



TOWARDS
A NATIONAL
COLLECTION



UKRI
Arts and
Humanities
Research Council

INTERIM REPORT

FOUNDATION PROJECTS

DEEP DISCOVERIES

PI: Lora Angelova, The National Archives, UK

The National Archives, UK | The University of Surrey |
Victoria & Albert Museum | Royal Botanic Garden Edinburgh |
Museum of Domestic Design and Architecture | Gainsborough Weaving Company |
Sanderson Design Archive

DECEMBER 2020

TABLE OF CONTENTS

Executive Summary.....	1
Abstract	2
Aims and Objectives.....	2
Partnership Structure.....	4
Staffing Structure	5
Covid-19 Impacts	5
Revised Programme	6
Events and Consultations	7
Research Approach	8
Early Research Results/Outputs	9
Next Steps	12
Contacts	12

Executive Summary

Project Summary: The Deep Discoveries foundation project is carrying out research into the application of computer vision (CV) search tools for enhancing the ability of general audiences and specialist researchers to discover visual collections in new and/or more effective ways.

Structure and Methodology: The project has four Core Participants: The National Archives (TNA), The University of Surrey (UoS), The Victoria and Albert Museum (V&A), The Royal Botanic Garden Edinburgh (RGE); and three Project Partners: The Gainsborough Weaving Company, The Museum of Domestic Design and Architecture, and Sanderson Design Archive. A user experience (UX) team has been assembled from TNA and V&A, which also includes a software developer. The project has two Work Packages. WP1, 'technology development', is led by UoS and involves the development of CV search algorithms necessary for delivering new and/or more effective ways for cross-collection searching of digitised visual content. WP2, 'user assessment and impact', is led by TNA, V&A, and RGE and involves assessment of current visual collection users and their habits and needs, in order to inform the design of a prototype search tool using the technology developed in WP1. The UX research team works iteratively with the UoS team in the research, development, and testing of the technologies and concepts.

Progress: The core participants and project partners have transferred nearly 20,000 images to UoS for use in the search algorithm development. UoS have developed CV algorithms (deep neural network models) to support the visual search task of the project. The model is trained to discriminate between the kinds of objects present in the contributed image, as well as to discriminate between the visual styles of those images. To aid in fine-tuning the model, UoS have also developed a platform for image annotation which may be leveraged to collect annotations for the next phase of the project.

The UX research team has conducted a literature review, and through a number of project meetings with the participants and project partners, has developed the content and structure of workshop 1 (of 2 planned) and a follow-up survey to be launched in January 2021. Results from this information gathering campaign have been described in a detailed report, which includes three 'problem statements' or user-centered scenarios in search and discovery, used to inform the front-end design of a prototype web-based search app that will employ the CV algorithms developed by UoS.

Covid-19 Impacts: None of the core participants have been furloughed; however there has been increased pressure on their time due to (1) many colleagues in their organisations being furloughed, and (2) an increasing demand on digital services staff across IROs which affected the UX research team. Several project partners were also furloughed, as were many of the visual collections experts we were seeking to connect with during the first phase of the research. Workshop 1 and user interviews were conducted virtually, limiting the number of participants in attendance and the kinds of interactions possible. We have requested a 3-month no-cost extension to the project to facilitate completion of the planned research.

Next steps: A survey on visual collection use, search, and discovery will be launched at the start of 2021. A further problem statement will be developed to inform the prototype development so that the search tool can deliver benefit both to general audiences and specialist researchers. Based on the proposed search and discovery scenarios, the UoS team will explore novel 'explainable AI' techniques to explain via visualizations (such as heat maps) which regions of each search result are considered to match the query image. This will form part of an iterative interface to help determine users' search intent and enable interactive refinement of search results under multiple competing definitions ('facets') of visual similarity. The technology will then be integrated into a prototype web-

based app developed by the UX research team, and tested with users during Workshop 2. A full report of the research and results from the project will be generated.

Abstract

Deep Discoveries aims to contribute to the creation of a unified national heritage collection by exploring the potential of computer vision (CV) search for content discovery within and between our nation's digitised image repositories. The research will lead to the design of a prototype platform enabling cross-collection image linking by harnessing the ability of CV and deep learning methods to identify and recognise specific patterns without the need for preliminary integrated descriptive metadata. Searching in this manner allows for content-linking based on attributes such as pattern, colour, and motif, and creates the opportunity for users to discover unforeseen connections between digital and digitised items in visual collections across the country. The research will employ a narrow image set featuring botanical motifs, to begin exploring the potential of CV technology to generate engaging, relevant search and discovery options for different types of users and for different types of tasks. To achieve this, research carried out by a user experience (UX) research team from two GLAM IROs will inform the work of computer vision scientists at the University of Surrey. The Deep Discoveries project is also working with three partner organisations representing different owners and creators of visual collections – a small university-based museum, a private design studio, and a commercial design archive – to open up participation in funded research to smaller organisations, to glean a better understanding of their needs, and to assess the opportunities and challenges involved in gathering visual collections for the purpose of employing CV-based search and discovery tools.

