when pressed. These participants noted that such a 'shuffle' function would present itself as a quick and accessible way to explore the works – particularly if the user doesn't know about

the collection.

4.3.5 Discussion

Participants provided perspectives on *A Place for Art* beyond usability issues and information seeking approaches: in particular, they discussed the how an app developed for a tablet device could be used to foster discovery and engagement with a museum's collection, and how *A Place for Art* – or any other collections app developed for the iPad – could be integrate seamlessly as part of the experience of exploring a large campus or museum collection.

Participants described their user experience of *A Place for Art* in terms of their familiarity with gestures, interactions and conventions from other iPad apps. In particular, participants stressed that the app should follow interaction design conventions from similar experiences, such as the ability to swipe left and right to view photos, pinching to zoom on detailed works, and the ability to support multiple orientations. The way the app presents the works on a pathway as a smooth, horizontal scrolling list was reported by participants to be one of the most positive and natural aspects of their interaction within the app, simply because it follows a familiar design convention from other apps. Conversely, the fact that the app employed an unusual gesture for its key functionality – the ability to create a new pathway from a given piece of work – was identified as a major usability issue.

Some participants likened the virtual experience of exploring pathways to the physical experience of wandering in a museum:

> "It feels like there's no sort of menu. There's no end. We'll just keep circling through. It is part of the uniqueness ... It really does remind me of a museum. You just walk around and look for things."

Other participants noted that the app should better emphasise the physicality of the objects and the spaces that they occupy by depicting the works on the map of a campus, or by depicting the size and scale of the objects through the use of a visual indicator. One of the participants noted that the app would work well if given to visitors as an accompaniment for a museum visit. Others have noted that – due to the expansive nature of the university campus and the location-aware capabilities of the iPad device, that the app could be used to construct a 'virtual tour' of the campus. Some participants also noted that the presentation of works in a digital format such as an iPad app was "so much easier than a book. This is way quicker to find relevant things that you like" and that the use of an app was considered to be "more environmentally friendly." Other participants praised the quality of the iPad's screen and noted that the works appeared to look better on the iPad's high

resolution screen [22] than they do in the print publication of the collection. Some participants raised the question of why an app – rather than a website – was used to showcase the works in *A Place for Art*. Others have commented that an app was necessary in order to provide its unique user experience, stating that it "should be simple to use" and that an app is "meant to do one thing and one thing really well." Participants in 2 of the evaluations have also commented that the app provides a very "Mac-like", and one of them drew comparison towards "The O" from the Museum of Old and New Art. [23] One of the participants noted that they were more likely to revisit an art app if they know it gets updated with new works on a regular basis.

Overall, the participants remarked that the app fosters a greater sense of awareness of the works that surround their everyday life on campus, and in particular were engaged by the diversity and historical provenance of the works and their strong associations with the Illawarra. In 5 out of the 12 evaluation sessions, participants often demonstrated familiarity with some of the works that they encountered on campus in their everyday lives, and responded positively to the idea that the app provides more information about the works. Participants emphasized that the app creates an awareness of the value and diversity of the University of Wollongong Art Collection.

[22] An iPad with a high density 'retina' screen was used for this study.

[23] 'The O' is a device given to visitors to the Museum of Old and New Art that uses location sensing to provide them with information about nearby objects as they walk through the gallery space. The device is actually a modified iPod Touch and retains a look and feel similar to that of an iPad or iPhone. See also: `http://www.mona.net.au/theo`.

4.4 The Brooklyn Museum Canvas

The *Brooklyn Museum Canvas* is a rich, web-based interactive visualisation [24] that allows users to explore objects from the Brooklyn Museum's online collections. [25] Using the museum's API, [26] 15,000 objects were harvested by the *CollectionWeb* framework to create concept-linked pathways that allow users to visualise their navigation history through works related by artist, theme and visual imagery.

The *Brooklyn Museum Canvas* embraces the design principles of *pathways*, *pliability*, *visual imagery*, *similarity-based suggestions* and in particular, the use of *visual momentum*. The interface emphasises smooth transitions across pathways of connected concepts, and prominently uses motion to convey navigation. Its primary interface element is a single, contiguous, linked pathway structure that employs the design principle of *parallelism* to highlight the connections that objects have with one another. Section 4.4.1 describe the design features of the *Brooklyn Museum Canvas*, and how they relate to the principles of *parallelism* and *landscapes*.

Unlike the *Virtual Museum of the Pacific* and *A Place for Art* case studies, where objects and attributes were manually extracted from their collections, the formal context for the *Brooklyn Museum Canvas* is automatically extracted from the Brooklyn Museum's API. Section 4.4.2 describes this extraction process, where metadata records are transformed into a formal context which then forms the app's concept layer for its 15,000 objects. In addition to the automatically extracted metadata, 12 *featured artists* were selected for meta-data tagging whose works were tagged with descriptive metadata relating to their themes, subject matter, visual imagery, people and places. The intention of the process, described in Section 4.4.3, was to determine how concept-linked pathways could be used to represent pathway narratives driven by richly annotated metadata that link works according to visual style, subject matter and iconography.

As a case study, the *Brooklyn Museum Canvas* demonstrates how a publicly available collection on a scale of thousands of objects can be harvested and extracted to create pathways of conceptually linked content. Having metadata automatically extracted from a data-source overcomes the costs of manually constructing and curating attributes from collection metadata – as demonstrated in the *Virtual*

[24] A prototype of the *Brooklyn Museum Canvas* is available at `http://collectionweb.cs.uow.edu.au/brooklynmuseum/vis/canvas/interact/`. The prototype currently only supports the Google Chrome Web browser. If unable to access the prototype, video demonstrations of earlier versions of the *Brooklyn Museum Canvas* are available at `https://vimeo.com/33519904` and `https://vimeo.com/65289685`.

[25] The Brooklyn Museum online collections are available at: `http://www.brooklynmuseum.org/opencollection/collections/`

[26] An API, short for 'Application Programming Interface', allows access to structured data from the museum's collection. Information and technical specifications for the Brooklyn Museum's API are available at: `http://www.brooklynmuseum.org/opencollection/api/`

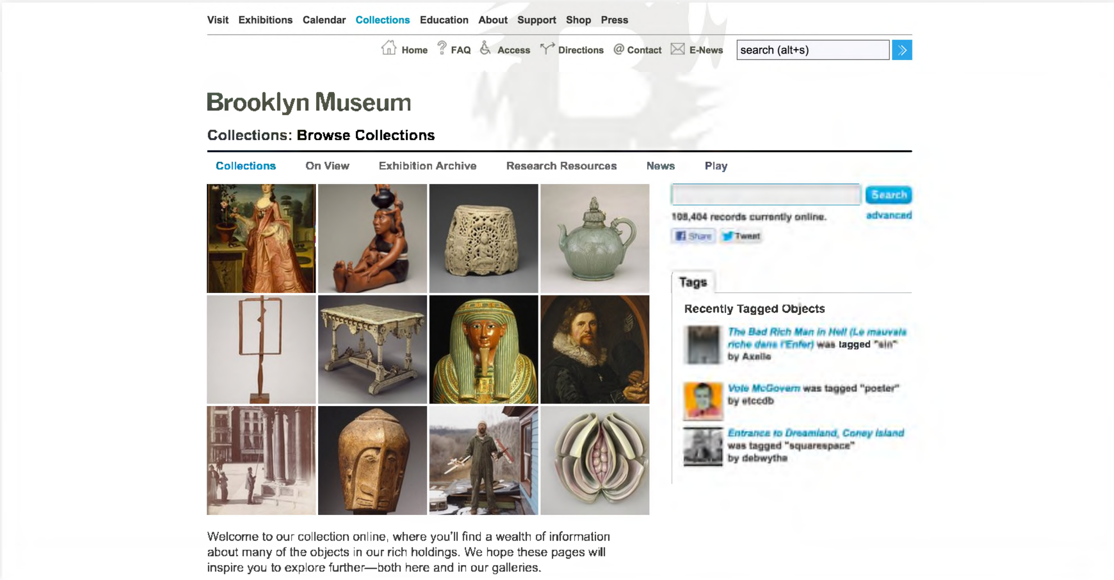

Figure 4.21: The home page of the Brooklyn Museum online collections site, available at : `http://www.brooklynmuseum.org/opencollection/collections/`.

Museum of the Pacific case study in Section 4.2.2. However a subset of objects in the Brooklyn Museum's collection were supplemented with additional metadata to demonstrate how objects could be linked thematically using concept-linked pathways. Further, as the *Brooklyn Museum Canvas* embraces the design principles of *pathways* and *visual momentum*, this case study presents an alternate design to a more conventional and publicly available online collections website. [27]

4.4.1 Interface design and user experience: parallelism and landscapes

The Brooklyn Museum online collections site, shown in Fig. 4.21: contains a welcome page highlighting featured objects, a search box, the social media 'Share' and 'Like' buttons and navigation links to each of its 12 collections with an option to browse by decade. Navigating to a collection page presents the user with thumbnails of objects from that particular collection, each linking the user to a single object page. The single object page, shown in Fig. 4.22, features an image and a description, where a user can view related works by clicking on other objects that share the same tags, or explore further by entering terms into its search box.

Figure. 4.23 presents the same object, but as it appears in the *Brooklyn Museum Canvas*. Again, the object's image and label are present on a gallery of horizontally arranged images. This gallery, called a 'pathway', can be explored by a 'camera'

[27] `http://www.brooklynmuseum.org/opencollection/collections/`

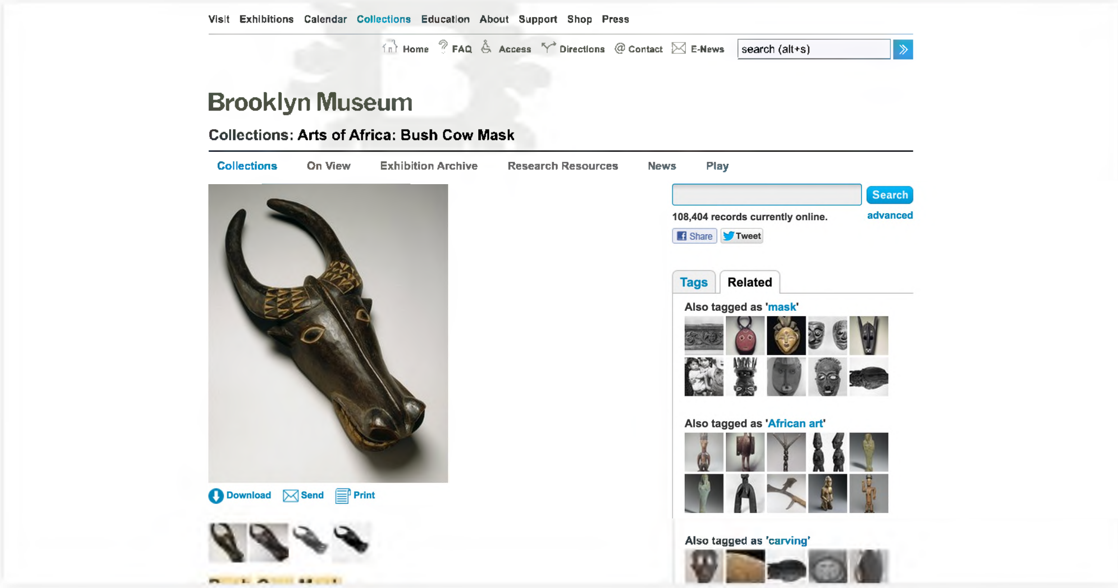

Figure 4.22: An object from the Brooklyn Museum online collections site, a bush cow mask from Bamenda, Cameroon.

that can be panned horizontally by the user. Each pathway visualises a formal concept: objects are represented by a series of images and labels, and attributes are represented by a natural language description affixed underneath. A user can pan the camera left and right by clicking on the arrows or by swiping their fingers on the trackpad. Unlike the Brooklyn Museum online collections site, where viewing related works requires the user to click on a link and experience a page refresh, the *Brooklyn Museum Canvas* allows users to view similar objects within the same view. Following the principles of *visual momentum* and *pliability* – the tight coupling action and response – a user can gently pan the camera through direct manipulation with a trackpad, or with a flick of their fingers, glide all the way to the end.

Each object on the pathway shows a high resolution image of the work accompanied by a label, which consists of a title, the artist, its medium, date and a description. At any point along the pathway, a user can 'inspect' an object by clicking on it, causing the camera to focus on the object and zoom in, as shown in Fig. 4.24. The interface smoothly zooms in without any delay or page refresh – the intention being that *visual momentum* within the app is retained. Further the action preserves the user's *orientation* on the pathway structure: the user is assured that when they are finished 'looking' at the object, the camera will again zoom out, returning it to its original position.

Unlike the *Virtual Museum of the Pacific* and *A Place for Art* case studies, some objects in the *Brooklyn Museum Canvas* are presented to a relative scale, where objects with higher width or height dimensions may appear larger on the pathway, such as the ones shown in Fig. 4.25. This is intended to represent partial

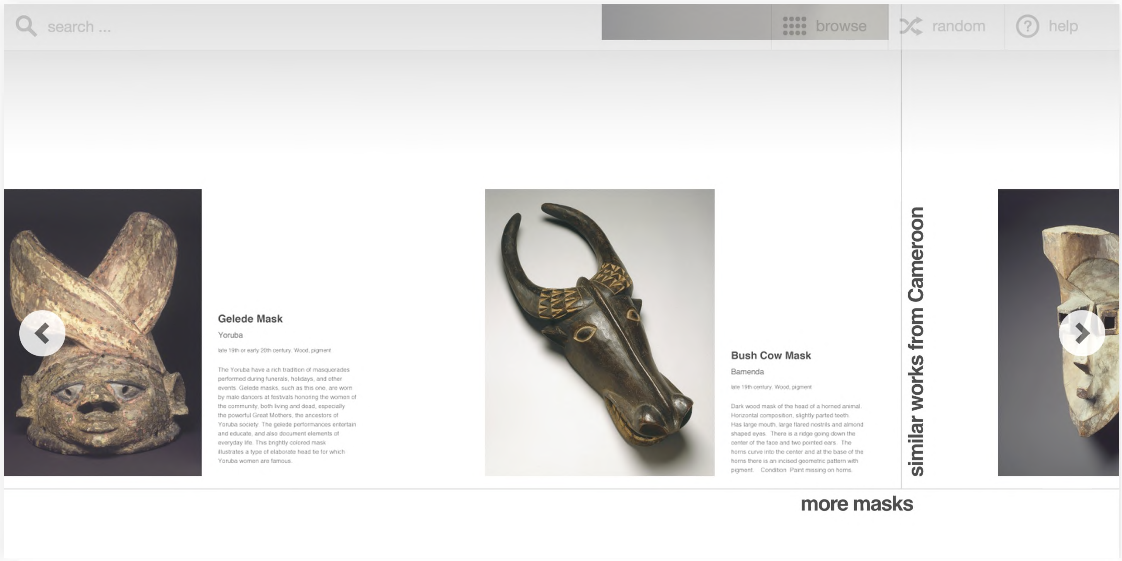

Figure 4.23: The bush cow mask from Fig. 4.22 as displayed on a pathway within the *Brooklyn Museum Canvas*.

visual knowledge about the size of the works in the collection. Although most records within the set of 15,000 objects have metadata relating their width, height or depth: there were some practical limitations addressed in presenting objects this way. For instance, only objects of specific types, such as those flagged with 'painting', 'drawing' or 'photograph' – any medium that is primarily two-dimensional – could be accurately scaled. In addition, the interface needed to employ upper and lower limits to how large or small an object could be scaled as the variance in size among objects was considered too large to legibly represent the smallest and largest objects in the collection to a true relative scale. While some two-dimensional works from the Brooklyn Museum's collection were successfully scaled this way, the results of the process are more evident when applied to anthropological collection of objects with consistent dimensional data, such as those from the *Virtual Museum of the Pacific* case study shown in Fig. 4.26.

