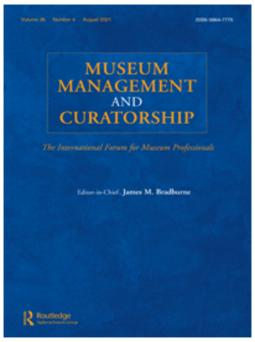

Museum Management and Curatorship

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rmmc20

This is not an apple! Benefits and challenges of applying computer vision to museum collections

Elena Villaespesa & Oonagh Murphy

To cite this article: Elena Villaespesa & Oonagh Murphy (2021) This is not an apple! Benefits and challenges of applying computer vision to museum collections, Museum Management and Curatorship, 36:4, 362-383, DOI: 10.1080/09647775.2021.1873827

To link to this article: https://doi.org/10.1080/09647775.2021.1873827

MUSEUM MANAGEMENT AND CURATORSHIP
2021, VOL. 36, NO. 4, 362–383
https://doi.org/10.1080/09647775.2021.1873827

Routledge
Taylor & Francis Group

This is not an apple! Benefits and challenges of applying computer vision to museum collections

Elena Villaespesa [a] and Oonagh Murphy [b]

[a]School of Information, Pratt Institute, New York, USA; [b]Institute for Creative and Cultural Entrepreneurship, Goldsmiths, University of London, London, UK

ABSTRACT
The application of computer vision on museum collection data is at an experimental stage with predictions that it will grow in significance and use in the coming years. This research, based on the analysis of five case studies and semi-structured interviews with museum professionals, examined the opportunities and challenges of these technologies, the resources and funding required, and the ethical implications that arise during these initiatives. The case studies examined in this paper are drawn from: The Metropolitan Museum of Art (USA), Princeton University Art Museum (USA), Museum of Modern Art (USA), Harvard Art Museums (USA), Science Museum Group (UK). The research findings highlight the possibilities of computer vision to offer new ways to analyze, describe and present museum collections. However, their actual implementation on digital products is currently very limited due to the lack of resources and the inaccuracies created by algorithms. This research adds to the rapidly evolving field of computer vision within the museum sector and provides recommendations to operationalize the usage of these technologies, increase the transparency on their application, create ethics playbooks to manage potential bias and collaborate across the museum sector.

ARTICLE HISTORY
Received 15 April 2020
Accepted 6 January 2021

KEYWORDS
Museum collections; artificial intelligence; computer vision; bias; ethics; data management

1. Introduction

When is an apple not an apple? Is the central question asked by Trevor Paglen in his 2019 exhibition From *'Apple' to 'Anomaly' (Pictures and Labels) Selections from the ImageNet dataset for object recognition* at the Barbican Centre, London. The exhibition examines how training sets, the large banks of photos used to train computer vision algorithms are tagged and categorized. The first work to greet the visitor as they enter the exhibition is a reproduction of Magritte's 1964 painting entitled *This Is Not an Apple*, the work by Magritte is a painting of an apple, with the words Ceci n'est pas une pomme – this is not an apple – painted across the top of the painting. Paglen has added an additional layer to the photographic reproduction of this painting, and included the categories, or

CONTACT Elena Villaespesa evillaes@pratt.edu
This article has been corrected with minor changes. These changes do not impact the academic content of the article.

tags that the machine vision training set, ImageNet applied to the painting when it was analyzed by its algorithm. These tags include the nouns 'red and green apple' (Crawford and Paglen 2019). In the exhibition catalog, Cook explains that these tags matter 'the applications to which these algorithms are being assigned are the actions that will shape our experience of the world' (Cook et al. 2019, 10). If we walk through the process of seeing in this example, we can begin to understand where the tensions lie in the application of computer vision technologies, or algorithmic ways of seeing. The machine sees an apple, on first look, the human eye sees an apple, the artist tells us it is not an apple. It's complicated, however, the human viewer can engage with this work as a surrealist provocation, the computer (or more specifically algorithm) struggles to 'see' beyond the literal, it is an apple, nothing more, nothing less. As such, Paglens work provides a helpful contextual foundation from which to begin to examine the relationship between computer vision, taxonomy, art and objects as culturally variable constructs (McKim 2019). This example helps to situate the conversation, around basic objects and simple nouns; however, museums deal with vast and complex collections, that engage with challenging and disputed histories, diverse cultures and contemporary society. As such, the challenges of applying computer vision to museum collections are equally complex. The problems that arise through the application of these technologies, if acknowledged, documented and critically engaged with, can be mitigated, and the opportunities of computer vision can be utilized to create a more robust, documented and discoverable collection (Murphy and Villaespesa 2020).

Computer vision can help visitors to engage with collections in new ways, and help curators to develop new insights into objects that they may not have had the time or resources to research in an analogue manner. It is for these reasons that computer vision is fast becoming a potential instrument to enrich museum collections data in a diverse range of ways, which ultimately can have an important impact on the user experience. This research paper seeks to examine the current practice of the usage of computer vision applied to museums' collections. More specifically, the work presented here investigates the current and potential use cases of computer vision for collection data, identifies the current opportunities and challenges of computer vision and maps the processes, organizational structure, operations, funding and evaluation of these types of initiatives. Due to the novelty of the subject and early stages of the productive usage of computer vision within a museum context, there is currently limited peer-reviewed research on the processes and challenges behind these initiatives. This paper presents in detail the experience of five museums with computer vision technologies and discusses the practical implications for the future of these technologies in museums. As such, this paper provides a rigorous foundation for future research in the area of computer vision technologies within the context of museum practice.