Aims and Objectives

The Deep Discoveries Foundation Project has three objectives which are interdependent and of equal priority and two objectives that are specific to the Towards a National Collection: Opening UK Heritage to the World SPF Programme:

1. Develop a software platform for matching content within and across collections, enabling visual records to be linked based on aesthetic properties such as pattern, style, and visual motifs.
2. Develop methodologies to survey the current users of digital image collections in the UK Heritage sector, evaluate who is excluded, what are the real and perceived barriers to user access and discovery of these collections, and find measurables that can be used to evaluate the success of the proposed technological advancement.
3. Convene workshops and surveys to explore current methods for engaging, evaluating, and diversifying audience participation; assess the ethical issues of deploying AI to image searching in digitised collections; gather data on image-resource discoverability, and define requirements for a visual search approach; and to showcase proof-of-concept technology.
4. Through the establishment of a national network around this theme, encourage and enable the integration of computer vision search technologies for visual collections in Discovery phase projects.
5. Produce a report on the State of the Art in Computer Vision Searching for Heritage Collections, which will be used to provide clear evidence-based policy recommendations on next steps in the field.

The project is structured around two work packages (WPs). WP1 involves technology development and deployment while WP2 revolves around current user assessment, development of methods to engage new and more diverse audiences, ethical questions, and technology uptake and impact assessment. The work plan and milestones are shown below; deliverables and associated roles and responsibilities are outlined in the following section.

Figure 1: Gantt chart illustrating project timeline; ‘break’ in dating illustrated with red line demonstrates COVID-19 delay. Timeline has been updated with the new project end date. Deliverables named in Table 2.

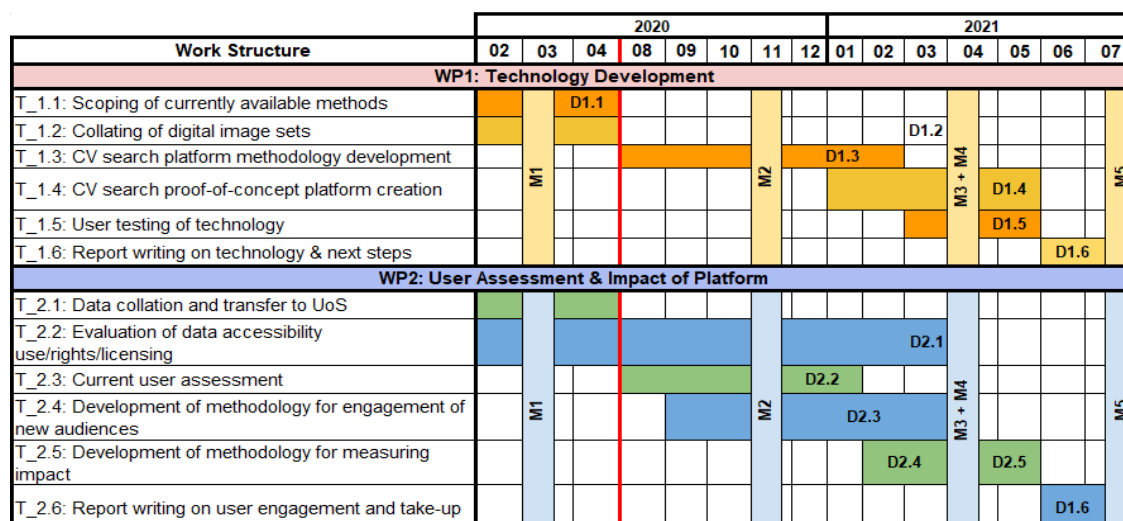


Table 1: Project milestones

MILESTONES

- M1** Project kick-off meeting (internal)
- M2** Workshop 1 + Project meeting 2
- M3** Technology developed to user-testing stages
- M4** Closing workshop + Project Meeting 3
- M5** Final report

Table 2: Responsibilities of core project participants

* TNA - The National Archives, UoS - University of Surrey, V&A - Victoria and Albert Museum, RBGE - Royal Botanic Garden Edinburgh

DELIVERABLE		LEAD*
WORK PACKAGE ONE		
D1.1	Evaluation of currently available technologies and decision-making around selection	UoS
D1.2	Report on ease of data access and state of format and availability	RBGE
D1.3	Proof-of-concept CV search platform creation	UoS
D1.4	Proof-of-concept platform user evaluation	TNA/V&A
D1.5	Methodology to evaluate technology impact on discovery	TNA/V&A
D1.6	Final Project report	TNA; All (input)
WORK PACKAGE TWO		
D2.1	Report on access to digitised visual collections, format, and IP/use to feed into final project report	RBGE
D2.2	Methodology to evaluate current user access	TNA/V&A
D2.3	Methodology to reach more diverse audiences, addressing ethics questions	TNA/V&A
D2.4	Methodology to evaluate technology impact on discovery	TNA/V&A
D2.5	Report on impact of user access and diversity	TNA/V&A
D2.6	Final project report	TNA; All (input)