Every pathway within the *Brooklyn Museum Canvas* has, adjoined to it, one or more similar pathways. Each adjoining pathway, or child node as described in Section 3.4, is presented at a 90 degree angle to the pathway that the user is currently viewing, as shown by the pathway labelled "similar works from Cameroon" in Fig. 4.23. A user can navigate through the pathway structure by clicking on the text labels, which at that point, the camera rotates 90 degrees to orient itself to the new pathway. Figure 4.27 shows how the camera rotates itself to align to the new pathway: during the course of the rotation, the text label of the user's previous pathway moves forward and aligns itself to the beginning of the new pathway. The user can also click on that text label and return to the previous pathway. In the *Brooklyn Museum Canvas*, the labels always move to last point of navigation, allowing the

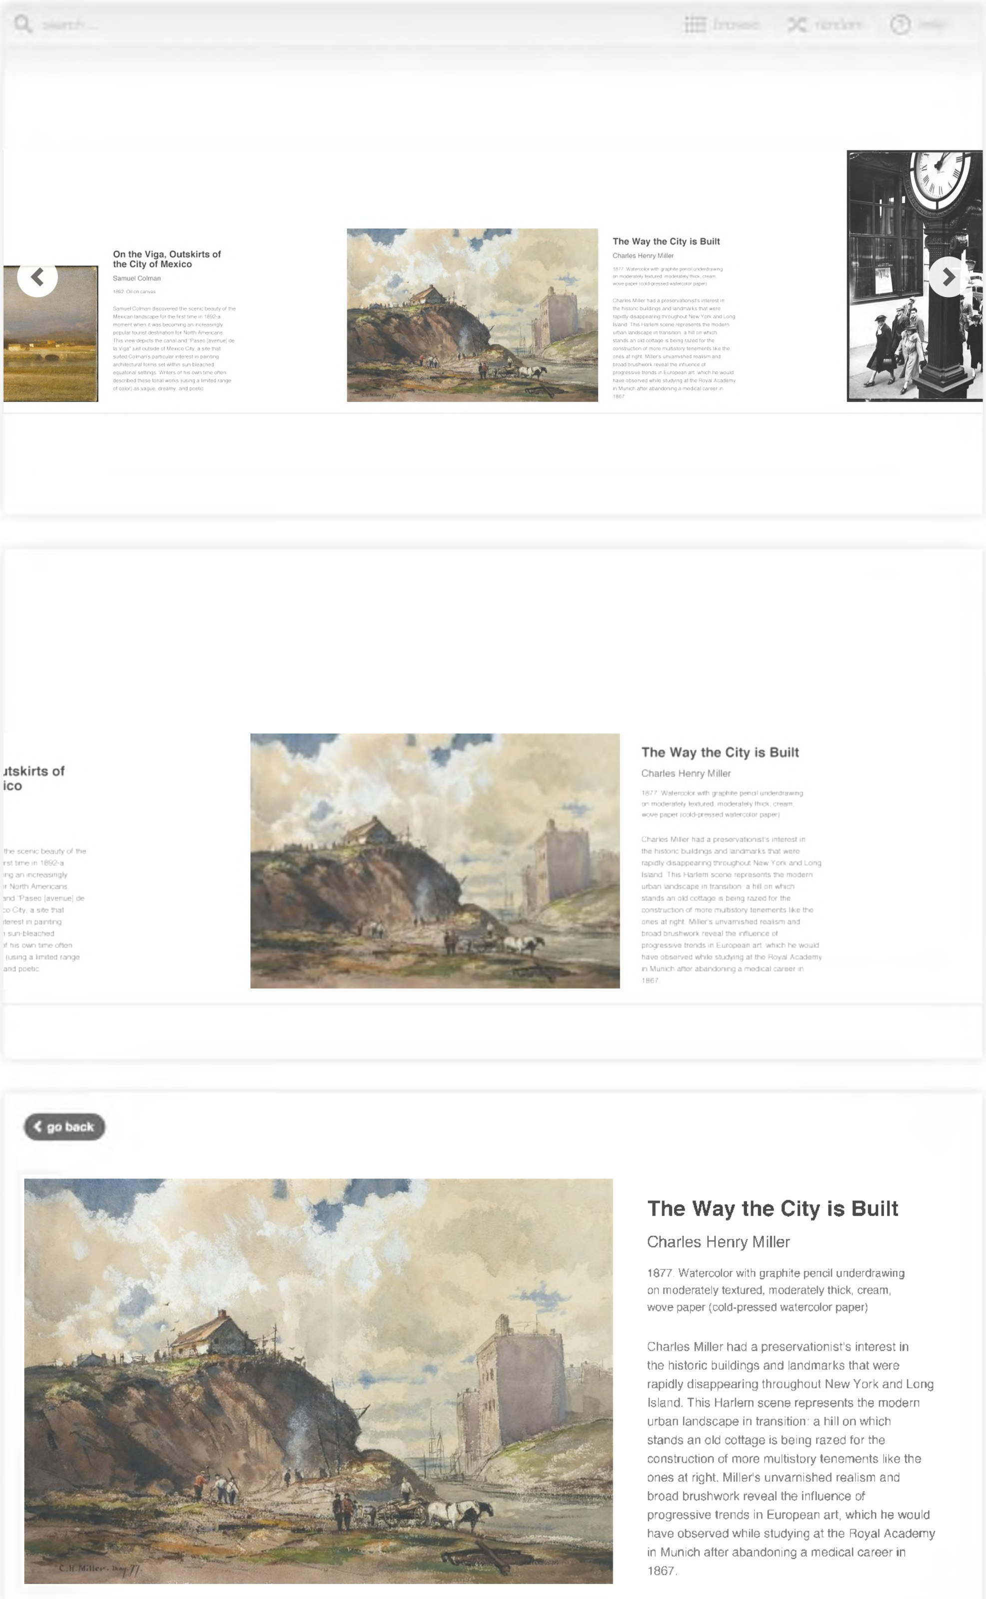

Figure 4.24: A series of frames that depict a user zooming in to an object on a pathway.

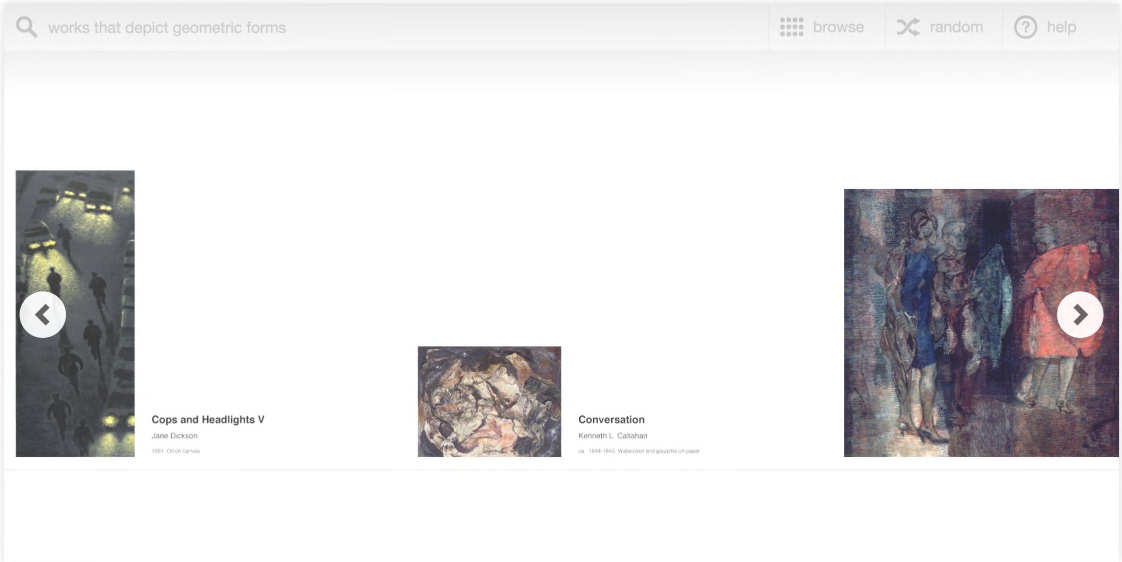

Figure 4.25: Two-dimensional works in the *Brooklyn Museum Canvas* are presented to a relative scale: imparting partial visual knowledge about the size of the works.

Figure 4.26: The user interface of the *Brooklyn Museum Canvas* showing images, metadata and pathways from the *Virtual Museum of the Pacific* case study, demonstrating how it could be used to show a collection of anthropological objects to a relative scale.

user to easily view the relationship between their current and previous pathway. In the same way concept-linked pathways are both generative and deterministic based on a user's journey, the *Brooklyn Museum Canvas* creates an ever-growing structure of concepts, allowing the user to explore similar concepts ad-infinitum, or backtrack throughout their navigation history and create new pathways.

When a user first visits the *Brooklyn Museum Canvas*, it displays a pathway based on a randomly selected formal concept from its context. Users can also select other pathways to branch from with the option of browsing by category, concept, medium, theme, or by the people, places and cultures depicted in the works (Fig. 4.28). Users can also search for pathways by entering keywords, and the interface matches and displays suggestions of semiconcepts as the user types. These suggestions are expressed in natural language, so that for instance, the search terms 'new york poverty' suggest the phrase "works that depict people in poverty in New York". By allowing the user to search for semiconcepts based on free text, the *Brooklyn Museum Canvas* can present formal concepts in a way that implies additional attributes based on the search terms entered. For instance, although the phrase "works that depict people in poverty in New York" is generated from attribute values { 'poverty', 'New York' }, the *Brooklyn Museum Canvas* presents the formal concept "photographs by Consuelo Kanaga that depict people in poverty and melancholy" – inferring the additional attributes 'photograph' and 'Consuelo Kanaga' – given that all works that depict people in poverty in New York are photographs by that artist. This attribute inferencing is based on the way formal concepts are formed from single or partial attribute sets, as described in Section 3.1.

Unlike the static mediums of books or print catalogues, or even the relatively static presentation of the Brooklyn Museum online collections site, *Brooklyn Museum Canvas* primarily uses the affordances of motion to convey the user's interaction with the information space: *zooming* to inspect an object or return to the pathway; horizontal *scrolling* to explore and compare conceptually similar objects of a formal concept; and *rotation* to convey a shift from one group of related objects to another. These affordances of *scrolling* and *rotation* – each respectively conveying object browsing and concept navigation – are also evident in the design of *A Place for Art* (Section 4.3).

In the *Brooklyn Museum Canvas*, the use of motion and visual momentum, along with its use of pathways and intersections to represent groupings and divergences, characterise the experience as a connected *landscape*, a concept introduced by Janet Murray [16] in Section 2.4.2. Murray argues that these characteristics are similar to the virtual landscapes presented in 3D role-playing games [16, p. 175], citing the *The Legend of Zelda, Ocarina of Time* as an example of a playing experience that embodies similar ideas of expansiveness, path-finding and open-ended exploration.

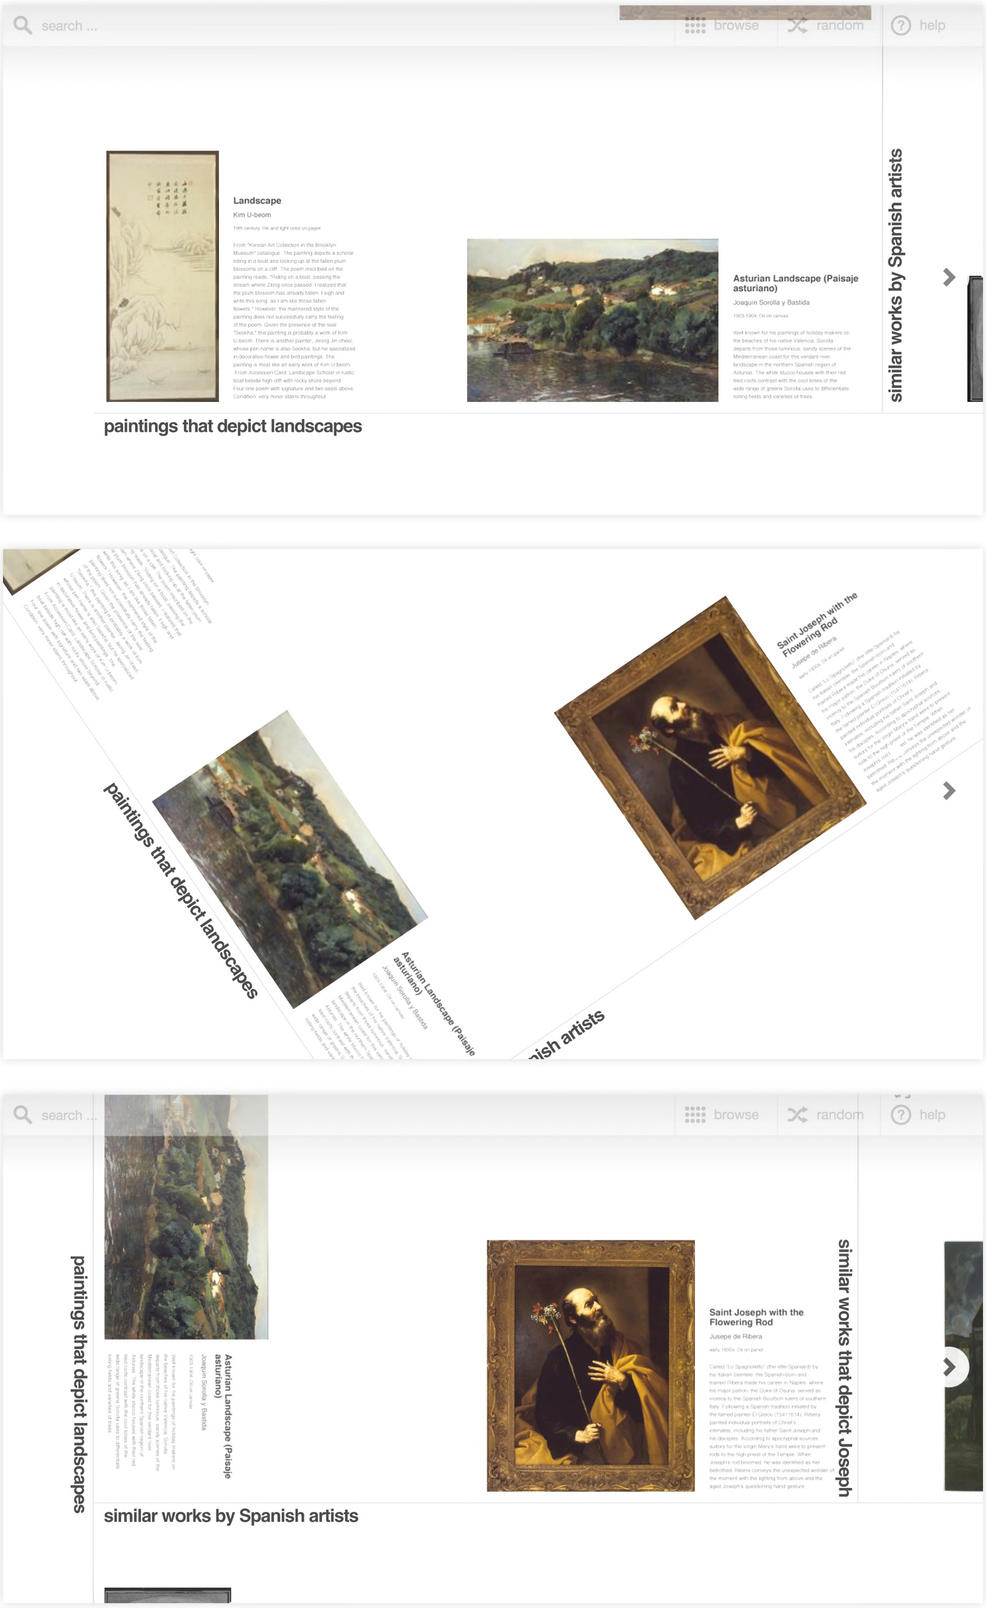

Figure 4.27: A series of frames that depict a user navigating from one pathway to another. The camera 'rotates' to orient the user on the new pathway. Note that the label 'paintings that depict landscapes' follows the camera to the new point of rotation, allowing the user to backtrack to their previous pathway.

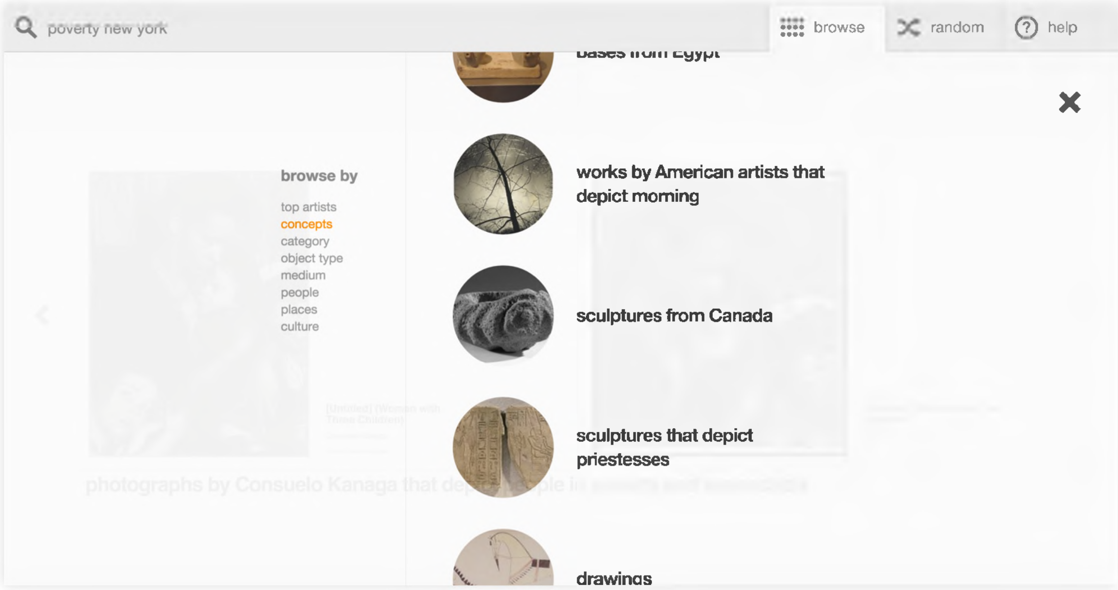

Figure 4.28: Browsing a selection of concepts within the *Brooklyn Museum Canvas*.

By contrast, the notion of *containers*, another concept introduced by the author, better describes the Brooklyn Museum's online collection site: each of its 12 collections is a container, itself composed of containers of object pages, where a navigation from one page to another is interrupted by a page refresh. The contrast between the *landscape* of the *Brooklyn Museum Canvas* and the *containers* of the *Brooklyn Museum Online Collection* can be seen when comparing Figs. 4.22 and 4.23, although, as shown in Fig. 4.22, the *containers* of links on the right-hand side, each providing a concise visual summary of similar objects as grouped by element, provide a richer and more dense selection of navigation options than the prominent and sparse pathway structures sprawled across the *Brooklyn Museum Canvas*.

The use of pathways and labels in this manner is also based on the design principle of *parallelism*, a term introduced by Edward Tufte that refers to the way one or more visual elements in are graphic or statistical display are connected – often using comparison or contrast – to illustrate differences in space or time, or to convey context and meaning [129, p. 82]:

> Parallelism connects visual elements. Connections are built among images by position, orientation, overlap, synchronisation, and similarities in content. Parallelism grows from a common viewpoint that relates like to like. Congruity of structure across multiple images gives the eye a context for assessing data variation. Parallelism is not simply a matter of design arrangements, for the perceiving mind itself actively works to detect and indeed to generate links, clusters and matches among assorted visual elements.

The visual layout of pathways presents such parallels: the horizontal positioning of similar works of art, labelled with natural language descriptions that connect one another as pathways intersect. Parallelism also has a temporal dimension that can be used to illustrate a narrative: the passage of a user's journey across concepts of artists, genres, mediums and styles. In Figure 4.27 for example, in observing the movement of the label 'paintings that depict landscapes' as it moves towards the new pathway 'similar works by Spanish artists', the placement of the label creates a new phrase that describes, in natural language, the conceptual connection between the two pathways. The connections between the two pathways are illustrated by a single stream of text partially broken by a 90 degree angle: such compositions are formed every time a user moves from one adjoining pathway to another within the structure.

Parallelism can also be used to describe narrative over a period of time. Edward Tufte illustrates by example [129, pp. 90–91] through his reproduction of *Rock 'N' Roll is Here to Pay: The History and Politics of the Music Industry* [130], a timeline of prominent artists, musical genres, and their expansion, divergences and evolution from the 1950s to the 1970s. In his description of this "streams-of-story", shown in Fig. 4.29, Tufte describes how major genres spawn and evolve from their musical parents. The tree-like structure of the visualisation and the similarity of musical divergences are not unlike the conceptual divergences seen as a user explores the concept-linked pathways within the *Brooklyn Museum Canvas*, although the interface does not comprehensively provide a per-year overview of artists as this visualisation does. Although the *Brooklyn Museum Canvas* inherits some elements parallelism through the use of visual composition and motion, Section 4.4.3 describes, in further detail, how the themes, subject matter and visual imagery of key artists within the collection are illustrated by user journeys through a series of concept-linked pathways.