2. Computer vision technologies for museum collections

Included under the large and diverse umbrella of artificial intelligence (AI), computer vision, also known as machine vision, comprises the computational methods to obtain information and properties from visual content such as images and videos (Davies 2012; LeCun, Bengio, and Hinton 2015). There are different applications of computer vision, for instance, image processing, features detection, 3D appearance modeling, motion analysis, object classification and facial recognition, among others (Szeliski 2010). Computer vision looks at individual pixels and the features that are derived from them, seeking patterns in their variations, that correspond with patterns the computer has previously encountered. For computer vision, a pattern could include, two eyes, a noise, a mouth, equals a face. As the algorithm reviews more images, it begins to develop a nuanced understanding of a pattern, and as such can recognize that whilst a human and a dog have the same features, they are distinct objects.

In 2011, Veirum et al. argued in this very journal that 'if something is not to be found on the Internet, it probably does not exist at all' at the heart of this provocation is not a rallying cry for digitization, but instead a rallying cry for the importance of discoverability of digitized objects (2011, 7). Indeed, it is the need for digitized objects to be easily found, or discovered that makes these visual processing algorithms so promising for the management of museum collections online (Majd and Safabakhsh 2017; French and Villaespesa 2019). Computer vision technologies can generate data from the digital images of collection objects at a very fast pace compared to the speed of producing these data manually by museum staff. There are several experiments which act as a helpful proof of concept for these technologies of particular note are projects at Hirshhorn, Auckland Art Gallery, Rijksmuseum (Mensink and van Gemert 2014; "Hirshhorn Eye - Hirshhorn Museum and Sculpture Garden | Smithsonian" n.d.; "Auckland Art Gallery's New Chatbot Demonstrates Artificial Intelligence to Give New Access to 17,000 Artworks | Auckland Art Gallery" n.d.). These projects have been documented through museum, and vendor websites; however, in order to get grips with the technologies behind these projects, it is important to establish a contextual foundation, or shared understanding of computer vision applications in museums.

Extracting the physical elements of an object such as color or shape can be considered to be a basic application of computer vision technologies. The results can be implemented on museum websites offering new ways of discovering the collection. For example, on the Cooper Hewitt Smithsonian Design Museum and Dallas Art Museum websites, users can filter by shape and direction of lines, on the Barnes Foundation website users can filter by space and light. While the Victoria & Albert Museum is taking a more cautious approach, testing a visualizer that allows the user to explore a small sample of their collections by color, visual texture and shape on an experimental section of their website.

A more complex application of computer vision technologies is finding similarities and connections across images through 'object recognition, facial recognition, colour and composition analysis' ("IK Prize 2016: Recognition – Exhibition at Tate Britain | Tate" 2016a). At present, museums rely on third-party algorithms, or off the shelf tools to utilize these technologies from technology companies such as Microsoft, Google and IBM. These sophisticated computer vision tools have been trained using millions of images to create an algorithm that can identify visual trends and patterns. An early example of this approach is the winner of the Tate IK Prize in 2016, *Recognition*, created by Fabrica, a communication research center based in Italy and JoliBrain AI specialists from France. *Recognition* utilized a range of AI platforms that cross-tabulated live news stories with the Tate Collection, the results of this analysis were streamed live in a gallery at Tate Britain with visitors asked to augment what the computer saw through their own feedback (or quality control) ("Recognition" 2016b).

Other novel applications of machine vision technologies within the cultural sphere include Google's X Degrees of Separation which finds relationships between objects from two collections, or the Google Arts & Culture app which allows users to upload a picture of their face, and find a portrait with similarities. ("Google Arts & Culture Experiments - X Degrees of Separation" 2018). Whilst novel applications gain traction on social media, these technologies are also being used to support curatorial research on art and objects. Saleh et al. from Cornell University developed a tool that compared artworks for similarities of technique, tone and content to identify similarities beyond known art movements, their work provided a template for the creation of new insights for art historians and curators that could help to develop new conversations about art and artists (2014).

Museums have also begun to explore how they can operationalize these technologies, to develop new insights into their collections databases. One key use of these technologies is in detecting elements depicted in a digital image of a museum object to generate subject tags. Browsing by subject tag is a valuable way for users to discover collection objects (indeed without them, these tags digital images remain undiscoverable within the database). As an example, computer-generated tags have brought together related objects that were otherwise distributed across many subcategories of the collection such as the Science Museum Group's collection of dolls which are in multiple curatorially defined categories (Figure 1). However, there are clear challenges raised around tags generated through computer vision in terms of accuracy which needs to be thoroughly considered (Bernstein 2017; Pim 2018; Choi 2019). The application of computer vision technologies in a museum context is further problematized, because beyond the accuracy of detection, subject tags are as the name suggests subjective. Subject tags that are commonly used within a museum are defined from a specific worldview, and the worldview of those that trained the machine may be different from that of the museum, as an institution, or indeed the worldview of their visitors (Davis 2020).