Partnership Structure

Core Participants:

The National Archives, UK (TNA): The Principle Investigator (PI), Co-Investigator 4 (CI4), Project Administrator (PA), UX Research Staff 2 (RS2) and Software/User Interface Developer (RS3) are based at TNA's Collection Care, Research, and Digital Services departments. The team leads on WP2 – user-centric research, as well as coordinating the project, and ensuring delivery. The organisation has supplied 7845 images from the Board of Trade Registered Designs collections. The collection has jpeg images with unique image id; other forms of metadata are not available. The images are primarily of textile design patterns, though not all, featuring a botanical theme.

University of Surrey (UoS): (CI1) and Post-Doctoral Research Associate (RS1) are based at UoS Centre for Vision, Speech and Signal Processing (CVSSP); the team leads on WP1 – technology development – by ingesting and keeping secure images provided by all partners, exploring existing and developing new CV algorithms for the project, and working closely with the UX team to develop the front-facing user interface.

Victoria and Albert Museum (V&A): CI3 and UX RS4 are based at V&A's research institute (VARI) and Digital Media department; the team is working on WP2 with RS2 completing the UX research team. The organisation has supplied a total of 426 images from their collections. The images primarily are of floral patterns and designs on paper, though not all, featuring a botanical theme. Other forms of metadata are not available for this collection.

Royal Botanic Garden, Edinburgh (RBGE): CI2 is based at the Plant Records department and manages data for the living and herbarium collections. He is working within WP2 on questions around visual collection data availability. The organisation has supplied 7435 images from various collections in a spreadsheet which has URL links to the images to download and associated metadata such as scientific names, photographer, collector name and number, and catalogue number and license links. The images are primarily botanical themed including leaf and flower sketches, photographs of leaves, mountains, bushes, flower/seed/fruit of plants.

Project Partners: Partners on the project attend Project Meetings described in the work plan and milestones table, participate in workshops and surveys and contribute valuable information around the use and state of their digitised visual collections.

Museum of Domestic Design and Architecture (MoDA): Part of the Middlesex University, MoDA's collections feature designs for wallpapers and textiles, many of which have a floral or botanical theme. Two staff members are contributing time in-kind worth £3000 to the project, as well as 1170 digitised images and the associated metadata. The annotations contain i) short description of image content (e.g. Design for a textile of red, blue and yellow flowers); ii) materials/technique used (e.g. Watercolour on detail paper) and iii) production date.

Gainsborough Weaving Company (GWS): The archive of the weaving studio features thousands of images of historic textile designs, as well as contemporary designs added as they are produced; however, few have been digitised thus far. One staff member is offering time in-kind worth £600, as well as 760 images to the project.

Sanderson Design Archive (SDA): The archive's collections feature hand-painted, flower-group designs and floral-patterned textiles and wallpapers. One staff member is offering time in-kind worth £1645 as well as approximately 950 images and the associated metadata which include i) image-type (Wallpaper/Textile/Handprinted-Wallpaper); and Collection name (e.g. Triad 1968-69 3rd Edition, riad 1970-71 Collection)

Staffing Structure

PI: Lora Angelova, Head of Conservation Research, TNA: project direction, leads on reports, liaising with Programme Director and foundation projects; ensures delivery of project to proposed timeline and budget

CI1: John Collomosse, Professor of Computer Vision, UoS CVSSP: direction and lead on visual search development; supervise PDRA RS1, contribute to final report

CI2: Robert Cubey, Deputy Herbarium Curator, RBGE: leads on RBGE data access; provide collection-specific advice; report on access to digitised graphic collections, format, and IP/use to feed into final project report

CI3: Joanna Norman, Director of the V&A Research Institute, V&A: leads on V&A data access; provide collection-specific advice; contribute to final report

CI4: Pip Willcox, Head of Research, TNA: facilitates links with partnership networks; contribute to final report

RS1, PDRA: Dipu Manandhar, Research Fellow, UoS CVSSP: leads on development of visual search technology, works closely with UX research team

RS2: Hari Chadrapal, UX Researcher, TNA: part of UX research team, leads on UX methodologies and user interviews, insight from TNA UX research expertise

RS3: Bernard Ogden, Research Software Engineer, TNA: part of UX research team, leads on UX methodologies and frontend software development

RS3: Jack Craig, Design Lead, V&A: part of UX research team, leads on UX methodologies, advising on search and discovery for V&A collection website and the challenges in enabling discovery with a large heterogeneous dataset

Project Administrator, Rachel Smillie, Head of Academic Partnerships, TNA: works closely with PI to aid in project delivery – scheduling meetings, sending out information and reminders, organising workshops

Contractor, initially envisioned as UX consultant, this role may be reassigned in the second half of the project. TBD at an upcoming project meeting and in consultation with Programme Director and Senior Researcher.