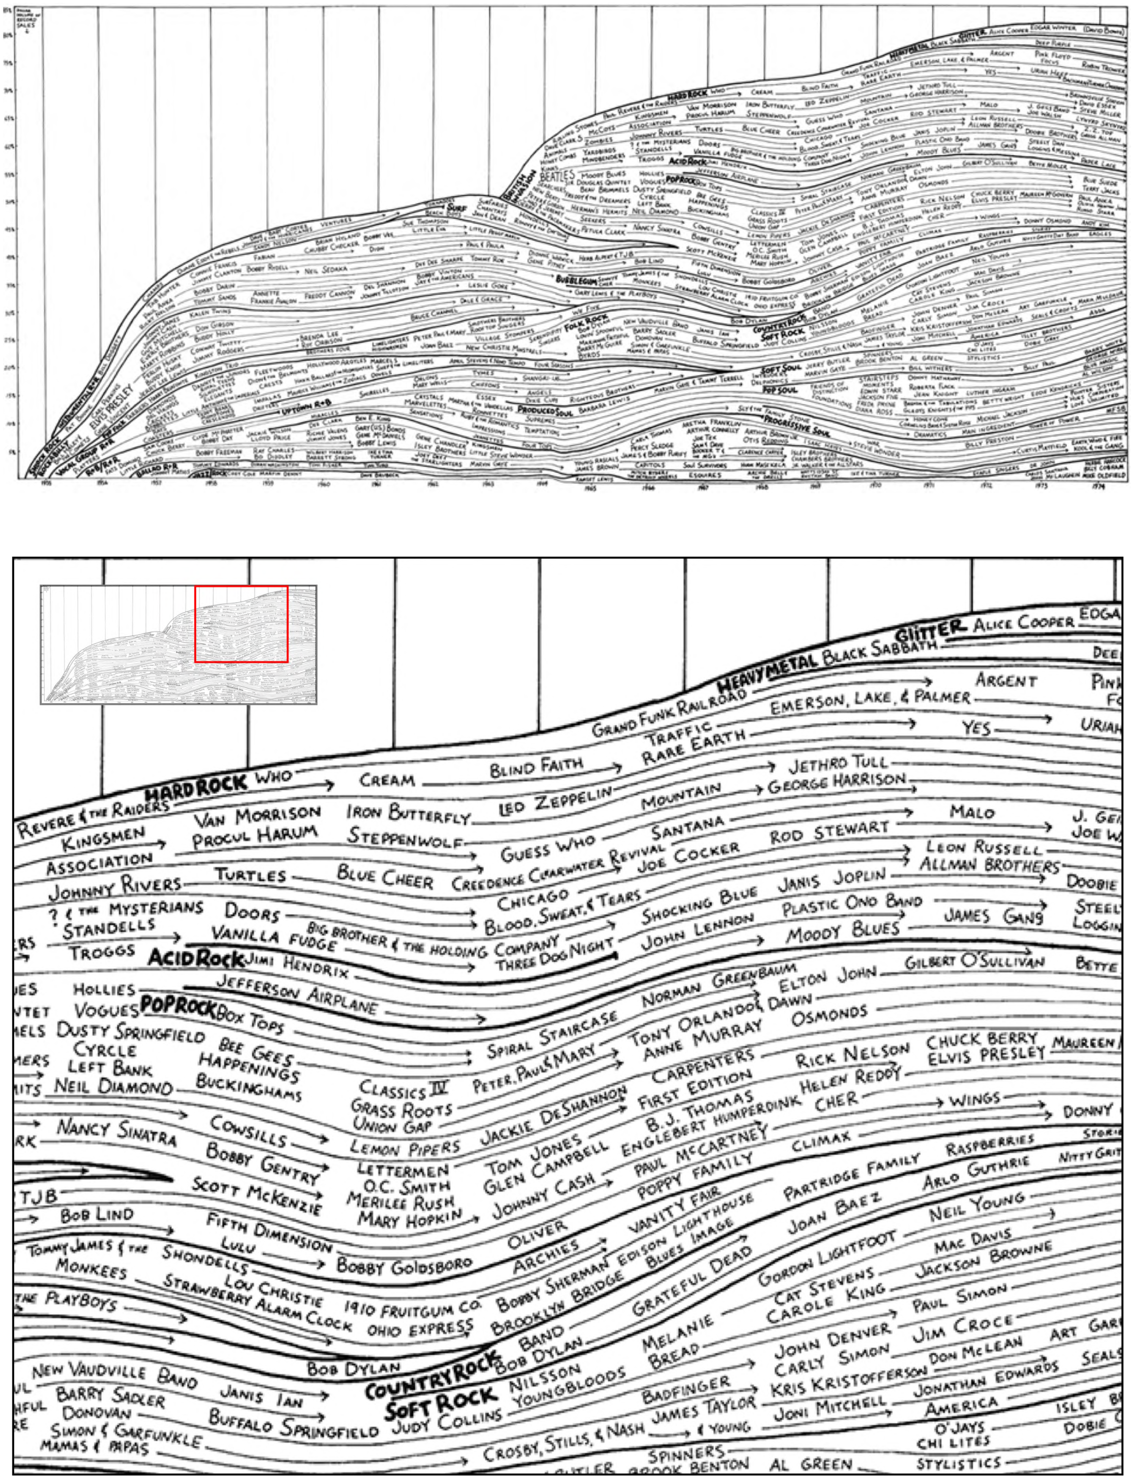

Figure 4.29: A visualisation of musical genres from the 1950s to the 1970s [130], used by Edward Tufte [129, pp. 90–91] as an example of how parallelism illustrates both their context and evolution on the horizontal axis, and a list of prominent artists for a given year on the vertical axis. The above figure shows the complete visualisation, with a detailed inset shown below.

4.4.2 Harvesting collection metadata

The objects for the *Brooklyn Museum Canvas* were harvested using the Brooklyn Museum's API. [28] The API implements the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) [131] [29] [30], a protocol that allows client applications and aggregators, known as *harvesters*, to collect and update metadata descriptions of digital resources from a *provider*. The object records from the API are formatted using the Dublin Core metadata standard, [31] a content description model that uses a fixed set of metadata elements, such as `title`, `creator` or `description` to describe a digital resources. The following is a snippet of the object record in Fig. 4.22 from the Brooklyn Museum's OAI-PMH API `GetRecord` method:


```
<record>
   <header>
      <identifier>o-2972</identifier>
      <datestamp>2011-09-13T17:34:08Z</datestamp>
      <setSpec>objects:african_art</setSpec>
      <setSpec>objects:full_images</setSpec>
   </header>
   <metadata>
      <object id="2971" title="Mask" accession_number="... >
         <images>
            ...
         </images>
         <artists>
            ...
         </artists>
         <geolocations>
            ...
         </geolocations>
      </object>
   </metadata>
</record>
```


[28] `http://www.brooklynmuseum.org/opencollection/api/`
[29] The OAI-PMH specification is available at: `http://www.openarchives.org/OAI/2.0/openarchivesprotocol.htm`
[30] Documentation for the Brooklyn Museum's implementation of OAI-PMH is available at: `http://www.brooklynmuseum.org/opencollection/api/docs/oai`
[31] The Dublin Core specification is available at: `http://dublincore.org/`

The record consists of a `<header>` element and a `<metadata>` element. The `<header>` element contains a repository identifier and a datestamp that indicates when the record was created, modified or deleted. OAI-PMH uses these datestamps to implement *selective harvesting* that allows a harvester to retrieve only the records that were created, deleted or updated after a specified time period. For instance, in order to ensure that the pathways in the *Brooklyn Museum Canvas* were kept up to date, *CollectionWeb* implements selective harvesting to incrementally update its record store on a timely basis, rather than re-download all of the records from the repository. Once the records are incrementally fetched, there is no need to recompute the formal concepts, as the collection's *concept layer* can also be incrementally updated with the new objects (Section 4.1.3).

The `header` section also contains two `setSpec` tags, each indicating a *set* that the object is a member of. In the above example, the tags indicate that the object is part of a collection of works that feature full-sized images, and that it's part of the 'Arts of Africa' collection. In OAI-PMH, selective harvesting can also be implemented using set membership criteria, so that only objects of a specific collection are harvested. Out of the 84,963 records available via the Brooklyn Museum's API, only 15,000 were extracted: this was due to the fact a) not all object records have images and b) the *Brooklyn Museum Canvas* was required to show the diversity of the museum's 12 collections, so the amount of objects selected for harvesting was reduced to 15,000 in order to flatten the relatively uneven distribution of objects from each collection, as shown in Table 4.5. The table also shows the number of objects harvested from each collection: the objects with the most descriptive metadata fields, such as those used for exhibition labels, were selected for harvesting.

Once the objects are harvested, the *CollectionWeb* framework constructs the formal context by parsing each record and extracting attributes from their `<metadata>` elements. Each `<metadata>` element contains an `<object>` element that lists, as its attributes, metadata fields such as the its title, medium and description. The `object` element also contains an `<artists>` element that list its creators or originating cultures, and a `<geolocations>` element that lists one or more locations associated with the work. Each of these elements are parsed using the appropriate parsers in the *CollectionWeb* framework. Using the Bush Cow Mask from Fig. 4.22 as an example, Table 4.6 lists the processes used to extract attributes from its metadata elements. The resulting 12 attributes, when combined, produce the following natural language statement (Section 4.1.2):

```
Bamenda bush cow mask from Bamenda, Cameroon made with pigment
and wood
```

Table 4.5: The number of objects harvested from each of the Brooklyn Museum's 12 collections.

Collection	Total Objects	Objects Harvested
Decorative Arts	16627	1638
Arts of the Americas	14132	1637
Contemporary Art	10062	915
Asian Art	8703	1638
Egyptian, Classical, Ancient Near Eastern Art	8645	1638
American Art	7038	1638
Photography	6366	1405
European Art	6220	1637
Arts of Africa	3962	1468
Arts of the Pacific Islands	1596	235
Arts of the Islamic World	1477	1040
Elizabeth A. Sackler Center for Feminist Art	135	111
Total	84963	15000

Table 4.6: The attribute extraction process for the Bush Cow Mask in Fig. 4.22.

Object Metadata	**Extraction Process**
Attributes of the `<object>` element:	
`classification="Masks"`	Converts to singular form, then maps to the `isTypeOf` predicate, resulting in attribute: `isTypeOf:mask`

`title="Bush Cow Mask"`	Maps to the `titled` predicate, resulting in attribute:
	`titled:Bush Cow Mask`
	Parsed using the *ObjectDescriptionParser*, resulting in attributes:
	`is:bush cow mask`
	`is:mask`

`medium="Wood, pigment"`	Noun phrases extracted using the *NounPhraseParser*, resulting in attributes:
	`hasMedium:wood`
	`hasMedium:pigment`

`collection="Arts of African"`	Maps to `hasCategory` predicate based on the name of the collection, resulting in attribute: `hasCategory:African`
Attributes of each `<artist>` element:	
`name="Bamenda"` `role="Culture"`	Given that `role="Culture"`, `Bamenda` maps to the `fromCulture` predicate, resulting in attribute: `fromCulture:Bamenda`
Attributes of each `<geolocation>` element:	

`name="Bamenda, Cameroon"`
`location-type="place made"`
	Given that `location-type="place made"`, `Bamenda, Cameroon` maps to the `madeIn` predicate, resulting in attributes:
	`madeIn:Bamenda, Cameroon`
	`madeIn:Cameroon`
	`associatedWithLocation: Bamenda, Cameroon`
	`associatedWithLocation: Cameroon`

Figure 4.30: *Meguro Drum Bridge and Sunset Hill, No. 111 from One Hundred Famous Views of Edo* by Japanese artist Utagawa Hiroshige. Image obtained from `http://www.brooklynmuseum.org/opencollection/objects/121723/Meguro_Drum_Bridge_and_Sunset_Hill_No._111_from_One_Hundred_Famous_Views_of_Edo`

Within the collection, objects that were classified as photographs, prints or sculptures had titles that described the themes and subject matter of a work, whereas objects that were classified as ceramics, vessels or masks had titles that provided a literal description of the work. For instance, the `Bush Cow Mask` in Table 4.6 is classified as a `Mask`, and its `title` – *Bush Cow Mask* – was parsed using the `ObjectDescriptionParser`, resulting in attributes `is:bush cow mask` and `is:mask`. By contrast, the work shown in Fig. 4.30 – a Japanese woodblock print titled *Meguro Drum Bridge and Sunset Hill, No. 111 from One Hundred Famous Views of Edo* – uses the `ImageSubjectParser` to extract attributes `depicts:views`, `depicts:location:Meguro Drum Bridge`, and `depicts:location:Sunset Hill`. Tables 4.7 and 4.8 list the metadata and extracted attributes for this object, with Table 4.7 listing the metadata of its `<object>` element, and Table 4.8 listing the metadata of its `<artist>` and `<geolocation>` elements.

Each `<artist>` and `<geolocation>` element generates attributes that describe

Table 4.7: The attribute extraction process of the `<object>` metadata fields for the Japanese woodblock print in Fig. 4.30.

Object Metadata	**Extraction Process**
Attributes of the `<object>` element:	
`classification="Print"`	Maps to the `isTypeOf` predicate, resulting in attribute: `isTypeOf:print`

`title="Meguro Drum Bridge and Sunset Hill, No.111 from One Hundred Famous Views of Edo"`	Maps to the `titled` predicate, resulting in attribute:
	`titled:Meguro Drum Bridge and Sunset Hill, No.111 from One Hundred Famous Views of Edo`
	Detection of keywords suggesting that the work is part of a series, resulting in attribute:
	`fromSeries:One Hundred Famous Views of Edo`
	Parsed using the *ImageSubjectParser*, resulting in attributes:
	`depicts:views`
	`depicts:location:Meguro Drum Bridge`
	`depicts:location:Sunset Hill`

`medium="Woodblock print"`	Noun phrases extracted using the *NounPhraseParser*, resulting in attributes:
	`hasMedium:woodblock print`
	`hasMedium:print`

`collection="Asian Art"`	Maps to `hasCategory` predicate based on the name of the collection, resulting in attribute: `hasCategory:Asian`

Table 4.8: The attribute extraction process of the `<artist>` and `<geolocation>` metadata fields for the Japanese woodblock print in Fig. 4.30.

Object Metadata	**Extraction Process**
Attributes of each `<artist>` element:	

`name="Utagawa Hiroshige"`
`nationality="Japanese"`
`role="Artist"`
	Given that `role="Artist"`, `Utagawa Hiroshige` maps to the `byArtist` predicate, resulting in attributes:
	`byArtist:Utagawa Hiroshige`
	`associatedWithAgent:Utagawa Hiroshige`

Attributes of each `<geolocation>` element:	

`name="Japan"`
`location-type="place made"`
	Given that `location-type="place made"`, `Japan` maps to the `madeIn` predicate, resulting in attributes:
	`madeIn:Japan`
	`associatedWithLocation:Japan`

specific relationships between the object and its artist or location, (e.g. `byDesigner :Tiffany Studios` or `designedIn:New York`), along with attributes that denote a general association (e.g. `associatedWithAgent:Tiffany Studios` or `associatedWith-Location:New York`), given that a work may be produced by one or more people, or associated with more than one location. Further, as demonstrated previously in Table 4.6, the `<artist>` element is also used to describe the originating culture of an anthropological work. For the woodblock print in Table 4.8, the framework recognised that the work was part of a series based on the detection of keywords such as `from`, `collection of` or `one of a`. Further, the presence of the `location:` element indicates that the Stanford Named-Entity Recogniser [109] – as implemented by the `ImageSubjectParser` – recognised the place names of `Meguro Drum Bridge` and `Sunset Hill` from the work's title. The presence of this `location:` *sub-predicate qualifier* allows for a more expressive natural language statement of the work and its associated concepts:

```
a print by Utagawa Hiroshige made with woodblock print
that depicts views of Meguro Drum Bridge and Sunset Hill
```

Rather than list the noun phrases of a work's title, the natural language statement provides a richer expression than of the work as it describes how elements of its subject matter are related. The following section reports on how a set of manually

added tags – metadata that describes both the subject matter of the works and the interpretation of their visual styles – are used to link pathways as *narratives*.

4.4.3 Exploring featured artists through narrative

There were a number of artists whose works were prominently represented within the Brooklyn Museum's online collection. Some of these artists were featured as part of physical exhibitions held by the museum, and as a result, the collection contains an extensive knowledge store of their works. Many of these works have, associated with their object records, full-text descriptions of their exhibition labels. Given that these artists were selected for exhibition, this section demonstrates, via a curated process of manual tagging, how concept-linked pathways can replicate an exhibition-like experience within the *Brooklyn Museum Canvas*.

12 of the most prominent artists from the collection – as determined by the number of works – were selected for manual tagging, although only two of these artists will be presented in this section by example. The tagging process aimed to capture the knowledge represented by the exhibition labels and the artist biographical data provided by the Museum's API. Based on a visual inspection of each work, the process also aimed to capture a cross-section of themes, subject matter and visual styles. Formal Concept Analysis was then used to construct concept-linked pathways that combined the explicit knowledge encoded within the museum's metadata, the tacit background knowledge of the object labels and descriptions, and the interpretive, visual knowledge of the works to recreate an interactive, curated experience of the collection. The representation of this knowledge within the *Brooklyn Museum Canvas* intends to parallel the "streams-of-story" depiction of artists, genres and their evolution in Fig. 4.29, although in this case (and rather analogously), the user can choose their own stream and dictate their own story.

The process of manual tagging required a tool that allowed a tagger to easily assign key terms to works of art. The tool needed to assist the tagger to re-use commonly entered tags, and it also needed to provide a visual summary of the collection so that the tagger can 'get a sense' of the visual style of the works. A tagging tool, shown in Fig. 4.31, was built specifically for this purpose. The tool provides live previews of each work along with a concise visual overview of all works from that collection. When a user hovers their mouse cursor over a single tag, its interface highlights all works displayed with that tag, such as the series of 'eye' themed photographs shown in Fig. 4.31. This allows the tagger to use visual cues and ensure that tags are applied consistently.