Covid-19 Impacts

COVID-19-related delays to the project primarily revolve around staff availability and affect WP2, and subsequently, progress on the overall project. None of the core project members have been furloughed; however, many staff members within the V&A and RBGE are on furlough, putting a large amount of stress and workload on our partners there. Furthermore, all three IROs – TNA, V&A, and RBGE – have seen a significant shift in demand to Digital services-related staff as noted by others across the sector.¹ From a sudden surge in digitally available material to new ways of interacting with audiences entirely online (e.g. virtual exhibitions and events), our UX research team has been stretched in unforeseeable ways. In order to advance with the design and training of the computer vision search algorithms, the UoS team requires input from the UX research team around user habits, practices, and tasks. To attain these, workshops and surveys need to be carried out; however, these have been delayed due to (1) staff availability to devote time for research on the project, and

¹ Network of European Museum Organisations 'Survey on the impact of the COVID-19 situation on museums in Europe' Final Report, December 2020. https://www.nemo.org/fileadmin/Dateien/public/NEMO_documents/NEMO_COVID19_Report_12.05.2020.pdf Acc. 11/12/2020

(2) furlough of many users and collection experts whom we would like to consult. Demands of a similar nature have also slowed progress on D2.1. In addition, some staff at our Project Partner organisations are furloughed, in one case, before images were transferred to UoS. We have asked for a 3-month no-cost extension, which will allow for the readjustment of our team to the new demands of post-COVID-19 working. The changes to the work programme and deliverable timings are reflected in the work plan Gantt chart and deliverables table.

COVID-19 disruption has also meant that our initial intent to co-host our first workshop with other TaNC Foundation Projects facing similar research questions was not successful. It goes without saying that the first workshop was hosted virtually, and for this reason was restricted to a much smaller number of attendees than would have been possible during an in-person event. This may be counter-intuitive, as virtual events tend to draw in larger audiences; however, the nature of the workshop, which necessitated discussion between the participants and breakout sessions, required careful curation in the digital space. We may instead opt to repeat the first workshop after the results have been analysed, or proceed with a follow-up survey to gather further information. A public event on ‘computer vision for heritage applications’ has been discussed with the University of Oxford’s Visual Geometry Group and the Alan Turing Institute’s new working group on CV for heritage. On a positive note, both the V&A and TNA were in the process of re-evaluating and redesigning their online collection sites; the research carried out by the UX leads in these organisations has been informed by these initiatives (which undoubtedly have shifted since the onset of the epidemic), dovetailing with the project’s aims to understand user’s engagement with visual collections online.

Revised Programme

Work Structure	2020												2021						
	02	03	04	08	09	10	11	12	01	02	03	04	05	06	07				
WP1: Technology Development																			
T_1.1: Scoping of currently available methods			D1.1																
T_1.2: Collating of digital image sets											D1.2								
T_1.3: CV search platform methodology development									D1.3										
T_1.4: CV search proof-of-concept platform creation		M1					M2				M3 + M4	D1.4			M5				
T_1.5: User testing of technology												D1.5							
T_1.6: Report writing on technology & next steps													D1.6						
WP2: User Assessment & Impact of Platform																			
T_2.1: Data collation and transfer to UoS																			
T_2.2: Evaluation of data accessibility use/rights/licensing											D2.1								
T_2.3: Current user assessment									D2.2										
T_2.4: Development of methodology for engagement of new audiences		M1					M2				D2.3				M5				
T_2.5: Development of methodology for measuring impact										D2.4		D2.5							
T_2.6: Report writing on user engagement and take-up														D1.6					

	DELIVERABLE	LEAD	START	END
WORK PACKAGE ONE				
D1.1	Evaluation of currently available technologies and decision-making around selection	UoS	02/2020	04/2020
D1.2	Report on ease of data access and state of format and availability	RBGE	02/2020	04/2021
D1.3	Proof-of-concept CV search platform creation	UoS	05/2020	03/2021
D1.4	Proof-of-concept platform user evaluation	TNA/V&A	01/2021	06/2021
D1.5	Methodology to evaluate technology impact on discovery	TNA/V&A	03/2021	06/2021
D1.6	Final Project report	TNA; All (input)	06/2021	07/2021
WORK PACKAGE TWO				
D2.1	Report on access to digitised visual collections, format, and IP/use to feed into final project report	RBGE	02/2020	04/2021
D2.2	Methodology to evaluate current user access	TNA/V&A	05/2020	02/2021
D2.3	Methodology to reach more diverse audiences, addressing ethics questions	TNA/V&A	09/2020	04/2021
D2.4	Methodology to evaluate technology impact on discovery	TNA/V&A	02/2021	06/2021
D2.5	Report on impact of user access and diversity	TNA/V&A	09/2020	06/2021
D2.6	Final project report	TNA; All (input)	06/2021	07/2021
MILESTONES				
M1	Project kick-off meeting (internal)	TNA	03/2020	
M2	Workshop 1 + Project meeting 2	TNA/V&A	09/2020	10/2021
M3	Technology developed to user-testing stages	All	04/2021	05/2021
M4	Closing workshop + Project Meeting 3	UoS/RBGE		05/2021
M5	Final report	TNA/All	06/2021	07/2021

Events and Consultations

	EVENT/CONSULTATION	DATE(S)	NO OF ATTENDEES/RESPONDENTS
M2	WORKSHOP 1 (virtual): 'Understanding current and potential users' goals and behaviours when working with Visual Collections online'	28.09.2020	11 + 7 DD facilitators
M2.1	Follow-up virtual workshop	TBD	
M2.2	Follow-up survey based on Workshop 1 input	01/02.2021	
M2.3	Interviews with Visual Collection Specialists	Ongoing	5
M2.4	Engagement event - CV in Heritage Collections webinar	03/04.2021	
D1.2/D2.1	Survey on access to digitised visual collections, format, IP; use	TBD 02.2021	
M4	WORKSHOP 2	TBD 04/05.2021	

Talks/Posters/Interviews

13.05.2020: Blog post 'Deep Discoveries: A new way of exploring and connecting digitised image collections' by Lora Angelova and Liz Fulton published on The National Archives website

09.2020: Interview of Bernard Ogden by Aberystwyth University Archive Administration MA student Chiara Fallone about different ways archival institutions are using emerging technologies

04.11.2020: Poster presentation about the project at The National Archives' Annual Digital Lecture by Carly Kind, 'The death of anonymity in the age of identity'. A follow-up Twitter chat around the posters was held on 09.11.2020

25.11.2020: Interview of Lora Angelova with input from John Collomosse and Dipu Manandhar by Giles Bergel for Europeana's AI in GLAMs interim report

28.01.2021: Invited talk and workshop participation of Bernard Ogden and Lora Angelova at the Archives in the UK/Republic of Ireland and AI (AURA) Network's second workshop.

Research Approach

The Collections Trust report '[Mapping digitised collections in England](#)'² tested the potential of image-recognition and analysis tools (Google Cloud Vision, Amazon Rekognition) for enhancing the performance of a prototype tool for cross-collection searching. The report also described Art UK's collaboration with the Visual Geometry Group at the University of Oxford in complementing the efforts of human taggers in creating metadata for search and discovery. The report notes that any framework for collection aggregation should 'demonstrate the potential' for image-recognition tools to be integrated.

The image recognition technology explored in Deep Discoveries functions in an entirely different way to the platforms explored in the Collections Trust report, and thus offers additional benefit to the Towards a National Collection programme aims. The above-mentioned tools function by carrying out object classification tasks – in effect, similar to auto-tagging images with labels (e.g. lorry, cat, dog etc.) and tend to operate at a semantic level i.e. tagging what is in an image. This can be useful for generating metadata about digital collections in an aggregator, and for enhancing keyword-based searching. However, this kind of search can be relatively limited and constrained by the (1) types of terms the system generates, (2) the users' intent, or task, in making use of the search tool, (3) the users knowledge of 'right term' to search. The algorithms developed by the UoS CVSSP team differ in two ways. First, they do not operate at a semantic level but instead analyse the visual style/aesthetic of the image rather than its content. Second, they do not classify/tag, but rather, use an image as input and return similar images based on their style. This can be thought of as a ranking operation ('return a list of images ranked by relevance') rather than a labelling operation ('tag all images containing x'). The search and discovery interface for this kind of algorithm consequently would not rely on keyword searching, but rather on image drag-and-drop, or selecting images that lead to others (e.g. Pinterest, Behance.Net), or even sketching by the user.

WP1 will develop a software platform for matching content within and across collections, enabling visual records to be linked based on aesthetic properties such as pattern, style, and visual motifs. The platform will require a scalable solution to cross-domain matching of visual content that is uniquely possible within the multi-disciplinary context of Deep Discoveries due to the diversity of content and end-user requirements of the collections. Consider a visual design identified within a single collection; similar patterns or motifs may exist within other collections – some may be sketches, some paintings, some photographs – it would be potentially useful to: a) recognize the relevance of this content, invariant to the depictive media or style; b) identify a diverse snapshot of

² Mapping Digitised Collections in England, Final report. Collections Trust, 2019. (<https://www.gov.uk/government/publications/mapping-digitised-collections-in-england>) acc. 11/12/2020

these variants to discover developmental directions of that motif e.g. for art historians or designers seeking to leverage the collections for inspiration. Delivering (a) requires new algorithms for cross-domain matching (T1.3, revised Months 3-12) of visual content at both a lower level e.g. structural similarity, and a higher level e.g. aesthetic look and feel of the piece holistically. Delivering (b) requires the engineering of an interactive search platform capitalising on (a) and delivering a user interface (T1.3, revised Months 12-15) suitable for exploring the diversity of the collections across different facets of variation for the match within the stock image data. Preceding (a) and (b) was a shorter piece of engineering work enabling ingest of visual content across participating collections (T1.2, Months 1-3).