Given that these tags described the subject matter and visual style of the works, a number of sub-categories were identified, such as tags that describe vi-

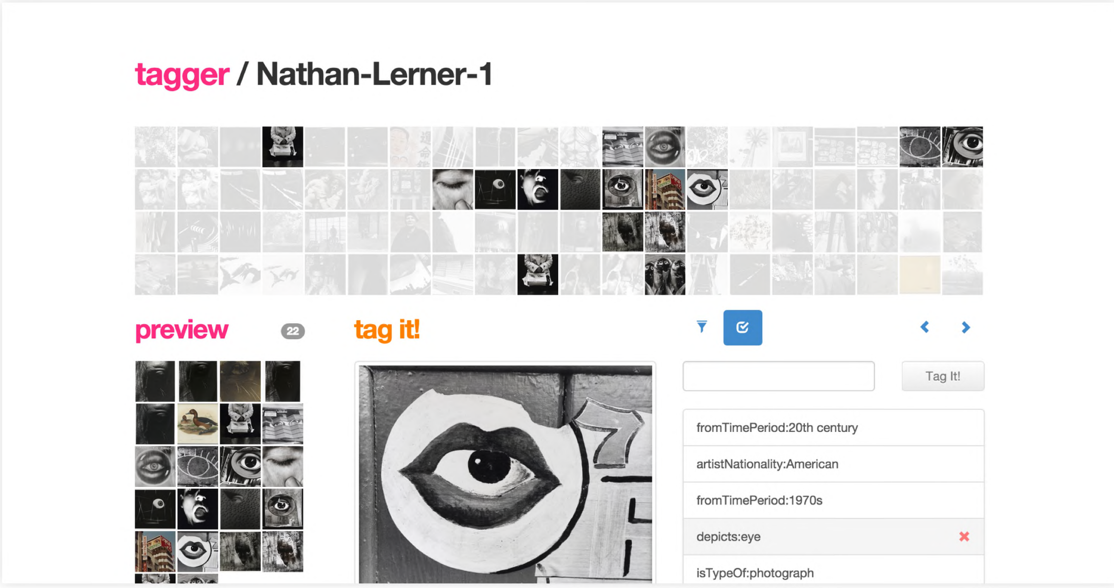

Figure 4.31: The tagging tool provided a visual overview of the collection, allowing the ability to tag works based on common motifs, visual styles and techniques.

sual elements, emotional themes depictions of people, places and locations. For the attributes generated by these tags, these sub-categories were annotated as *sub-predicate qualifiers*, so for instance, the tag `depicts:person-subject:field worker` refers specifically to the fact that the field worker in the image is a person, and the tag `depicts:visual:the glare of the Sun` refers specifically to the fact that it's a prominent visual feature or stylistic trope. These qualifiers allow for a richer expression of the narrative of the user's journey, as they explore works as linked by theme and visual style.

The narrative that's revealed to the user is highly dependent on the themes and styles that were evident in the curated selection of works for that artist. These narratives are demonstrated by example via the analysis of the metadata of two prominent artists from the collection, each demonstrating a stream of connected pathways as they would appear in the *Brooklyn Museum Canvas*.

Consuelo Kanaga

Consuelo Kanaga was an American photographer and photojournalist who produced works from the 1930s and 1940s. Intending to capture the essence of humanity in every photo she took, she was well known for her cityscapes, still lifes, portraits and photographs of African Americans. [132] There are approximately 1500 of her photographic prints within the collection, making her by far the most prominent artist in terms of the amount of works collected.

Her portraiture of African Americans was identified as one of the most subjects

associated with her works, as shown in Table 4.9. The table also highlights other recurring themes, such as animals, natural and rural scenes, and portraits of women and children. An analysis of the most common visual styles, shown in Table 4.10 revealed that her works were often shot in strikingly high contrast, with many of her works featuring shadows, reflections, silhouettes and the Sun's glare. Her works also explored depictions of family and poverty, such as *[Untitled] (Mother with Children, New York)* in Fig. 4.1 (p. 68), arguably one of her most iconic works.

In Section 4.1.2, it was demonstrated how concept-linked pathways could be used to explore related themes, such as families in poverty, melancholic portraits of women, and streetscapes of New York. This section will expand on these themes by demonstrating how pathways can be used to illuminate a *narrative* of works, beginning with the formal concept titled *photographs by Consuelo Kanaga that depict people in melancholy.* This narrative, as illustrated in Fig. 4.32, consists of a series of diverging concepts, each being a child concept of the one above it (see Section 3.4.1). Alongside each concept are images of its objects that collectively provide a visual explanation of the similarities – and gradual shift – of the themes in her works. The figure provides a visual overview of the imagery and narrative that a user would experience as they navigate through the successive pathways of the *Brooklyn Museum Canvas.*

Evident in this narrative are the common threads that prominently characterise Kanaga's portraiture: the first four concepts in Fig. 4.32 highlight the implicit connections between anguish, poverty and the urban setting of New York. The mathematics of *concept similarity* (Section 3.3) dictate that these connections are made through the shared attributes of these formal concepts. The second half of the narrative links the works based on style, rather than theme: the visual similarities of the objects presented in each concept – although not as evident in their natural language descriptions – are a product of a tagging and curating process that relied heavily on a visual inspection and comparison of the works. This resulted in a connection between the silhouetted imagery of the concept titled *more photographs by Consuelo Kanaga* and the shadowy imagery prevalent in the following concept *more works by American artists that also depict shadows.* A similar likeness in visual similarity is shown between her own self-portraiture, and that of Andy Warhol.

The divergences of themes and styles across different artists is also evident in Fig. 4.33, a "streams-of-story" representation of the narrative in Fig. 4.32. The visualisation follows the principles derived from Edward Tufte's [129, pp. 90–91] exposition of *Rock 'N' Roll is Here to Pay: The History and Politics of the Music Industry* [130] in Fig. 4.29 – but instead of showing an evolution of musical genres and artists, the visualisation reflects the emerging themes, user journeys and narratives that spawn from a single concept. The left-to-right axis depicts a sequence

Table 4.9: The most frequently occurring tag attributes of works by Consuelo Kanaga.

Attribute	No. Objects
`depicts:presentation:portrait`	147
`depicts:nature`	86
`depicts:person-subject:woman`	73
`depicts:person-subject:child`	72
`depicts:landscape-subject:rural setting`	69
`depicts:person-subject:African American`	62
`depicts:sepia hues`	50
`depicts:animal:animal`	44
`depicts:person-subject:girl`	42
`depicts:water`	41
`depicts:person-subject:young girl`	40
`depicts:person-subject:man`	39
`depicts:visual:high contrast imagery`	39
`depicts:visual:reflections`	38
`depicts:tree`	32

of pathway divergences within the *Brooklyn Museum Canvas*, highlighting three different trajectories from a root concept. The highlighted pathway from Fig 4.32 is darkened and shown in bold text. Its concepts are presented in parallel to the artists that span across each concept, each shown as coloured branch in the top half of the visualisation. For example, the connection between the fifth concept – *more photographs by Consuelo Kanaga* – and its following concept – *more works by American artists that also depict shadows* is shown in parallel to the emerging branches of artists Nathan Lerner and Kenyan Cox as they both share visual and stylistic cues to Kanaga's photographs. The other pathways present in the visualisation also give rise to the emerging themes in Kanaga's work, such as the connection between her portraiture of elderly women and rural scenes in the second pathway; and the connection between the urban scenes of poverty in New York, San Francisco, and related works by Nathan Lerner.

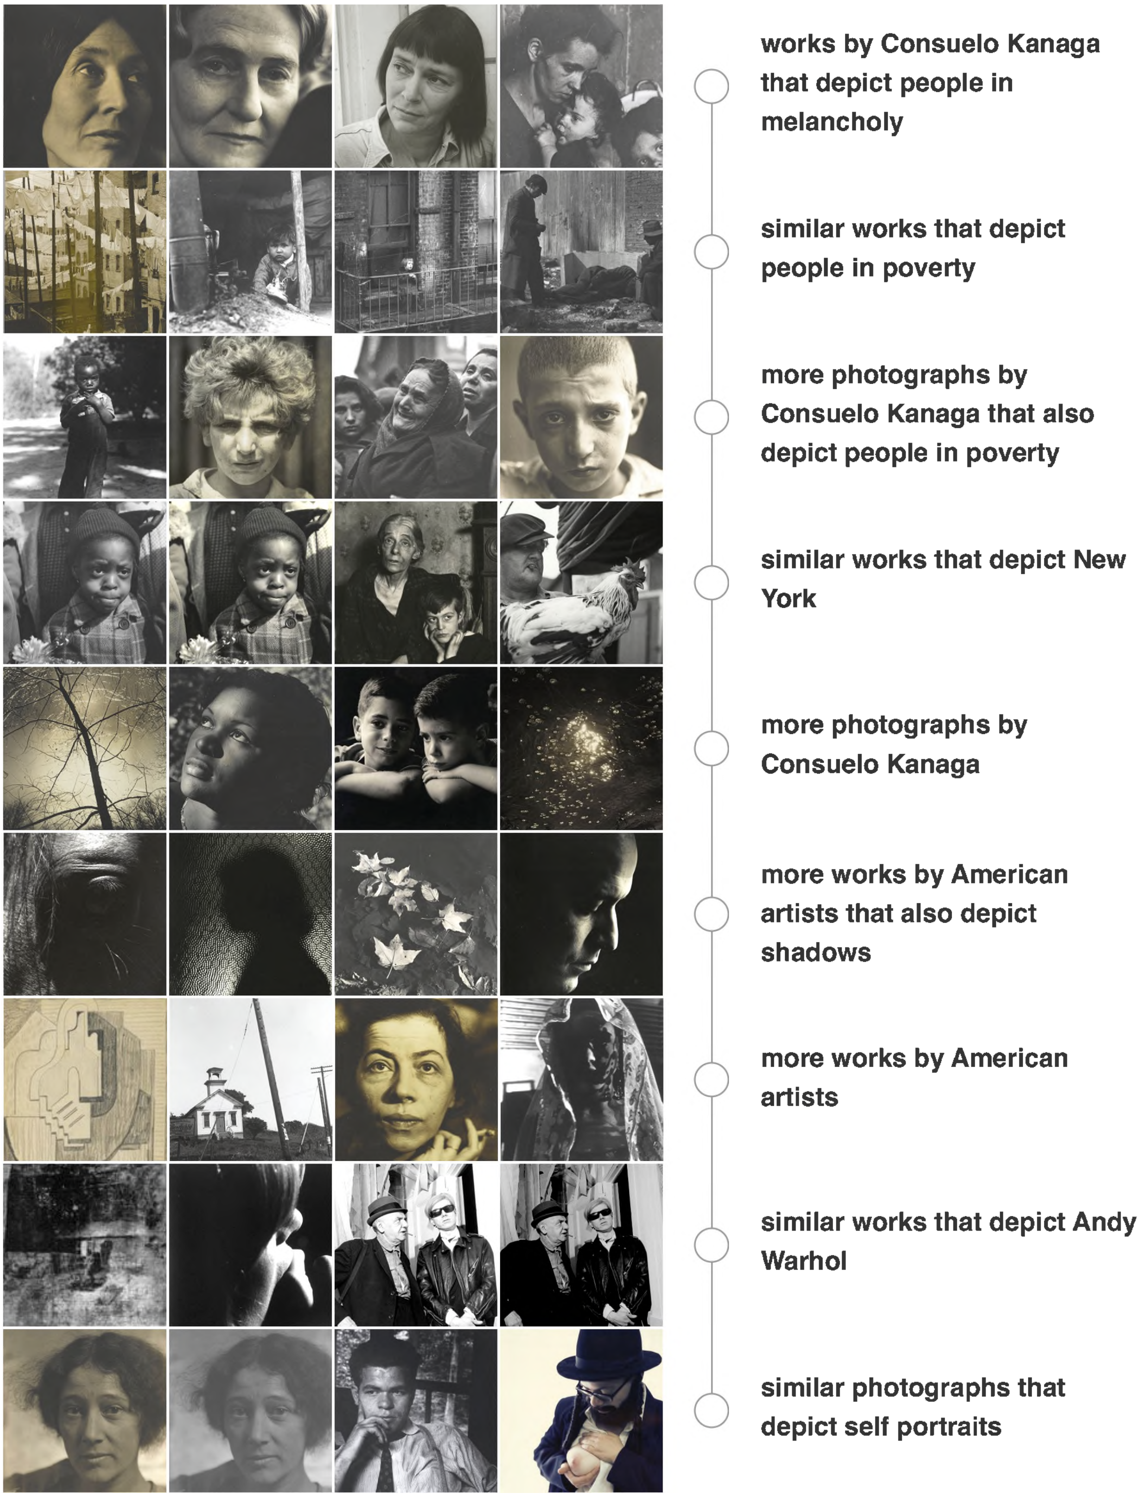

Figure 4.32: A narrative of pathways of the works by Consuelo Kanaga.

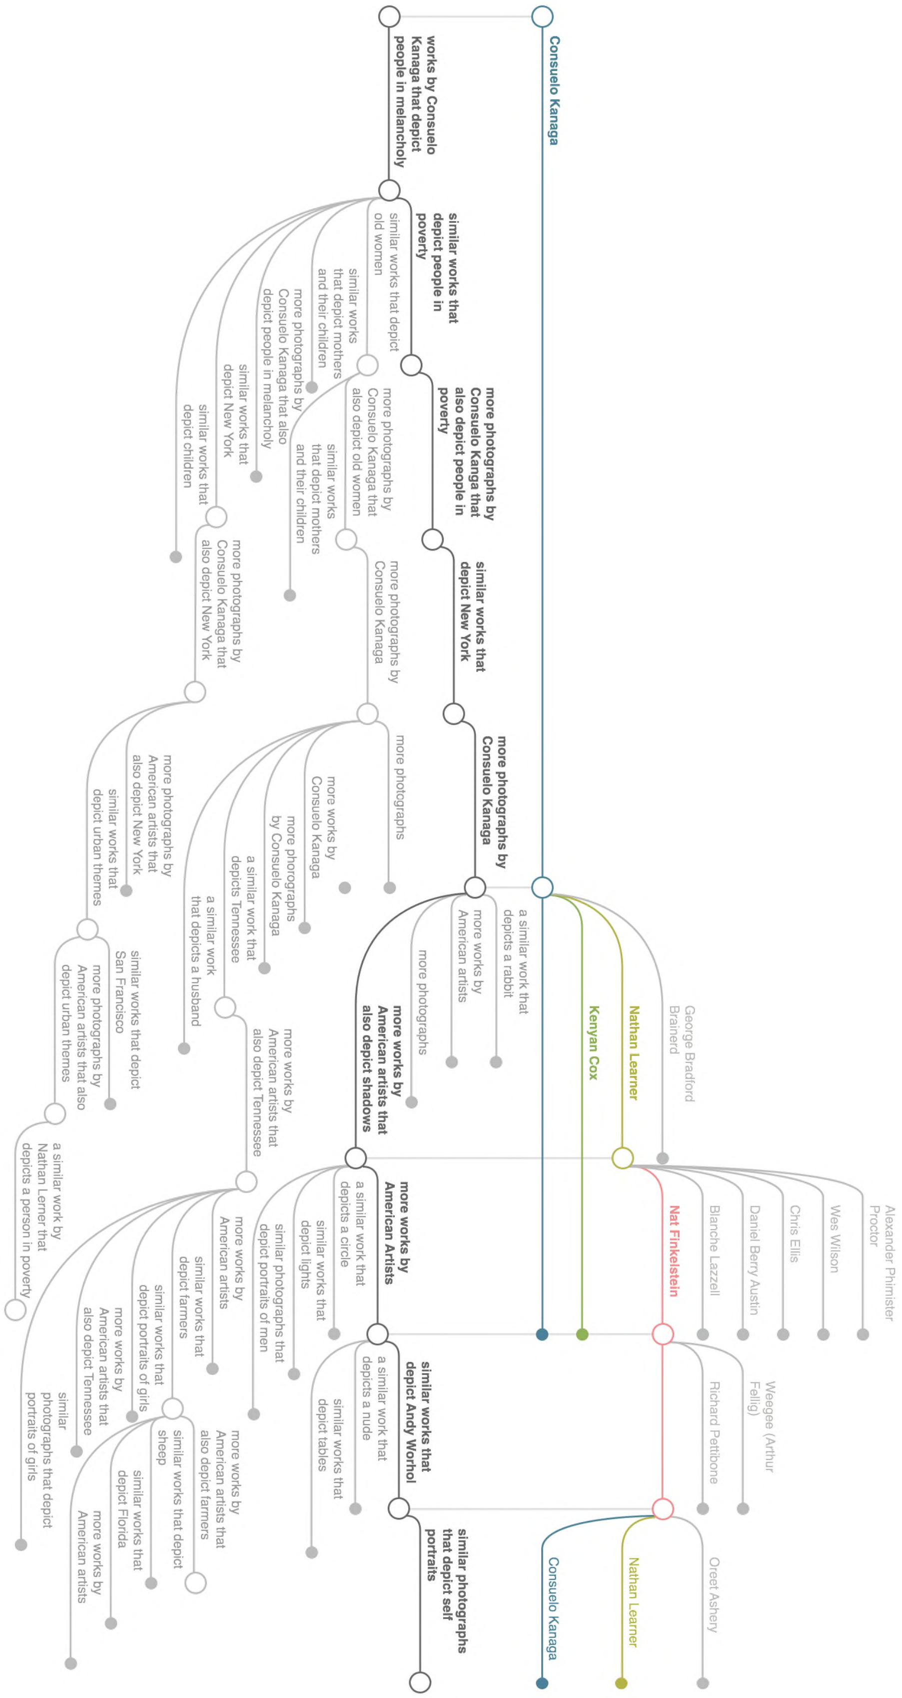

Figure 4.33: A "streams-of-story" visualisation of connecting pathways and concepts from Consuelo Kanaga's works.

Table 4.10: The most frequently occurring visual style tag attributes of works by Consuelo Kanaga.

Attribute	No. Objects
`depicts:visual:high contrast imagery`	39
`depicts:visual:reflections`	38
`depicts:visual:the glare of the Sun`	31
`depicts:visual:shadow`	27
`depicts:visual:silhouette`	24

Nathan Lerner

Like Kanaga, Nathan Lerner also documented the poverty and destitution that was rife in urban neighbourhoods during the period of the Great Depression, and as seen in Kanaga's "streams-of-story" representation in Fig. 4.33, they share a stylistic overlay in their use of dark, shadowy imagery. Lerner was well known for his technical innovations and interplay with lightness and darkness, projections, geometric forms and use of viscous fluid in photography [133], as evidenced by the common themes shown in Table 4.11.

Despite photographing a wide variety of subject matter such as tunnels, cityscapes, amusement park rides and mannequins, Lerner's photography is characterised with geometric forms such as lines and meshes. As shown in the narrative of concepts in Fig. 4.34, many of his works share common visual elements. Curiously, due to these shared attributes and a primarily visual means of interpreting Lerner's work, the narrative draws the connection between the criss-crossing lines of the concept *similar works that depict lines*, their mesh-like forms in the following concept, and the mannequins – many of them veiled or behind meshed obstructions - in the one that follows. A similar narrative of his visual style follows in a "streams-of-story" depiction of connecting pathways in Fig 4.35, with the bottom-most pathway linking works that feature light, darkness and "fluid forms." These connections arise from the primarily visual knowledge embedded in the works: knowledge that was interpreted and then transcribed as attributes during the tagging process.

The narratives of artists Consuelo Kanaga and Nathan Lerner – as presented here by example – demonstrate that the linking of objects, and the chain of concepts that follows, can be based on the explicit data model of artists, objects and their entities as dictated by the museum, and the implicit, visual interpretation of the works, combined with the background knowledge of each artists as obtained via a process of meta-data tagging. Further, the parallels between the visual interpretive

Table 4.11: The most frequently occurring tag attributes of works by Nathan Lerner.

Attribute	No. Objects
`depicts:eye`	16
`depicts:line`	13
`depicts:visual:shadow`	12
`depicts:location:Chicago`	12
`depicts:visual:lines`	12
`depicts:visual:light and darkness`	11
`depicts:visual:abstract forms`	11
`depicts:lighting`	10
`depicts:visual:geometric forms`	10
`depicts:visual:high contrast imagery`	10
`depicts:face`	9
`depicts:location:New York City`	9
`depicts:visual:shapes`	8
`depicts:visual:mesh`	8
`depicts:visual:textured surfaces`	8

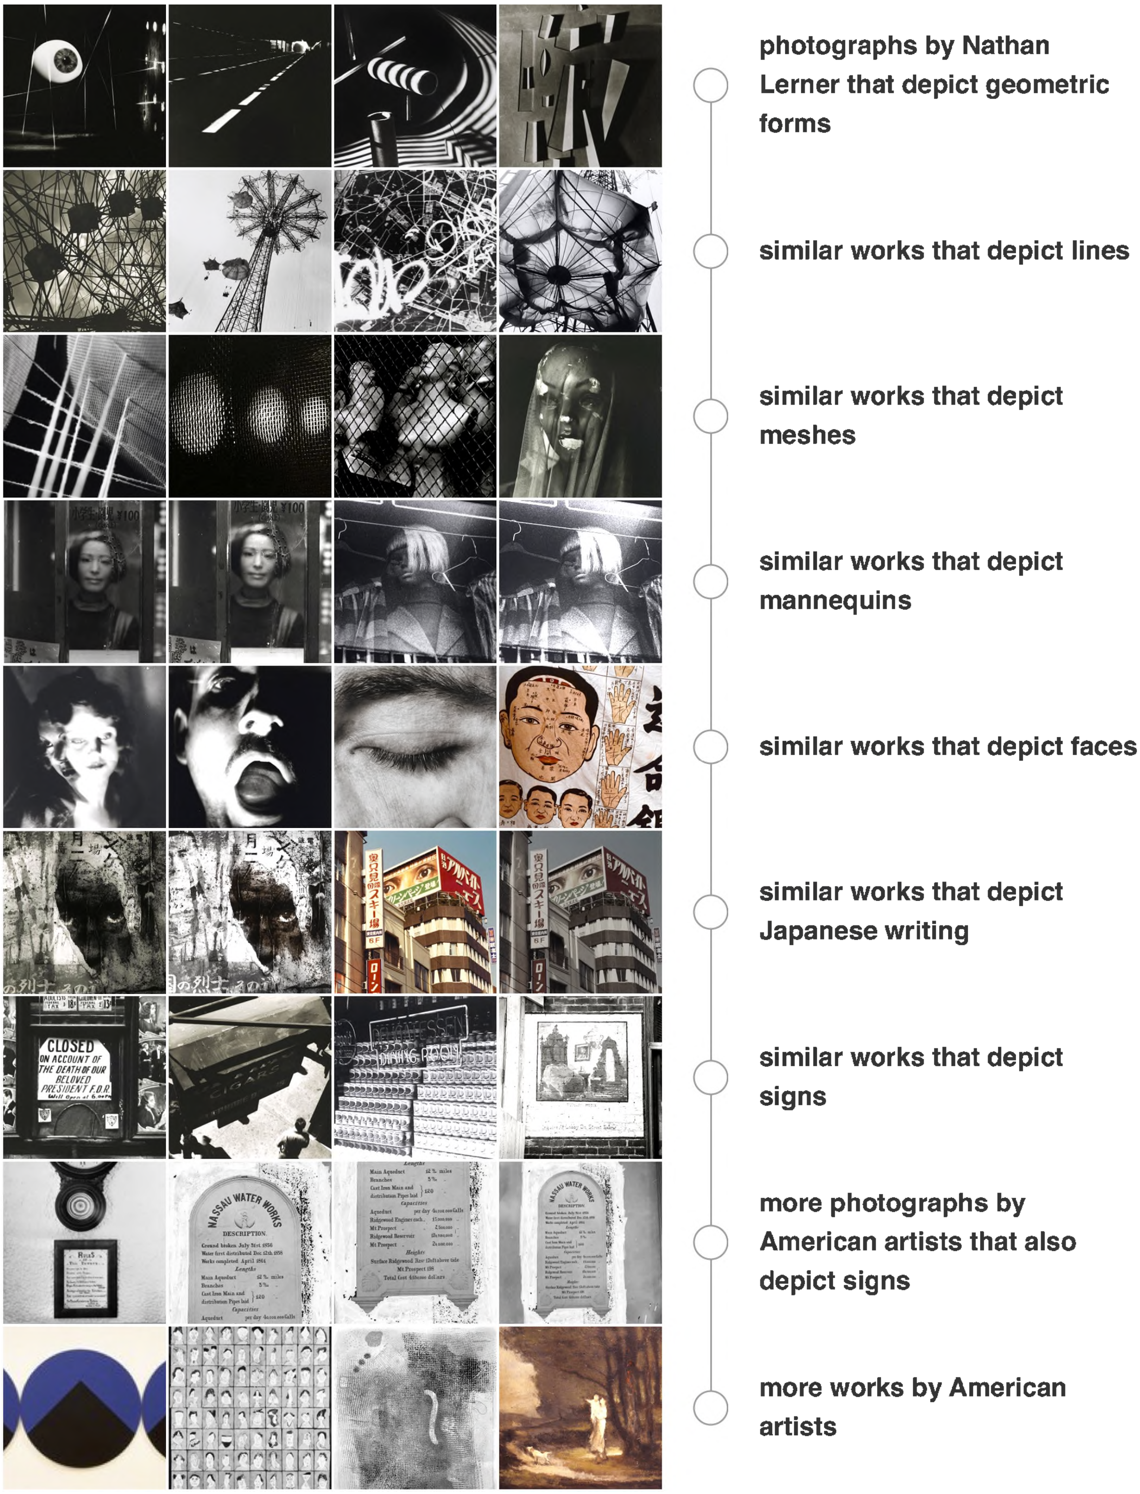

Figure 4.34: A narrative of pathways of the works by Nathan Lerner.

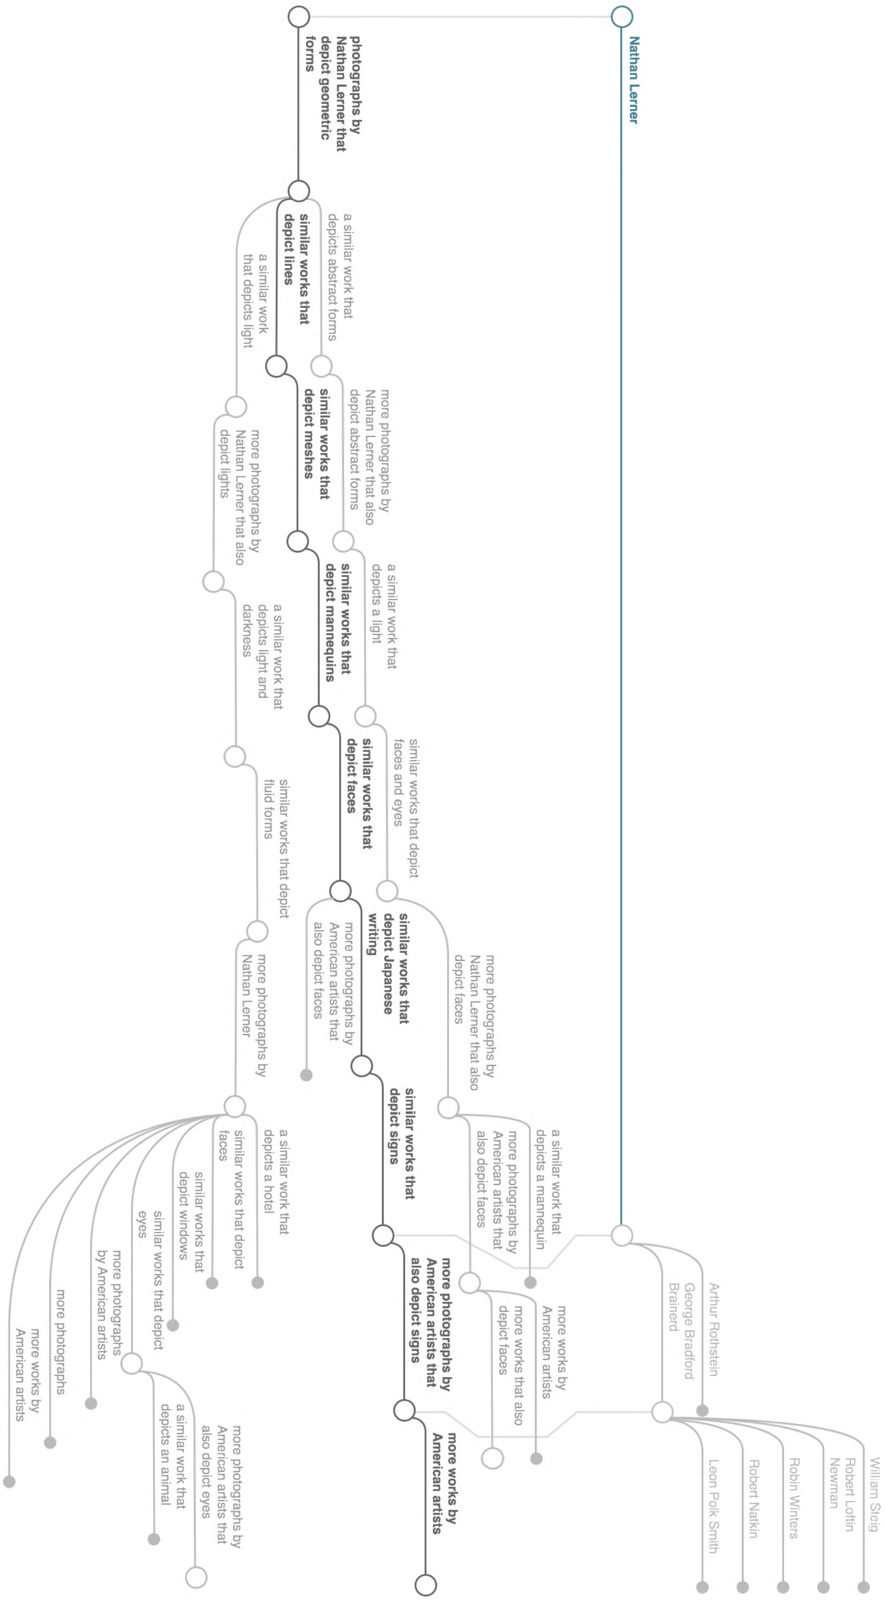

Figure 4.35: A "streams-of-story" visualisation of connecting pathways and concepts from Nathan Lerner's works.

knowledge of the works are made more apparent when linked concepts are presented alongside imagery of the works, as shown in Figs. 4.32 and 4.35 and within the *Brooklyn Museum Canvas*. This sections demonstrates the use of Formal Concept Analysis as a means of linking and analysing subjective and emotional knowledge: a trait that supports Wille's thesis in the use of Formal Concept Analysis to capture and express human thought and interpretation [17].

Chapter 5

Results, Analysis and Discussion

This chapter links the observations from the case studies in Chapter 4 back to the design principles in Section 2.5. Section 5.1 summarises the key findings of the user experience evaluations of the *Virtual Museum of the Pacific* and *A Place for Art* case studies, while Section 5.2 describes each design principle as they are realised through the analysis of the concept designs, their case studies and user experience evaluations.

5.1 Key Findings from User Experience Evaluations

Due to the intentionally taskless nature of these systems, the user experience evaluations for the *Virtual Museum of the Pacific* and *A Place for Art* case studies were formative in nature. The interfaces were assessed based on the participants' reactions to the interaction paradigm, their own personal connection to the collections, the organisational, societal and anthropological issues surrounding the objects, the representation and contextualisation of objects within the interface, and the participants' own personal, emotive, subjective judgements, which included judgements of aesthetics and, particularly within the *A Place for Art* case study, demonstrations of affect, emotion and playfulness.

While the external user experience evaluation of the *Virtual Museum of the Pacific* was semi-scripted in nature – and hence reflected a more task-based approach – all evaluations recognise information seeking as an inherently complex experience that encompasses a wide range of motivations and emotions. The interfaces were considered as cultural artefacts in their own right that are open to reflection and critique [15, 31]. While the organisational, user and technical considerations of these two studies differ sufficiently to warrant their own interpretation, a number of common themes have emerged from both evaluations:

5.1.1 Meaningful engagement with high quality content

This was considered to be the most prominent theme to emerge from the evaluations. In particular, participants from all evaluations generally connected with the content in a meaningful way, and participants from both praised the quality of the imagery, photographs and textual descriptions. They also responded positively to the design of both interfaces that place great emphasis on the imagery of the objects. Participants from both evaluations had a desire to visually inspect the works by panning and zooming, and many held expectations that this is a standard feature that should be made available on all museum websites. Both participant groups reflected that, as a whole, each respective interface highlighted the *diversity* and *value* of their collections.

5.1.2 The importance of context

Secondary to the primacy of high quality content was the context and the meaning that could be derived from it. The majority of participants from both evaluations wanted to know more about the provenance of the objects: the participants in the *Virtual Museum of the Pacific* evaluation wanted to know more about geographical

origins of each artefact – especially among those who self-identified as Pacific Islander who demonstrated very strong, personal connections to the works. For those who participated in the *A Place for Art* study, participants wanted to know more about the artists behind the works, and suggested that the app should provide access to artist biographies and links to their social media accounts. They also demonstrated a particular connection with the physicality of the works that surrounded them on campus.

5.1.3 Non-didactical, non-hierarchical methods of navigation

Museum staff from the internal evaluation of the *Virtual Museum of the Pacific* case study emphasised that the collection needs to be presented in non-didactic ways and employ non-hierarchical navigation methods, and it was anticipated that the *Virtual Museum of the Pacific* could be used as a means of presenting collections in a non-linear manner. However, some of the external participants of the study – the Pacific Islanders and researchers – found that the interface overwhelming and at certain points, disorienting, with no way to 'go back.' In light of this feedback, the user interface for *A Place for Art* was vastly simplified – its implementation of a concept pathway as a lateral, branchable gallery of objects was intended to clearly convey whether a user was moving 'forwards' or 'backwards' through the collection. The more positive response to pathway approach in its evaluation was also a result of its leisurely, non-goal oriented context, rather than semi-structured task-based approach of the *Virtual Museum of the Pacific* evaluation. Participants who interacted with *A Place for Art* recognised it as a highly serendipitous interface, describing the experience as "flowing", "meandering" and "flaneur-like." Other participants remarked that pathways allow them to "choose their own story" and "create their own paths" and responded positively to the fact that the app provides unexpected connections. However, participants were divided on the way the app presents groupings of objects in such a loosely-structured manner – while some praised the 'category-less' nature and fact that works can exist in a number of different concepts, others expected a more structured approach, such as the ability to view an fixed list of artists or mediums – a 'table-of-contents' style view, and then explore individually within the categories.

Based on user feedback from both evaluations, it was recommended that more conventional navigation options, such as a listing of an artists, mediums or locations, would be made available for users as a means of resolving the issue of 'getting lost' within the pathway structures. Further, participants from both evaluations used, or recommended the use of, the ability to backtrack through previous pathway

navigations as a way of overcoming their sense of disorientation within the collection.

5.1.4 Terminology, metadata and interpretation

Museum staff from the *Virtual Museum of the Pacific* evaluation highlighted the need for online museums to embody multiple representation and perspectives of the works. Although its pathway-based navigation allowed the content to be presented in multiple, overlapping contexts rather than a single, fixed, hierarchical classification structure, staff and curators from the museum pointed out that the app's navigation design still reflected its 19th-century classification system. This was due to the way metadata was extracted for this study: although attributes were selected based on the research of each individual object, they were constrained to the museum's *Archaeology and Anthropology* thesaurus. Further, all objects were tagged with their universal *object category* and *type* identifiers, resulting in the two-tiered classification structure dominating the structure of concepts within the concept lattice. However, another taxonomy – the hierarchy of geographical place names – also influenced the connections between concepts. The ability navigate between differing or conceptually related object types across various geographical domains was regarded quite positively by participants who identified themselves as Pacific Islander. Although these geographical place names followed a Western hierarchical convention for naming countries, provinces and regions – the system can also represent concepts based on indigenous representations of place and context. These multiple views can be selected individually, or combined using perspectives (see Section 4.2.1).

The interpretation of concepts within *A Place for Art* were far more subjective. During the evaluation, participants reported an understanding of how different concepts were represented and connected to one another via the use of pathway structures. However, although some participants demonstrated an awareness of the connections between different concepts, they questioned the authenticity or accuracy of these connections. The issues of accuracy concerning object metadata – which corresponded to issues in accuracy in the way pathways were represented – stemmed from the inherently interpretive process of analysing the visual and thematic qualities from the works as presented in the *A Place for Art*. In one case, the visually, but arguably interpretive, description of a set of artworks was considered by the participant pairs to be inaccurate. The problem was compounded when the system branched to a different pathway that, in the minds of the participants, was not completely incongruous with their own expectations of similarity – a notion of similarity that was primarily *visual* rather than *conceptual*: "I pick something thats got clean lines, and then I get things that are abstract and messy.". In such a case, the app provided a didactic response to the user's request for more similar

artworks – effectively ‘telling’ the visitor what the next pathway should be. Other participants, however, preferred to be ‘surprised’ by these connections.

5.1.5 Aesthetics and visual momentum

Participant groups from both case studies praised the visual treatment of the interfaces, responding positively to their minimal designs and the emphasis that is placed on content, rather than interface elements. Participants from the *A Place for Art* evaluations responded positively to the prominence of visual imagery within the app – and the fact that these works can be directly manipulated using tightly connected, responsive touch gestures. Participants described the app as being “very pleasing on the eye” and its movement “like moving through a fluid.” Although participants held similar views in the *Virtual Museum of the Pacific* case study, some noted that its dark minimal visual design did not reflect the colour or the diversity of the Pacific.