WP2: Based on experience from our own institutions as well as the Project Partner organisations (which include private companies), we will modify the questions proposed in the Collections Trust report to create a survey to assess the state of digitised collections and access thereof in the context of creating a vision search platform (D1.2/2.2, revised Months 9-14). Through insight from Workshop 1 and a follow-up survey, we will explore the current users of the selected collections and outline several search scenarios representative of the kinds of new audiences and engagement practices we are aiming to achieve. Based on these users and target behaviours, we will create a set of questions that can be iteratively used to test the proof-of-concept technology during its development (T2.3 & T2.4, revised Months 8-14). After a prototype platform is created, it will be publicised and tested by users in order to assess the new tool's ability to generate new and insightful search and discovery results and engage new audiences or existing audiences in new ways (D2.5, M4, revised Months 14-16).

Early Research Results/Outputs

Progress

The first tasks of the UX research team in the Deep Discoveries project was to gain an overview, from collections experts, holders, or curators, on the kind of search and discovery their current users are seeking, the types of tasks they can and cannot (but would like to) perform, and the kinds of new audiences they believe a computer vision search tool can engage with their collections. For example, some researchers and operators of commercial design studios are interested in finding exact matches of images and patterns in order to trace origin or ensure that their intellectual property is protected, respectively. Other researchers, members of the public, creatives, as well as customers of commercial design studios may be interested in non-identical matches that create an opportunity for discovery through unexpected, but visually similar search returns. These different agendas may mean that the prototype under development by the project should be able to provide different kinds of search options, as well as function in a way that guides non-expert audiences or researchers to understand how these different types of searching operate. Questions around user intent and training of the search algorithm also require the team to investigate the meaning of terms such as 'motif' or 'style' for different user groups.

After the initial collection of images from the core participants and partners and the recruitment of the UX researchers, a literature review was carried out by the team and the collated resources were circulated among the group for discussion. Workshop 1 was organised as a pilot in surveying visual collection experts from our project partners and several other design archives. Simultaneously, a number of interviews were conducted with people working in various fields of design to glean more detailed information. Based on insight from these events, three problem statements describing different user/search task scenarios were created. A follow-up survey (<https://www.smartsurvey.co.uk/s/DeepDiscoveries/>) based on the results from these interviews, the first workshop, and the needs of the team at UoS will be made public and advertised through professional and social media channels by all partners, the TaNC programme, and other TaNC Foundation project organisations. The results from the survey will aid in addressing the needs of

users, creators, and collection holders, and in establishing the kind of user interface the prototype might have.

Our software developer has created a simple site (<https://tanc-ahrc.github.io/DeepDiscoveries/>) using the github template generated by the IIF TaNC Foundation Project. The site has a summary of the project, an updated timeline, team profiles, and progress updates, and blog links. A basic interface (Fig. 2) for the backend algorithms has also been created (part of Tasks 1.3 and 1.4). This tool is for our project use, and is currently not live online, though it has been used during the workshop and interview phase of the research. To test the technology, the user can drag and drop an image into the box, and see the results that the computer vision search returns based on the UoS team's currently available algorithms. The images returned are from those collected from our core participant and partner organisations (Tasks 1.1, 1.2, 2.1 and D1.1, 1.3), but do not include those from our commercial partners, for proprietary reasons, even though these images are being used to train the search algorithm on the backend.