5.2 Realising Design Principles

This section presents a descriptive and analytic evaluation of the case studies in Chapter 4 with respect to the design principles in Section 2.5.

5.2.1 Museums as collectors, creators and disseminators of knowledge

The role of the museum is to collect, create and disseminate knowledge about their objects [8] and the material, social, historical and cultural contexts that surround them [6]. This design principle is realised via the use of *CollectionWeb* framework (Section 4.1), which uses Formal Concept Analysis to mine and expose the conceptual relationships between objects.

Formal Concept Analysis mathemetises knowledge and is based on the philosophical logic of human thought and is thus seen fit as a framework creating, linking and augmenting tacit and interpretive knowledge. [17, 97] In a museum setting, the technique provides the mathematical structures for an individual object as a *formal object*, metadata as its *formal attributes* and collections or parts thereof as *formal contexts* (Section 3.1). The construct of the *formal concept* is bound by a set of objects that may possess a certain set of attributes, or a set of attributes shared by a group of objects. Concepts are inherently relational (Section 3.2) and spatial, (Sections 3.3 and 3.4) and thus can be used to describe the *context* of an object, and link concepts with other concepts.

A *formal context* may represent an entire collection, or part of a collection. *Conceptual scales*, implemented as *perspectives* within the *Virtual Museum of the Pacific* (Section 4.2.1), can be used to delimit or combine knowledge from different parts of the collection, or delimit and combine particular interpretations of those collections based on the ability to filter and combine attributes as well as objects. They can be used to represent a number of cultural interpretations of a collection, such as the example reported within the *Virtual Museum of the Pacific* case study where a visitor could choose to browse the Pacific collection under a Western interpretation or an indigenous interpretation, or a combination of the two.

The *Virtual Museum of the Pacific* also allows users to add tags to objects. Tags, which can be represented as *formal attributes*, can be used to both describe the object and shape its relationship with other objects, given that an object's attributes influence and determine its member concepts and relationship with other formal concepts. The idea of using social tagging to shape both the content of an object and its relationships with other objects is a concept that's been explored within the literature. [5, 11, 134]

Formal Concept Analysis can also be used to link subjective, emotional and interpretive knowledge [17]. In the *Brooklyn Museum Canvas* case study, objects were tagged with visual, interpretive and subjective attributes. Concept pathways were then used to describe a *narrative*, linking objects and artists by common themes (Section 4.4.3). The case study also demonstrated that the conceptual relationships produced by Formal Concept Analysis were a reflection of the knowledge of the tagger or curator. In the *Brooklyn Museum Canvas* case study, the narrative of concepts primarily reflected the visual and thematic interpretations of the objects, whereas in the *Virtual Museum of the Pacific* case study, the presentation of navigation options within the user interface primarily reflected in its 19th century classification system.

5.2.2 Highlighting context and connectedness within museum collections

The idea of the virtual museum was originally conceived as a way of showing *connectedness* – the ability to introduce new perspectives on the collection by connecting works of art that are not normally otherwise accessible together. [33, p. 3] This design principle is also realised by the use of Formal Concept Analysis – *formal concepts* can be used to group related objects together that can be connected to other concepts in a *concept lattice* (Fig. 3.2 in Section 3.2) or the *concept pathway* (Fig. 3.6 in Section 3.4). *Connectedness* is also used as a visualisation metaphor – objects are visually 'joined' and connected to one another on the pathways within the *A Place for Art* and *Brooklyn Museum Canvas* interfaces. The concept is also realised through narrative and the natural language expression capabilities of the *CollectionWeb* framework, given the ability to not only describe concepts of related objects, but also describe, in narrative, how those concepts are connected to one another. The *narratives* of concept pathways within the *Brooklyn Museum Canvas* case study also show how different artists are connected to one another via divergent themes, as evidenced in the "streams-of-story" visualisations in Figs. 4.33 and 4.35 in Section 4.4.3.

Connectedness was a quality that also consciously influenced the curatorial and interpretive processes of object selection and annotation. For instance, the process of selecting objects for the *Virtual Museum of the Pacific* case study ensured that there was a significant overlap of object types, materials and geographical provenances so that the collection could represent a diverse cross-section of material Pacific culture and provide opportunities for new connections to form across multiple facets and dimensions. This intent was also reflected in the visual and subjective tagging processes for the *Brooklyn Museum Canvas* case study that used a specially built tool that assisted the tagger in ensuring that themes were applied consistently so

that they could overlap and connect with one another.

Participants in both user experience evaluations also expressed a desire to connect with the works: in the *Virtual Museum of the Pacific* evaluation, participants were overwhelmingly interested in the geographic provenances of the objects and wanted to connect with other objects from the same region. Participants in the *A Place for Art* evaluation expressed a desire to connect with artists via social media channels, and emphasised highlighting the *social connections* that surround each artist – "The best way for this to work is to have it connected: you want artists connected, you want teachers connected; you want students connected – it just keeps going." Some participants, demonstrating a connection to the physicality of works on campus, wanted to visit the artworks in person after viewing them within the app.

5.2.3 Pathways as a mechanism that semantically structures content

The *pathway* metaphor has been posited a way of semantically structuring content within museum and cultural collections. [12, 90]. This design principle is realised via the use of *concept lattices and their neighbourhoods* (Section 3.2) and *concept pathways* (Section 3.4) as a means of linking objects and their equivalencies. The *concept pathway* was conceptualised for the specific purpose of realising this design principle, replicating the branches and divergent possibilities that one would encounter while navigating pathways within a connected landscape (Section 2.4.5). Central to the idea of these structures is the notion of *directionality* – as outlined in Chapter 3 – concept lattices and concept pathways are not used solely the purposes of data analysis and visualisation, but as a means to create structures that can be traversed and explored. Figs. 3.2 and 3.3 provide examples of directionality, highlighting examples of conceptual structures and the user journeys that traverse them. *Directionality* is also conveyed in the way the *CollectionWeb* framework expresses natural language phrases of concepts relative to one another (e.g. "works that depict people in poverty" $\rightarrow$ "similar works from New York that also depict urban themes"). Under the principle of *directionality*, a set of concepts that consists of individual steps within a concept lattice or a concept pathway can be construed as a *narrative*, as shown in Sections 4.1.2 and 4.4.3.

Although the *Virtual Museum of the Pacific* intended to depict the pathway metaphor via navigation of upper and lower neighbours within a concept lattice, the concept of the pathway as an interface element was made more explicit in the design of the *A Place for Art* and *Brooklyn Museum Canvas* interfaces. From the users' perspective, participants of the *Virtual Museum of the Pacific* case study noted

that its interface allowed them to easily comprehend groups and clusters of objects. The realisation of the pathway metaphor as a way of structuring the navigation experience was acknowledged by participants from the *A Place for Art* study, who noted that the groups of works were connected 'by theme', with some noting that multiple paths can be created from a single work.

5.2.4 Pathways as a metaphor for creative exploration

The *pathway* metaphor has also been posited a way of affording creative exploration. [12, 90] This design principle is realised in the way *conceptual neighbourhoods* and *concept pathways* present the user with multiple navigation options. The *concept pathway*, in particular, offers opportunities for divergent exploration in the sense that new pathway branches always present the user with concepts that contain new objects (Section 3.4). This raises the possibility of encountering new and unexpected connections that can spur from existing points of interest – a quality that characterises serendipitous interfaces [74].

In the *A Place for Art* user evaluation, participants described the experience as "flaneur-like" as they would "unravel pathways." Users often responded to this idea with a level of emotional affect, saying that the app provides "a level of unexpectedness" and that it "encourages the passion in me." Some participants interpreted a pathway journey as a sense-making process: "I want to use it more and see where the connections are."

The interfaces of both the *Virtual Museum of the Pacific* and *A Place for Art* apps explicitly avoided any notion of *linearity*: the idea that a journey through the collection should always have fixed start and end points. According to the research conducted by Goodale et al. [12], linearity was not seen as an ideal principle, as it was seen that it could limit the potential of pathways as devices for exploration. In spite of this, users of both evaluations reported that they felt 'lost' as they navigated through the collection, and some preferred to begin from a well-defined, and fixed, 'home state', such as a screen that highlights the main artists, mediums or geographical areas of a collection.

5.2.5 Designing for the information flaneur

The *information flaneur* a personification of a curious, creative and critical information seeker that is based on the literary persona of the *flaneur*. Rather than seek to fulfill knowledge gaps or solve problems, the *information flaneur* approaches a situation at a leisurely pace, whose actions and directions are motivated by his or her creative mind. [15] This form of information seeking can be described as directionless, yet serendipitous and curious, and is based on an information seeking model

that combines horizontal exploration with vertical immersion. This design principle is realised through the use of non-linear information spaces, a consideration of pleasing design and aesthetics, and an incorporation of serendipitous elements that enable connections that lead to the unexpected [74].

Given the importance of emotion and motivation in designing for the *information flaneur*, both the content and the emotion of the participants' feedback was considered, resulting in an analysis of the most positive and negative experience of the app in terms of affect. As reported in Section 4.3.2, the most positive experiences, as dictated by the participants, were the app's aesthetics and visual design, a personal engagement with the collection, and the perception that the app offers free and open exploration of its artworks. A number of comments made by the participants that reflected their desire to explore and be surprised by the sometimes 'unexpected' connections that the app makes, although not all participants were comfortable with this and instead preferred more conventional methods of navigation.

The information flaneur is characterised by a creative mind and their capacity to make sense through multiple facades and interpretations. The ability for Formal Concept Analysis to represent objects in multiple, overlapping contexts, reflects this conceptual view – a view also shared by some participants from the *A Place for Art* user experience evaluation. As noted in Section 4.3.2, it was reported that some participants enjoyed discovering new connections, likening the experience as "moving through a spiral", although some felt 'lost' as they were navigating through the collection, and required a mental overview before they could adequately begin their journey. Others, noting they could draw multiple paths from a single work, felt that they could just get a sense of collection by simply moving around and experimenting, bridging their perceptions and experiences back to the sensemaking view of the *information flaneur* – as the flaneur explores, he or she gets a sense of what city life is about.

Orientation

Dörk et al. [18] also stress the importance of such interfaces to provide a sense of orientation to users. The negative experiences of some participants from both the *A Place for Art* and *Virtual Museum of the Pacific* evaluations who reported a sense of disorientation suggest that the interfaces did not sufficiently provide enough overview or orientation to the user. In the *Brooklyn Museum Canvas* case study, the issue attempted to be addressed through the placement of fixed, perpendicular paths, each providing a contextual navigation link to an adjoining pathway node within its structure, although the interface still does not provide an overview or 'birds-eye view' of the entire visualisation.

Visual Momentum

Visual momentum refers to the way interfaces convey smooth transitions from one context or state to another delay or without breaking the users sense of visual or conceptual perception. In the *Virtual Museum of the Pacific* case study, visual momentum is manifested via a grid of objects that dynamically grow and shrink in response to user queries. In the *A Place for Art* case study, visual momentum is manifested through touch interactions – users can browse related artworks on a pathway that 'slide' back on forth, and the actual path animates by 'disconnecting' and 'reconnecting' to another path. In the *Brooklyn Museum Canvas* case study, visual momentum is maintained through the use of a movable 'camera' that zooms, pans and rotates to convey the affordances of object inspection, path browsing and path navigation, respectively.

Serendipity

The *Virtual Museum of the Pacific*, *A Place for Art* and *Brooklyn Museum Canvas* interfaces all inherit qualities of serendipitous digital environments [74] – they all emphasise prominent visual imagery; provide links to tangentially related concepts via the use of tag clouds and textual links (*The Virtual Museum of the Pacific*), visual, connecting elements (*The Brooklyn Museum Canvas*) and visual information surrogates (*A Place for Art*); and allow new and tangential connections to form via the navigation of neighbouring concepts. In their guidelines for serendipitous interfaces that appeal to the *information flaneur*, Dörk et al. [15] suggest that the placement of novel or unusual resources that relate to a user's previous interactions could increase serendipity in that it leads users down unexpected paths. *Concept pathways* follow this principle in that new pathway connections must always display concepts that contain new objects – increasing the possibility that the user may encounter an unexpected pathway.

5.2.6 Pliability

Pliability is characterised as the degree that an interface, as realised through its interaction, feels involving, malleable and serendipitous [135]. The specific aspects of pliability covered here describe its visuo-tactile dimensions: the tight coupling of user action and interface response, and the conflation of interface and contents.

The challenge in assessing *pliability* stems from its perceptual qualities – a quality that can only be recognised in the present moment – the 'here and now' – of interaction. Although pliability is not an inherent quality manifested *within* an artefact, some of the design features of the *Virtual Museum of the Pacific*, *A Place for*

Art and *Brooklyn Museum Canvas* interfaces were intended to engender a sense of pliability in their use.

In the *Virtual Museum of the Pacific*, users move through a concept lattice via navigation of its conceptual neighbours. Every time a user moves to a neighbouring concept, the expanded set of objects animates as it 'grows' and 'retracts.' The intention was to engender a sense that the user could manipulate the grid of images by including or excluding terms. Some users reported being able to sense a connection between different types of objects from different locations via this interaction method, as demonstrated by example in Fig. 4.13.

In the design of *A Place for Art*, a conscious decision was made to reduce the number of navigation options and focus on a far more direct interaction using touch gestures, rather than a 'point-and-click' interface. The vision, as implemented in the iPad app, was to allow visitors to interact with a sliding pathway, and, using their fingers, 'branch' out an connect objects with other objects. The user evaluations in Section 4.3.2 reported positive aesthetic experiences from users – describing the experience as "meandering", "whirling and twirling" and like "moving through a fluid." Much consideration was given to the design of the interaction that, when an artwork was double-tapped on, it caused that work to be 'popped out' and displaced from its position in the pathway, where the camera then rotates 90 degrees and a new pathway is drawn in. At the point, the user can then swipe left on the new pathway, or tap on the old pathway (now arranged in a perpendicular manner) to 'fold' the new work back in. This interaction – although technically complex and costly to implement – was to further reconcile the users perception of the continuity of two navigation points, and to provide a realistic, sensual experience of being able to 'shape' the path that's in front of them.

Similar elements of continuous, fluid interaction, guided the design and implementation of the *Brooklyn Museum Canvas*. However, rather than have users directly branch out and 'create' path structures, the pathways within the *Brooklyn Museum Canvas* are generated based on a user's current point within the information landscape along with their previous navigations. In this sense, the *Brooklyn Museum Canvas* presents more of a 'wandering tour' approach – one that was intended, through its visual and interaction design – to replicate the 'white cube' aesthetic of a gallery and the experience of wandering through it. The experience is controlled via a use of a camera that can be panned, rotate and zoomed across interconnecting pathways. The fluidity of the experience is contrasted with the conventional layout and navigation design of the Brooklyn Online Collections site: while the site presented the same information about the objects and also provided a way of exploring objects via its 'Related Tags' feature (see Fig. 4.22), it followed standard website design conventions that treated objects as individual records that required a visual

refresh every time the user wanted to move from one page to another. In demonstrating the concept of *pliability* as applied to geographical information, Löwgren [135] demonstrates a similar comparison between the static pages of an older version of the Eniro mapping software, and the expansive, sometimes dizzying experiences of crossing terrains in Google Earth.

Chapter 6

Conclusion

This chapter concludes the thesis and consolidates its research outcomes. In Section 6.1.2, the research is summarised as a design theory that describes pathway-based navigation within online museum collections. Following from that, Section 6.2 presents a discussion and reflection of key points from this study.

6.1 Research Outcomes and Contributions

6.1.1 Summary of research outcomes and contributions

Drawing from museological and creative information seeking perspectives, a set of design principles were devised to support the pathway metaphor as a mechanism that semantically structures content and affords creative exploration within museum collections. Following a design science and concept-driven interaction design approach, these principles informed the development of a framework for extracting museum knowledge along with three different interfaces that support the exploration of objects from three different museum collections.

CollectionWeb is a software framework that was developed to support conceptually enriched navigation within museum collections. The framework implements Formal Concept Analysis to generate conceptual structures that can be explored and traversed by an end user. To demonstrate scalability, a performance evaluation of one of its key algorithms was tested on data-sets from actual museum databases. To address the specific design principle of the pathways as a mechanism that semantically structures content, the conceptual structure of the *concept pathway* (Section 3.4) was also devised.

Three separate interfaces were designed and implemented as working software: *The Virtual Museum of the Pacific*, *A Place for Art* and the *Brooklyn Museum Canvas*. These three interfaces all relied on the *CollectionWeb* framework to generate pathways and express them in natural language. User experience evaluations were

conducted for the *Virtual Museum of the Pacific* and *A Place for Art* interfaces, each highlighting the organisational context and the users' expectations, evaluations and interpretations of the interfaces. The design of the *Brooklyn Museum Canvas* is inspired by the additional design principle of *parallelism*. The case study also experimented with how connecting pathways can be used to express *narratives*. In order to observe how visual and interpretive knowledge could be represented using concept pathways, a set of works from two featured artists were interpreted and manually tagged – the resulting pathway narratives have been described and reported.

The design principles were revisited in Section 5.2, where key elements of each case study were addressed and described in the form of a descriptive, analytic evaluation. A summary of the key issues arising from both user evaluations was also presented. The inception of design principles from the literature – along with the creation and evaluation of the artefacts – contributed to the development of a design theory.

6.1.2 A design theory

In some disciplines, such as practice-based [136] and interaction design research [20, 31, 86], the designerly knowledge of an artefact is realised through its form, function, and materiality. Scholars such as Cross [137] state that such knowledge exists within the actual artefact itself. In recognition of this, Gregor and Jones [32] describe a way to build a design theory from this knowledge. This framework is used to articulate the research contributions of this thesis as a design theory. According to Gregor and Jones [32], the criteria for a design theory are as follows:

- *Purpose and scope* – the context, background and motivation of the theory.
- *Principles of form and function* – key design elements.
- *Artifact mutability* – a description of how the artefact(s) would work in different instantiations. From an interaction design perspective, this refers to how an artefact changes, adapts or responds to different data or use contexts.
- *Testable propositions* – a set of questions to be asked or a set of principles of instrumentality, form and materiality to be addressed through an artefact's instantiation and/or evaluation.
- *Justificatory knowledge* – kernel theories that support the *purpose and scope*.
- *Principles of implementation* – the way key *principles of form and function* are implemented.

- *Expository instantiation* – the artefact itself.

Specific components of this theory have been addressed in sections throughout this thesis. Chapter 2 covers the *purpose and the scope* of the theory, addressing literature that describes the role of knowledge from a museological perspective along with exploratory perspectives on information seeking, play pleasure and aesthetics. Using the kernel theories of information landscapes, (Section 2.4.2) the information flaneur, (Section 2.4.3), pliability (Section 2.4.4) and pathways, (Section 2.4.5) the key design principles are described (Section 2.5), which represent its *principles of form and function.* Given that the research follows a concept-driven approach, these design principles are also *testable propositions* and are realised by three different interfaces: *The Virtual Museum of the Pacific*, *A Place for Art* and the *Brooklyn Museum Canvas*, which are all based on different instantiations of the *CollectionWeb* framework – demonstrating *artifact mutability.* Each of the case studies in Chapter 4 describes how the design principles are implemented, with a review of key design features with respect to the principles in Section 5.2. Finally, the design artefacts – the pieces of working software – represent the *expository instantiation* of the design theory.

6.2 Discussion and Reflection

In addition to the outcomes and design theory drawn from this research, a number of reflections are drawn, all based on an analysis of the key findings from Chapters 4 and 5:

- Regardless of the interaction design and visualisation methods employed, visitors in an exploratory context respond positively to high quality content and prominent visual imagery, a finding that was consistent with prior research [37, 74].

- The importance of providing context cannot be understated – participants readily connected with key contextual elements of the objects that felt important to them. The case studies in this thesis mostly provided context in the form of *object-to-object* or *concept-to-concept* relationships. While such relationships may be meaningful to visitors who could articulate their preferred contextual dimension – geographical provenance in the case of the *Virtual Museum of the Pacific* case study, or artists in the case of the *A Place for Art* case study – such elements may also be also be intangibly or experientially represented through other means, such as the ability to experience the physicality of the works, or the opportunity to deep-link and interact with artists

via social media, as was reported in the *A Place for Art* case study. Although such questions have been addressed in other venues [2, 37], further research could be used to suggest how, and to what degree, serendipitous digital environments in online collections facilitate meaning-making that goes beyond the exploratory capabilities of the interface.

- As reported in the *Virtual Museum of the Pacific*, *A Place for Art* and *Brooklyn Museum Canvas* case studies, museum metadata that has been processed with knowledge representation frameworks are still subject to the same issues of terminology, interpretation and access as metadata represented by any other means. Some participants who took part in the *A Place for Art* user evaluation study reported that some of the pathway descriptions did not match their own interpretations of the artworks or the way they were connected to other paths. The gap between the knowledge represented by the app, and the users' own interpretive knowledge of the works could be resolved through social tagging [11, 39] – although key challenges still arise in normalising tagging data to a point that would make it usable for Formal Concept Analysis. [138, 139]

- The issue of *subjective* and *interpretive* knowledge was further explored within the *Brooklyn Museum Canvas* case study in that, although Formal Concept Analysis could be used to provide implicit connections between groups of objects and promote serendipity, those connections are reflective of the interpretive knowledge processes of the person (or communities) analysing, interpreting and tagging the objects – an idea that has also been explored in other research. [5, 39] Formal Concept Analysis, in and itself, cannot create new knowledge, but it can be harvested in a way that allows new insights to form in the mind of the user who is interpreting or interacting with that information space. This approach forms part of the information seeking and analysis activities known as *conceptual knowledge processing* [140] – an approach that has been explored in other case studies. [17, 98]

- Although exploratory interfaces deliberately lack design principles [30], there were a number of usability related issues that were consistent across both user experience evaluations when taking into account the context of a museum visitor. Despite being presented as serendipitous interfaces, users still requested – and sometimes preferred – options of navigation or exploration in what they considered to be established interaction design conventions for online museums such as search boxes, an 'index' or 'contents' style page, or a visual overview of the collection. A lack of overview, which in some cases lead to a lack of orientation within the information space, was consistently considered the most negative aspect of the user experience. This finding was

consistent with the need to provide a sense of overview and context for the *information flaneur.* [15]

- On a related issue, users' perceptions of exploratory interfaces can be varied or even polarizing. As reported in the *A Place for Art* case study – some participants responded to the idea of unstructured exploration in an overwhelmingly positive manner, whereas others required a more structured overview and become frustrated when they could not make a sense of the collection or determine *why* the pathways were structured in a such a manner. A possibility arising from these disparities may be the psychological profiles of participants in relation to their views and expectations on uncertainty and serendipity – a potential avenue that could warrant further research.

- The conceptualisation, representation and evaluation of the pathway metaphor for navigation could have been built and tested in a purely instrumental fashion – for example, by creating simple, functional, 'Web 1.0' style prototypes that employed minimal considerations to visual design. However, the decision to include the elements of visual momentum and aesthetics in their design was in recognition of the interfaces as aesthetic [13] and cultural artefacts [16] in their own right, where such elements were to be interpreted by users in addition to their instrumental qualities. Some participants viewed these artefacts through the lens of form and materiality [141] - for instance, within the *A Place for Art* evaluation, several participants made comments reflecting the aesthetics of their experience, the 'smoothness' of the pathways, and their personal engagement with the works from the collection.

These discussion points provide a final reflection and analysis of the work conducted in this thesis, linking some of the key findings to more tangential areas of research in the literature. It is anticipated that these points provide the basis of more research that uncovers new design possibilities that connect and engage audiences with museum collections.

Bibliography

(1) F. Cameron, "Wired collections - the next generation", *Museum Management and Curatorship*, 2001, **19**, 309–312.

(2) B. Colquhoun, Museums and the Web 2013: Proceedings, ed. N. Proctor and R. Cherry, Archives & Museum Informatics, Portland, OR, 2013.

(3) B. Ennis-Butler, S. Hinton, and M. Whitelaw, Proceedings of the 2011 Conference of the Australian Council of University Art and Design Schools, ed. G. Bull, 2011.

(4) L. Manovich, J. Douglass, and F. Kuester, "Cultural Analytics in Large-Scale Visualization Environments", *Computer*, 2011, **44**, 39–48.

(5) S. Cairns, "Mutualizing Museum Knowledge: Folksonomies and the Changing Shape of Expertise", *Curator: The Museum Journal*, Jan. 2013, **56**, 107–119.

(6) S. Pearce, in *Interpreting Objects and Collections*, ed. S. Pearce, Routledge, London, 1994.

(7) W. E. Washburn, "Collecting information, not objects", *Museum News*, 1984, **62**, 5–15.

(8) G. F. MacDonald and S. Alsford, "The Museum as Information Utility", *Museum Management and Curatorship*, 1991, **10**, 305–311.

(9) G. Anderson, *Reinventing the museum: Historical and contemporary perspectives on the paradigm shift*, Rowman Altamira, 2004.

(10) M. Ross, "Interpreting the new museology", *Museum and Society*, 2004, **2**, 84–103.

(11) S. Chan, Museums and the Web 2007: Proceedings, 2007.

(12) P. Goodale, P. Clough, N. Ford, M. Hall, M. Stevenson, S. Fernando, N. Aletras, K. Fernie, P. Archer, and A. D. Polo, Proceedings of the 2nd European Symposium on Human-Computer Interaction and Information Retrieval, 2012, vol. i, pp. 75–78.

(13) L. E. Udsen and A. Jø rgensen, "The aesthetic turn: unravelling recent aesthetic approaches to human-computer interaction", *Digital Creativity*, 2005, **16**, 205–216.

(14) J. Löwgren, "Pliability as an experiential quality: Exploring the aesthetics of interaction design", *Artifact*, 2007, **1**, 70–82.

(15) M. Dörk, S. Carpendale, and C. Williamson, Proceedings of the 2011 annual conference on Human factors in computing systems, 2011, pp. 1215–1224.

(16) J. Murray, *Inventing the Medium*, MIT Press, Cambridge, MA, 2012.

(17) R. Wille, in *Formal Concept Analysis*, ed. B. Ganter, G. Stumme, and R. Wille, Springer Berlin / Heidelberg, 2005, vol. 3626, pp. 47–70.

(18) M. Dork, S. Carpendale, and C. Williamson, "Visualizing explicit and implicit relations of complex information spaces", *Information Visualization*, Nov. 2011, **11**, 5–21.

(19) A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research", *MIS Quarterly*, 2004, **28**, 75–105.

(20) E. Stolterman and M. Wiberg, "Concept-Driven Interaction Design Research", *Human-Computer Interaction*, Apr. 2010, **25**, 95–118.

(21) A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research", *MIS Quarterly*, Mar. 2004, **28**, 75–105.

(22) M. Akesson, K. Kautz, and C. Eriksson, ICIS 2010 Proceedings, 2010.

(23) D. Hovorka, in *Incommensurability and Multi-paradigm Grounding in Design Science Research: Implications for Creating Knowledge*, ed. J. Pries-Heje, J. Venable, D. Bunker, N. Russo, and J. DeGross, Springer Berlin-Heidelberg, 2010, vol. 318, pp. 13–27.

(24) K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A Design Science Research Methodology for Information Systems Research", *Journal of Management Information Systems*, Dec. 2007, **24**, 45–77.

(25) B. Kuechler and V. Vaishnavi, "On theory development in design science research: anatomy of a research project", en, *European Journal of Information Systems*, Oct. 2008, **17**, 489–504.

(26) J. Pries-Heje and R. Baskerville, "The design theory nexus", *MIS Quarterly*, Dec. 2008, **32**, 731–755.

(27) M. K. Sein, O. Henfridsson, S. Purao, M. Rossi, and R. Lindgren, "Action design research", *MIS Quarterly*, Mar. 2011, **35**, 37–56.

(28) W. Aalst and A. Kumar, "XML-Based Schema Definition for Support of Interorganizational Workflow", *Information Systems Research*, 2003, **14**, 23–46.

(29) S. Greenberg and B. Buxton, Proceedings of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08, ACM Press, New York, New York, USA, 2008, p. 111.

(30) D. A. Norman and R. Verganti, Proceedings of the 5th Conference on Designing Pleasurable Products and Interfaces, 2011, pp. 1–19.

(31) H. Jung and E. Stolterman, Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems - CHI EA '11, ACM Press, New York, New York, USA, May 2011, p. 399.

(32) D. Jones and S. Gregor, *The Anatomy of a Design Theory*, 2007.

(33) W. Schweibenz, Knowledge Management und Kommunikationssysteme, Workflow Management, Multimedia, Knowledge Transfer. Proceedings des 6. Internationalen Symposiums f{ü}r Informationswissenschaft (ISI 1998), ed. H. Zimmermann and B. Schramm, 1998, pp. 3–7.

(34) E. Hooper-Greenhill, *Museums and the Shaping of Knowledge*, Routledge, 1992.

(35) G. MacDonald and S. Alsford, in *A Different Drummer: Readings in Anthropology with a Canadian Perspective*, 1987, pp. 41–47.

(36) S. Styliani, L. Fotis, K. Kostas, and P. Petros, "Virtual museums, a survey and some issues for consideration", *Journal of Cultural Heritage*, Oct. 2009, **10**, 520–528.

(37) M. Skov, Ph.D. Thesis, Royal School of Library and Information Science, 2009.

(38) J. Trant and B. Wyman, Proceedings of the International World Wide Web Conference, Edinburgh, Scotland, 2006.

(39) J. Trant, *Tagging, Folksonomy and Art Museums: Results of steve.museum's research*, tech. rep., University of Toronto, 2009.

(40) P. Marty, "Museum Websites and Museum Visitors: Digital Museum Resources and their Use", *Museum Management and Curatorship*, 2007.

(41) B. J. Soren and Canadian Heritage Information Network, "Best practices in creating quality online experiences for museum users", *Museum Management and Curatorship*, 2005, **20**, 131–148.

(42) G. Schreiber, A. Amin, L. Aroyo, M. van Assem, V. de Boer, L. Hardman, M. Hildebrand, B. Omelayenko, J. van Osenbruggen, A. Tordai, J. Wielemaker, and B. Wielinga, "Semantic annotation and search of cultural-heritage collections: The MultimediaN E-Culture demonstrator", *Web Semantics: Science, Services and Agents on the World Wide Web*, 2008, **6**, 243–249.

(43) L. Aroyo, N. Stash, Y. Wang, P. Gorgels, and L. Rutledge, in *The Semantic Web*, ed. K. Aberer, K.-S. Choi, N. Noy, D. Allemang, K.-I. Lee, L. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber, and P. Cudré-Mauroux, Springer Berlin / Heidelberg, 2007, vol. 4825, pp. 879–886.

(44) Y. Wang, N. Stash, R. Sambeek, Y. Schuurmans, L. Aroyo, G. Schreiber, and P. Gorgels, "Cultivating Personalized Museum Tours Online and On-Site", *Interdisciplinary Science Reviews*, 2009, **34**, 139–153.

(45) C. Andujar, A. Chica, and P. Brunet, "User-interface design for the Ripoll Monastery exhibition at the National Art Museum of Catalonia", *Computers & Graphics*, 2012, **36**, 28–37.

(46) M. Ljiljana, M. Aleksic, and M. Tankosic, 2nd International Conference on Sustainable Tourism and Cultural Heritage, Brasov, Romania, 2013, pp. 322–326.

(47) D. Goldfarb, M. Arends, J. Froschauer, and D. Merkl, "Revisiting 3D information landscapes for the display of art historical web content", *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology - ACE '11*, 2011, 1.

(48) J. A. Wise, J. J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow, Proceedings on Information Visualization, IEEE, 1995, pp. 51–58.

(49) K. Kwastek, Proceedings of the 10th Conference on Computers and the History of Art (CHART '03), 2003.

(50) M. Whitelaw, *Towards Generous Interfaces for Archival Collections*, 2012.

(51) L. Manovich, "Media Visualization: Visual Techniques for Exploring Large Media Collections", *Media Studies Futures*, 2011, 1–21.

(52) F. Amigoni and V. Schiaffonati, "The minerva system: A step toward automatically created virtual museums", *Applied Artificial Intelligence*, 2009, **23**, 204–232.

(53) G. Hoptman, in *Sociomedia: Multimedia, Hypermedia, and the Social Construction of Knowledge*, ed. E. Barrett, The MIT Press, 1992.

(54) M. Sharples, P. Lonsdale, J. Meek, P. Rudman, and G. Vavoula, Proceedings of 6th Annual Conference on Mobile Learning, mLearn 2007, ed. A. Norman and J. Pearce, Melbourne, Australia, 2007, pp. 238–244.

(55) P. Pazmino and L. Lyons, Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '11, ACM, New York, NY, USA, 2011, pp. 895–904.

(56) P. Mantyjarvi, F. Paternò, Z. Salvador, and C. Santoro, Proceedings of the 8th conference on Human-computer interaction with mobile devices and services, ACM, New York, NY, USA, 2006, pp. 191–194.

(57) R. Grinter, P. Aoki, M. Szymanski, J. Thornton, A. Woodruff, and A. Hurst, Proceedings of the 2002 ACM conference on Computer Supported Cooperative Work, ACM, 2002, pp. 146–155.

(58) S. Hsi and H. Fait, "RFID enhances visitors' museum experience at the Exploratorium", *Communications of the ACM*, 2005, 60–65.

(59) E. Hornecker and M. Stifter, Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments, ACM, 2006, pp. 135–142.

(60) K. Winesmith, Ph.D. Thesis, Unversity of New South Wales, 2008.

(61) N. Flor, "Side-by-side collaboration: a case study", *International Journal of Human-Computer Studies*, Sept. 1998, **49**, 201–222.

(62) N. Miyake, "Constructive interaction and the iterative process of understanding", *Cognitive Science*, June 1986, **10**, 151–177.

(63) L. L. Downey, "Group Usability Testing: Evolution in Usability Techniques".

(64) P. Dourish and V. Bellotti, Proceedings of the 1992 ACM conference on Computer-supported cooperative work - CSCW '92, ACM Press, New York, New York, USA, Dec. 1992, pp. 107–114.

(65) J. Kjeldskov and J. Stage, "New techniques for usability evaluation of mobile systems", *International Journal of Human-Computer Studies*, May 2004, **60**, 599–620.

(66) M. Lew, L. Guzman, and A. Demler, "Politics, feelings and social media", 2013.

(67) M. Skov and P. Ingwersen, Proceedings of the second international symposium on Information interaction in context - IIiX '08, ACM Press, New York, New York, USA, 2008, p. 110.

(68) C. Fulton, "The pleasure principle: the power of positive affect in information seeking", en, *Aslib Proceedings*, May 2009, **61**, 245–261.

(69) R. Huizinga, *Homo Ludens: A Study of the Play Element in Culture*, Beacon Press, 1971.

(70) P. Wright and J. McCarthy, *Technology as Experience*, MIT Press, 2004.

(71) T. Barker, Proceedings of the 15th Biennial Conference of the Film and History Association of Australia and New Zealand, Sydney, Australia, 2010.

(72) W. W. Gaver, J. Bowers, A. Boucher, H. Gellerson, S. Pennington, A. Schmidt, A. Steed, N. Villars, and B. Walker, Extended abstracts of the 2004 conference on Human factors and computing systems - CHI '04, ACM Press, New York, New York, USA, Apr. 2004, p. 885.

(73) W. Gaver, "Designing for Homo Ludens", *I3 Magazine*, 2002, 3–6.

(74) L. McCay-Peet and E. Toms, iConference 2013 Proceedings, 2013, pp. 688–691.

(75) L. Björneborn, "Serendipity dimensions and users ' information behaviour in the physical library interface", *Information Research*, 2008, **13**, 1–13.

(76) A. Thudt, U. Hinrichs, and S. Carpendale, Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems, 2012, pp. 1461–1470.

(77) Y. Qu and G. W. Furnas, "Model-driven formative evaluation of exploratory search: A study under a sensemaking framework", *Information Processing & Management*, Mar. 2008, **44**, 534–555.

(78) V. L. O'Day and R. Jeffries, Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '93, ACM Press, New York, New York, USA, May 1993, pp. 438–445.

(79) G. Marchionini, "Exploratory search", *Communications of the ACM*, Apr. 2006, **49**, 41.

(80) K. Andrews, M. Pichler, and P. Wolf, Proceedings of the 2nd IEEE Symposium on Information Visualization, InfoVis '96, 1996, pp. 62–63.

(81) J. L. Jensen, "Augmentation of Space: Four Dimensions of Spatial Experiences of Google Earth", *Space and Culture*, Jan. 2010, **13**, 121–133.

(82) E. Ruffaldi, C. Evangelista, V. Neri, M. Carrozzino, and M. Bergamasco, Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts, 2008, pp. 113–119.

(83) S. Greene, G. Marchionini, C. Plaisant, and B. Shneiderman, "Previews and overviews in digital libraries: Designing surrogates to support visual information seeking", *Journal of the American Society for Information Science*, 2000, **51**, 380–393.