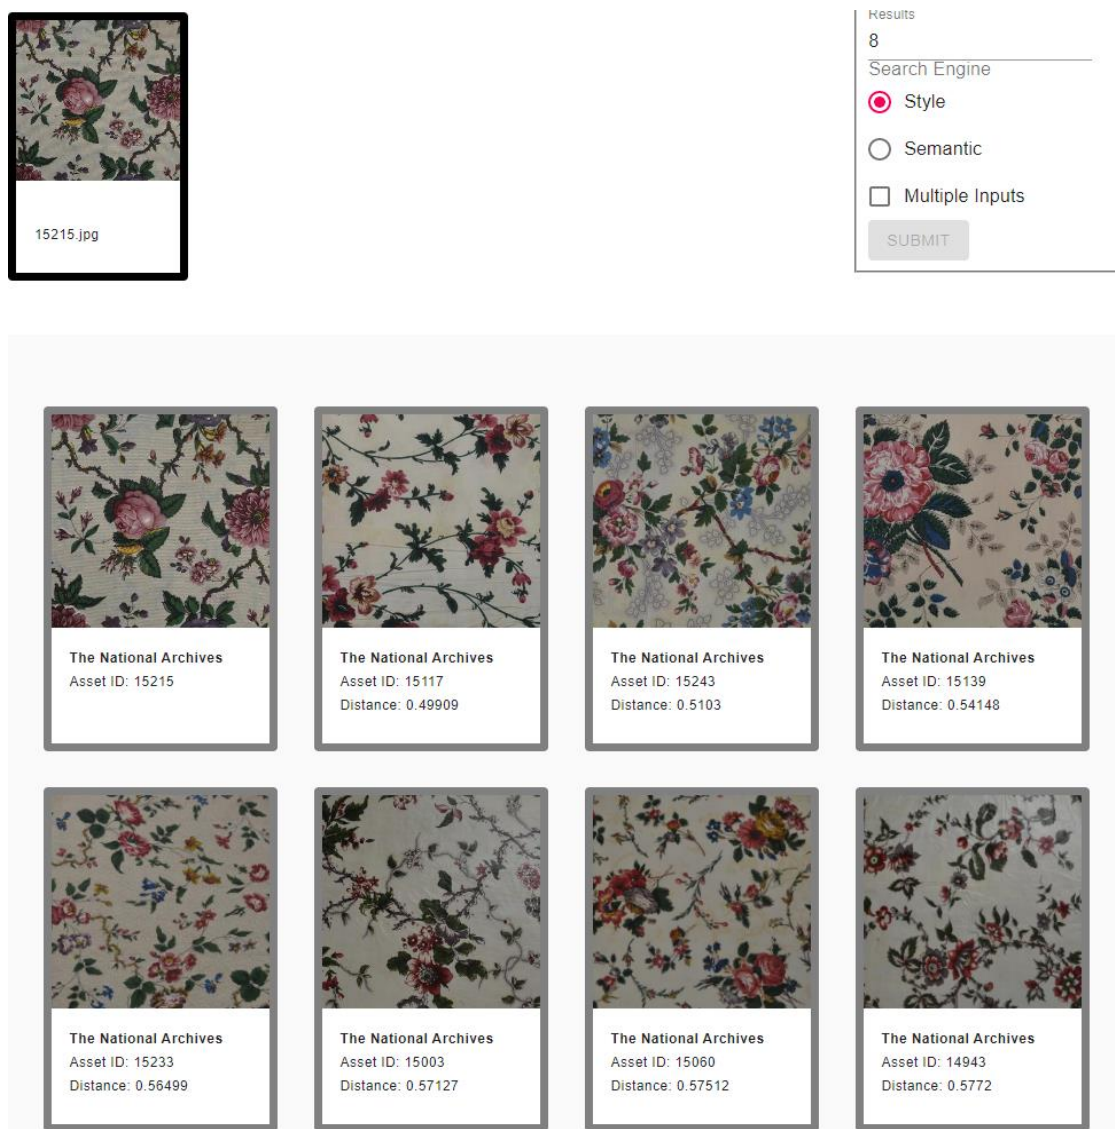


Figure 2. Screenshot of CV search interface used to test the algorithms created by UoS

Through the collected information from a literature review, workshop and interview responses, the UX research team generated a detailed report of the findings, as well as two iterations of a visual search web-based app prototype design. Through conversations with the UoS team, a third prototype will be necessary to accommodate a novel technological challenge that has emerged through the research: the ability for the CV AI to 'explain itself' through the results generated in

response to a user query, in order to allow users to refine the way that the technology carries out a search through interacting with the results. This process presents two advantages: (1) a challenging research question for the technology development, in generating a novel CV algorithm and refinement process that has not yet been demonstrated for visual searching, and (2) through allowing the user to tailor both the search query and modify the search aspects of the results generated in order to create a more targeted search for their specific needs, we can bypass the need to create strict definitions of ambiguous terms like 'style' or 'motif', which have different meanings to different users, and during different tasks (see ethics section below).

Initial surveying suggests, as has been highlighted by others, that vast amounts of the collections we would like to search across are either not digitised, or are only now beginning to be digitised; the format and quality of available images varies from organisation to organisation, and metadata type (if at all available) is also varied. This information will be elaborated on through our Deliverables 1.2 and 2.1. This is problematic when trying to create a proof-of-concept platform, as many of the images we have access to do not link back to a specific website or catalogue entry in the organisations' digital repositories. For the purposes of the project, we will describe these issues in more detail in our report; however, for the purposes of achieving a networked national collection, this will be a very large barrier that may need to be addressed in the next phase of the projects (e.g. if we are to use CV search, what do we want to search across, what do we want to be able to link back to?). If an aggregator for collections is created, current collection holders from smaller and less-resourced organisations about to embark on digitisation campaigns would benefit from workflows around generating images and metadata in formats that allow CV searching. Surveying has also suggested barriers around proprietary collections, or collections that have been digitised by 3rd party commercial partners. Some organisations would like their collections to be searchable and discoverable by new audiences and researchers; however, that would interfere with their commercial ambitions.