(84) M. Dörk, S. Carpendale, and C. Williamson, "Fluid Views: A Zoomable Search Environment", 2012, 233–240.

(85) S. Hinton and M. Whitelaw, Eletronic Visualisation and the Arts (EVA 2010), London, UK, 2010.

(86) D. Fallman, "The Interaction Design Research Triangle of Design Practice, Design Studies, and Design Exploration", en, *Design Issues*, July 2008, **24**, 4–18.

(87) C. Ahlberg and B. Shneiderman, Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94, ACM Press, New York, New York, USA, Apr. 1994, pp. 313–317.

(88) A. Foster, "A nonlinear model of information seeking behavior", *Journal of the American Society for Information Science and Technology*, 2004, **55**, 228–237.

(89) B. Shneiderman, Proceedings 1996 IEEE Symposium on Visual Languages, IEEE Comput. Soc. Press, 1996, pp. 336–343.

(90) M. Hall, E. Agirre, N. Aletras, R. Bergheim, K. Chandrinos, P. Clough, S. Fernando, K. Fernie, P. Goodale, J. Griffiths, O. L. de Lacalle, A. de Polo, A. Soroa, and Stevenson, in *Theory and Practice of Digital Libraries*, ed. P. Zaphiris, G. Buchanan, E. Rasmussen, and F. Loizides, Springer, 2012, vol. 7489, pp. 500–503.

(91) L. McCay-Peet and E. Toms, "Measuring the dimensions of serendipity in digital environments", *Information Research*, 2011, **16**, 1–7.

(92) R. Wille, in *Formal Concept Analysis*, ed. S. Ferre and S. Rudolph, Springer Berlin / Heidelberg, 2009, vol. 5548, pp. 314–339.

(93) C. Carpineto and G. Romano, *Concept data analysis: Theory and applications*, J. Wiley, 2004.

(94) M. Kaytoue, S. Kuznetsov, A. Napoli, and S. Duplessis, "Mining gene expression data with pattern structures in formal concept analysis", *Information Sciences*, 2011, **101**, 1989–2001.

(95) U. Priss, "Formal concept analysis in information science", *Annual review of information science and technology*, 2006, **40**, 1–22.

(96) R. Wille, "Conceptual landscapes of knowledge: A pragmatic paradigm of knowledge processing", *Classification in the Information Age*, 1999, ed. W. Gaul and H. Locarek-Junge, 344–356.

(97) R. Wille and B. Ganter, *Formal Concept Analysis: Mathematical Foundations*, Springer-Verlag, Berlin, 1999.

(98) J. Ducrou, Ph.D. Thesis, University of Wollongong, 2007.

(99) J. Ducrou, in *Conceptual Structures: Knowledge Architectures for Smart Applications*, Springer-Verlag, 2007, pp. 496–500.

(100) J. Ducrou, B. Vormbrock, and P. Eklund, Proceedings of 14th International Conference on Conceptual Structures, Springer, 2006, pp. 203–214.

(101) B. Ganter and R. Wille, in *Applications of Combinatorics and Graph Theory to the Biological and Social Sciences*, Springer-Verlag, 1989, pp. 139–167.

(102) A. Formica, "Concept similarity in Formal Concept Analysis: An information content approach", *Knowledge-Based Systems*, Feb. 2008, **21**, 80–87.

(103) A. Formica, "Ontology-based concept similarity in Formal Concept Analysis", *Information Sciences*, Sept. 2006, **176**, 2624–2641.

(104) J. Saquer and J. Deogun, "Concept aproximations based on rough sets and similarity measures", *International Journal of Applied Mathematics and Computer Science*, 2001, **11**, 655–674.

(105) E. Galbrun and P. Miettinen, "From black and white to full color: extending redescription mining outside the Boolean world", *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 2012, **5**, 284–303.

(106) N. Ramakrishnan, D. Kumar, B. Mishra, M. Potts, and R. Helm, Proceedings of the tenth ACM SIGKDD international conference on Knowledge Discovery and Data mining, ACM, 2004, pp. 266–275.

(107) E. Horvitz, "Uncertainty, action, and interaction: In pursuit of mixed-inititive computing", *Intelligent Systems: Trends and Controversies*, 1999, **14**, 17–20.

(108) D. Roth and D. Zelenko, Proceedings of the 17th International Conference on Computational Linguistics, 1998, pp. 1136–1142.

(109) J. R. Finkel, T. Grenager, and C. Manning, Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005), 2005, pp. 363–370.

(110) E. Panofsky, *Studies in Iconology: Humanistic Themes in the Art of the Renaissance*, Harper & Row, New York, 1972.

(111) C. Carpineto and G. Romano, "A lattice conceptual clustering system and its application to browsing retrieval", *Machine Learning*, 1996, **24**, 1–28.

(112) R. Godin, R. Missaoui, and H. Alaoui, "Incremental concept formation algorithms based on Galois (concept) lattices", *Computational Intelligence*, 1995, **11**, 246–247.

(113) P. Krajca, J. Outrata, and V. Vychodil, "Parallel algorithm for computing fixpoints of Galois connections", *Annals of Mathematics and Artificial Intelligence*, 2010, **59**, 257–272.

(114) J. Outrata and V. Vychodil, "Fast Algorithm for Computing Fixpoints of Galois Connections Induced by Object-Attribute Relational Data", *Information Sciences*, 2012, **185**, 114–127.

(115) S. Andrews, in *Conceptual Structures for Discovering Knowledge*, ed. S. Andrews, S. Polovina, R. Hill, and B. Akhgar, Springer Berlin Heidelberg, 2011, vol. 6828, pp. 50–62.

(116) M. Kirchberg, E. Leonardi, Y. S. Tan, S. Link, R. K. L. Ko, and B. S. Lee, in *Formal Concept Analysis*, 2012, vol. 7278, pp. 164–179.

(117) J. Outrata, Proceedings of the 10th International Conference on Concept Lattices and their Applications, ed. M. Ojeda-Aciego and J. Outrata, La Rochelle, France, 2013, pp. 261–274.

(118) S. Kuznetsov, "A fast algorithm for computing all intersections of objects in a finite semi-lattice", *Automatic Documentation and Mathematical Linguistics*, 1993, **27**, 11–21.

(119) S. Kuznetsov, PKDD 1999, 1999, pp. 384–391.

(120) J. Outrata, "Incremental Close-by-One", 2013.

(121) D. van der Merwe, S. Obiedkov, and D. Kourie, in *Proceedings of the International Conference on Formal Concept Analysis*, Springer Berlin Heidelberg, 2004, vol. 2961, pp. 205–206.

(122) E. M. Norris, "An Algorithm for Computing the Maximal Rectangles in a Binary Relation", *Revue Roumaine de Mathématiques Pures et Appliquées*, 1978, **23**, 243–250.

(123) F. Krishnabhakdi-Vasilakis, A. Lawson, T. Wray, P. Goodall, and P. Eklund, Proceedings of the International Conference on Education and New Learning Technologies, Barcelona, Spain, 2010.

(124) J. Trant, "Studying Social Tagging and Folksonomy: A Review and Framework", *Journal of Digital Information*, 2009, **10**.

(125) P. Brawn, *Usability Test Report, The Virtual Museum of the Pacific*, Technical Report, Usability, Eyetracker, 2011.

(126) D. Norman, *The Design of Everyday Things*, Currency Doubleday, New York, 1988.

(127) T. Winograd, *Bringing Design to Software*, Addison-Wesley, Reading, MA, 1996.

(128) L. Williams, R. Kessler, W. Cunningham, and R. Jeffries, “Strengthening the case for pair programming”, *IEEE Software*, 2000, **17**, 19–25.

(129) E. Tufte, *Visual Explanations*, Graphics Press, Cheshire, CT, 1997.

(130) S. Chapple and R. Garofalo, *Rock 'N' Roll is Here to Pay: The History and Politics of the Music Industry*, Chicago, Illinois, 1977.

(131) C. Lagoze, H. van de Sompel, M. Nelson, and S. Warner, *The Open Archives Initiative Protocol for Metadata Harvesting, Version 2.0*, 2014.

(132) B. H. Millstein and S. M. Lowe, *Consuelo Kanaga: An American Photographer*, University of Washington Press, 1992.

(133) K. Lerner, *Nathan Lerner - Biography*, 2014.

(134) D. Weinberger, “Taxonomies to Tags: From Trees to Piles of Leaves”, *Release 1.0*, 2005, **23**.

(135) J. Löwgren, in *Encyclopedia of Human-Computer Interaction*, ed. M. Soegaard and R. Dam, The Interaction Design Foundation, Aarhus, Denmark, 2008.

(136) L. Candy, *Practice Based Research : A Guide*, tech. rep., University of Technology, Sydney, 2006.

(137) N. Cross, “Designerly Ways of Knowing: Design Discipline Versus Design Science”, en, *Design Issues*, July 2001, **17**, 49–55.

(138) J. Klavans, R. Stein, S. Chun, and R. D. Guerra, Museums and the Web 2011: Proceedings, ed. J. Trant and D. Bearman, 2011.

(139) S. Golder and B. Huberman, “The Structure of Collaborative Tagging Systems”, *Journal of Information Science*, 2006, **32**, 198–208.

(140) R. Wille, in *Formal Concept Analysis*, ed. R. Missaoui and J. Schmidt, Springer Berlin Heidelberg, 2006, vol. 3874, pp. 1–29.

(141) E. Stolterman, H. Jung, W. Ryan, and M. Siegal, “Device Landscapes: A New Challenge to Interaction Design and HCI Research”, *Journal of Korean Society of Design Research*, 2013, **26**, 2–29.

(142) T. Wray, P. Eklund, and K. Kautz, ICIS 2013 Proceedings, Milan, Italy, 2013.

(143) P. Eklund, T. Wray, P. Goodall, and A. Lawson, “Design, information organisation and the evaluation of the Virtual Museum of the Pacific digital ecosystem”, *Journal of Ambient Intelligence and Humanized Computing*, 2012, **3**, 265–280.

(144) T. Wray and P. Eklund, Proceedings of the 1st International Workshop on Semantic Digital Archives, ed. L. Predoiu, S. Hennicke, A. Nurnberger, A. Mitschick, and S. Ross, 2011, pp. 109–120.

(145) P. Eklund, T. Wray, and J. Ducrou, "Linking Objects and their Stories: An API For Exploring Cultural Heritage Using Formal Concept Analysis", *Journal of Emerging Technologies in Web Intelligence*, 2011, **3**, 239–252.

(146) T. Wray and P. Eklund, Proceedings of the 9th International Conference on Formal Concept Analysis, ed. P. Valtcheck and R. Jaschke, Springer-Verlag, Nicosia, Cyprus, May 2011, vol. 6628, pp. 251–266.

(147) P. Eklund, P. Goodall, and T. Wray, National Digital Forum, Wellington, New Zealand, Sept. 2010.

(148) P. Eklund, A. Lawson, P. Goodall, and T. Wray, Museums and the Web 2010: Proceedings, ed. J. Trant and D. Bearman, Archives & Museum Informatics, Mar. 2010.

(149) P. Eklund, T. Wray, and J. Ducrou, Proceedings of the International Conference on Management of Emergent Digital EcoSystems, ACM, New York, NY, USA, 2009, pp. 36–245.

(150) T. Wray and P. Eklund, Proceedings of the 7th International Conference on Concept Lattices and Their Applications, ed. M. Kryszkiewicz and S. Obiedkov, Sevilla, Spain, Oct. 2010, pp. 139–150.

(151) P. Eklund, A. Lawson, T. Wray, P. Goodall, and M. V. Olffen, Pacific Art in the 21st century: Museums, New Global Communities and Future Trends, PAA 2010 International Symposium, Cook Islands, Aug. 2010.

Appendix A

A Place for Art: Pre-Evaluation Survey

The following are the results of the survey administered to the participants who undertook the *A Place for Art* user experience evaluation. The intention of the survey was not to form conclusive results, but rather, profile the participants for the subsequent user experience evaluation.

Some participants did not complete the survey or did not answer all of the questions. In such cases, there answer has been marked as N/A. A total of 24 individual participants took the survey.

How old are you?	*N*
18 – 25	13
25 – 34	8
35 – 44	3
45 – 55	0
over 55	0

How often do you visit museums and galleries?	*N*
Less than once a year	6
Between less than once a year and 3 times a year	5
Between three times a year and once a month	7
More than once a month	6

What is your current occupation, background or university major?	*N*
Graphic design	6
Visual arts	5
Psychology	2
Curatorial practice	2
Digital media	2
Computer science	1
Law	1
Industrial design	1
Librarian	1
Bio-medical science	1
Engineering	1
Content producer	1

How would you describe your relationship with the University of Wollongong?	*N*
I am not associated with the university	8
I am a current student at the university	13
I am a former student at the university	2
I am a current or former staff member at the university	1
I work for or am associated with a community or student organisation at the university	0
I am associated in some other way with the university	0

If you work or study at the University of Wollongong, how would you best describe your awareness with the public art that is displayed on campus? You may select more than one option.	*N*
I do not, or have not, noticed any paintings, sculptures or other works of art on campus	0
At least once I have noticed a work of art, such as a paintings or sculpture that is on display on campus, but I have not paid much attention to it	4
At least once I have noticed a work of art, such as a paintings or sculpture that is on display on campus, and I have read its label	10
At least once I have noticed a work of art, such as a painting or sculpture that is on display on campus, and I have felt curious by it	10
At least once I have noticed a work of art, such as a paintings or sculpture that is on display on campus, and I wanted to know more about the work	12
At least once I have noticed a work of art, such as a paintings or sculpture that is on display on campus, and I wanted to know more about the artist who made the work	11
Not applicable (I do not study or work at the University of Wollongong)	8

The University of Wollongong has released a book called A Place for Art that features selected works from the University's Art Collection. Prior to this survey, have you heard anything about the book?	*N*
I've never heard anything about A Place for Art	18
I know about the book through friends and family	3
I know about the book because I've seen it on the University of Wollongong's website	2
I know about the book because I've heard about it elsewhere within the media	1
I know about the book because I went to a launch that promoted the book	0
I own a copy of A Place for Art	0

What devices or technologies have you used to access online collections? You may select more than one option.	*N*
A Web browser from a desktop or laptop computer	16
A Web browser from a smartphone or tablet device	7
A tablet app downloaded from the app store or similar	1
A smartphone app downloaded from the app store or similar	5
I have never accessed an online collection	7

How have you previously accessed, or found out about an online collection? You may select more than one option.	*N*
Via a Google search (or similar) on the Web	13
As promoted and linked on the website of the museum, library or archive	10
As promoted and linked from the other marketing materials within the museum, such as signage and brochures	6
As linked via a Twitter, Facebook or social media	8
As linked via e-mail	2
As linked via a single record of the object, which led me to the collection of works	4
I have never accessed an online collection	7

If you were to explore or search for works from a museum's collection, what would be your preferred way of doing so? You may select more than one option.	*N*
Via a website as viewed from a desktop or laptop computer	20
Via a website as viewed a tablet device	13
Via a website as viewed on a smartphone device, such as (or similar to) an iPhone	10
Via an app downloaded from the app store (or similar) as viewed on a tablet device	14
Via an app downloaded from the app store (or similar) as viewed on a smartphone device	6

If you were to explore or search for works from a museum's collection, what would be your preferred way of doing so? You may select more than one option.	*N*
Via a website as viewed from a desktop or laptop computer	20
Via a website as viewed a tablet device	13
Via a website as viewed on a smartphone device, such as (or similar to) an iPhone	10
Via an app downloaded from the app store (or similar) as viewed on a tablet device	14
Via an app downloaded from the app store (or similar) as viewed on a smartphone device	6

Regardless of your preferred technology platform, what would best describe your reasons for accessing a museum's online collection? You may select more than one option.	*N*
To learn more about the works of art I am already familiar with	16
To learn more about artists that I am already familiar with	19
I can choose which objects or artists I can learn about	16
I would be curious just to try it out	11
As a way of accessing more information about the works I have seen within an exhibition	17
As a way of accessing additional material about the works, such as video content or audio interviews from the artist	17
As a way of exploring the collection	18
As a way of discovering new works	19
As a way of sharing favourite works with my friends on social networks, such as Facebook or Twitter	8

Appendix B

Previous Research Contributions

Although this monograph represents a single body of work, the research objectives, the methods used to address these objectives and the case studies that follow have been represented in a number of publications that have also been co-authored by the author of this thesis. For each work in the list of seminal publications that follows, a brief summary of the work is provided, along with how that work addresses the research objectives in Section 1.1.

- A conference paper [142] titled *Pathways through Information Landscapes: Alternative Design Criteria for Digital Art Collections* presented at the 2013 International Conference on Information Systems describes how *design science* and *concept-driven interaction design* can be used to shape a design artifact and build a design theory that describes pathway-based navigation within digital art collections. Research objectives 1 and 3 are addressed in its description of the pathway metaphor; its description of how Formal Concept Analysis can be used to create pathways of content within digital art collections; and the presentation of *A Place for Art* as a case study (Section 4.3).
- A journal paper [143] titled *Design, information organisation and the evaluation of the Virtual Museum of the Pacific digital ecosystem* published in the Journal of Ambient Intelligence and Humanized Computing also describes how Formal Concept Analysis can be used to create information spaces for navigation and exploration within a digital art collections. In addressing research objectives 1, 3 and 4, it describes the *Virtual Museum of the Pacific* case study (Section 4.2) as a practical demonstrator of a Formal Concept Analyis based system for browsing museum content while reporting on its impact to its organisational environment, and it also reports on the results of two user experience evaluations.

- A conference paper [144] titled *Concepts and Collections: A Case Study using Objects from the Brooklyn Museum* presented at the Proceedings of the 1st International Workshop on Semantic Digital Archives investigates some of the scalability concerns of using Formal Concept Analysis to organise and structure collection content – concerns that are investigated further and addressed in Section 4.1.3. It also describes work that forms the basis of the *CollectionWeb* framework in Section 4.1, and introduces elements that form the basis of the *Brooklyn Museum Canvas* case study (Section 4.4). Research objectives 1 and 2 are addressed through its discussion of applying Formal Concept Analysis to create connections among artworks in collections, and its discussion on how it could be applied to multiple collections with varying sizes in a scalable manner.

Other prior research contributions include works that further describe the theory and application of Formal Concept Analysis to museum collections [145–150], the *CollectionWeb* framework [148] and the work as presented from museological perspectives [123, 151]. Although this thesis is the result of the culmination or evolution of these published works, they do not entirely represent the scope of the work presented in this thesis.