Technology progress and output

We have developed CV algorithms (deep neural network models) to support the visual search task of the project. These have been integrated via Amazon Web Services as a RESTful web service to form the backend of the web-based search application developed by the UX team. The model is trained to discriminate between the kinds of objects present in the contributed image, as well as to discriminate between the visual styles of those images. We further plan to fine-tune these models using the metadata/annotations from project partners to improve the search experience. To this end, we have developed a platform for image annotation which may be leveraged to collect annotations for the next phase of the project. The next phase of the project will explore novel 'explainable AI' techniques to explain via visualizations (such as heat maps) which regions of each search result are considered to match the query image, and why. This will form part of an iterative interface to help determine users' search intent and enable interactive refinement of search results under multiple competing definitions ('facets') of visual similarity.

Ethics and access

Two questions around these topics have emerged during the first phase of the project. Although the computer vision search technology is based on deep neural networks, guidance on what constitutes a 'good' and 'bad' return to a search query will be needed through a user annotation activity. In order to carry out this task, we need to firmly decide on the definitions around concepts like 'pattern' or 'motif' and whose (as in, which user group's) definitions are selected can be problematic. Additionally, the people who carry out the annotation work need to be representative, as neither employing our own relatively homogeneous project participants, nor a call for anonymous volunteers via online crowdsourcing platforms would provide a diversified range of helpers. A second issue to be noted in the final report revolves around accessibility. By its very nature, computer vision search is a technology reliant on visual input. Hence, when speaking of diversifying audiences or ways of engaging, we must consider users who may be visually impaired, and instead

rely on Image Descriptions (ID). If organisations would like to use CV search, more work should be focused on ensuring IDs are available for visual collections, and research into linking these with the visual search results should be carried out. Here, again machine learning can be employed to expedite this process.³

Partnerships and engagement

The Deep Discoveries Project Partners have provided thousands of images to the project, as well as joined for project meetings and the first workshop. During the scoping phase of the project, we also connected with the Manchester Central Library Archive, The Wilson Arts & Crafts Archive in Cheltenham, The Whitworth in Manchester, The Warner Textile Archive in Braintree, the Leeds University Library, the University of Brighton Design Archives, and the Scottish Borders Textile Heritage Research Group at Heriot-Watt University. Many of these organisations were not available for participation in the first workshop due to the ongoing furlough situation. However, we will continue engaging with them in future surveys and follow-up workshops.

Publicity generated by the project and the programme has resulted in new connections - we have been contacted by other researchers working on computer vision and design archives or heritage collections: Wendy Andrews and James Campbell (University of Cambridge); Tim Weyrich (UCL) and Karina Rodriguez Echavarria (University of Brighton); Anita Quye (University of Glasgow); Giles Bergel (University of Oxford Visual Geometry Group). We held a Deep Discoveries introductory meeting and information exchange with Art UK Director Andrew Ellis and Tech Manager Terry Gould. The growing network of researchers interested in computer vision applications in heritage has led to the application of a TaNC Discovery-phase project, Patterns of Home, that builds on the Deep Discoveries foundation project and includes several of our projects core participants and project partners.

Next Steps

The next task in Tech Development Work Package is to focus on Task 1.3, which will require user annotation of the returned images in the search (e.g. is this a good match?). The UX research team is working closely with UoS to develop a method for moving this part of the research forward and establishing types of search and discovery that users would like to achieve based on the pilot workshop. To gather further information, the UX research team will also deploy a survey as described above. Following the algorithm development, the team will be working to engineer the interactive search platform with user interface. This technology can then be tested by project partners and core participants, and presented at Workshop 2 to a wider audience for testing and input. A method for assessing the success of such a tool in creating a new way for users to search and discover images in collections needs to be developed during the second half of the project. This assessment can occur during the prototype testing stages. A report on the state of digitised collections and access thereof in the context of creating a vision search platform needs to be produced, which will feed into the final report of the project.

Contacts:

Rachel Smillie – Rachel.Smillie@nationalarchives.gov.uk

Lora Angelova – Lora.Angelova@nationalarchives.gov.uk

John Collomosse – J.Collomosse@surrey.ac.uk

Robert Cubey – RCubey@rbge.org.uk

Joanna Norman – J.Norman@vam.ac.uk

³ Brownlee, Jason. How to Develop a Deep Learning Photo Caption Generator from Scratch (2019) <https://machinelearningmastery.com/develop-a-deep-learning-caption-generation-model-in-python/> Acc. 11/12/2020