In order to utilize computer vision in a way that serves their visitors, and, or users, museums are faced with the challenge of creating a solid foundation for emerging practice that acknowledges the opportunities these technologies present, whilst mitigating their limitations.

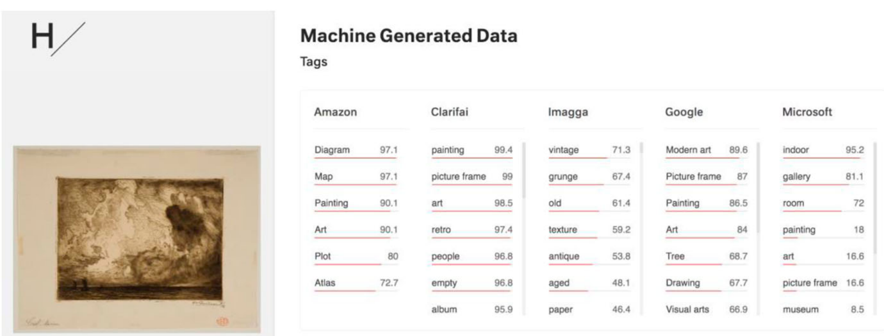

Figure 1. Screenshot of the Harvard Art Museums – AI Explorer website displaying the results of the machine-generated tags for the artwork *Sky; Marine* https://ai.harvardartmuseums.org/object/264992.

3. Methodology

The discussion presented in this paper is informed by the presentation of case studies during the action research project *Museums + AI Network* and a series of in-depth semi-structured interviews conducted between October 2019 and January 2020. The museum professionals interviewed for this study were participants in the *Museums + AI Network* and had been selected to be so because they had used computer vision technologies within their own museums. The *Museums + AI Network* brought together a range of senior museum professionals and prominent academics from the United Kingdom and the United States to develop the conversation around artificial intelligence, ethics and museums ("The Museums+AI Network" 2019). Two professionally focused workshops took place during the project, one in June 2019 at Goldsmiths, University of London, and one in September at the School of Information at Pratt Institute, New York. To dig deeper into the case studies presented during the two workshops, a series of interviews were conducted with five museum professionals that are using these technologies with a focus on their collection data. These museum professionals have a deep understanding of the collection information systems and website technical operations. The interviewees were: Dan Brennan (Museum Application Developer at Princeton University Art Museum, USA), Jennie Choi (General Manager of Collection Information at the Metropolitan Museum of Art, USA), Shannon Darrough (Director, Digital Media at the Museum of Modern Art, USA), John Stack (Digital Director at the Science Museum Group, UK) and Jeff Steward (Director of Digital Infrastructure and Emerging Technology at Harvard Art Museums, USA).

The semi-structured interviews consisted of 12 questions which focused on the opportunities and challenges that AI brings to museums, the projects the interviewees have worked on using computer vision, the resources and funding employed and the ethical implications that arise during these initiatives. Moreover, at the end of each interview, there were two questions that explored how interviewees anticipate the usage of these technologies in the near future (see list of questions in Table 1). Interviews lasted ~45 min each.

Table 1. Questions of the semi-structured interviews.

1. Please tell me about your role and your work with collection data at your museum?
2. What opportunities do you think AI brings to museums?
3. How are you currently using AI in your museum? Please describe the projects you are currently doing.
4. What is the purpose – creative or business need for your current use of AI?
5. What are the resources employed on these projects? How many staff are working on this? What department does this work sit within? Are you collaborating with external partners to develop AI?
6. How is your current AI work being funded? What opportunities and challenges do you see around financially supporting AI projects in the future?
7. What have been the main challenges (internal and external) in working on AI?
8. Is ethics something you talk about or strategically engage with when planning digital projects? If so, what are the ethical implications that these projects involve? How is this defined? What is the museum position and work in this area?
9. How are you evaluating the success of your AI initiatives? What are the key performance indicators that you are tracking to measure the impact of AI?
10. What are your future plans with AI? How do you see AI in 5 years in the museum sector?
11. If time and money was not an issue, what is your dream AI project to work on?
12. Please share anything else we did not ask and you would like the world to know.

Using qualitative data analysis software (QRS NVIVO 12), both the interview transcripts and the notes from the Museums + AI Network meetings were coded and analyzed using grounded theory methodologies in order to search for themes in their experience applying computer vision (Corbin and Strauss 2014). The first step in the process involved labeling the interview transcript and meeting notes and breaking the text down into detailed codes. The next step in the analysis was to find the contextual links between the codes to group them together to then create themes based on these categories. These themes are presented in the Results section of this paper along with interviewees' quotes and concrete examples of their computer vision projects.

4. Application of computer vision to museum collections: five case studies

This section presents a summary of the projects the five museums who are examined in this study have been working on in relation to the usage of computer vision to interpret their collections. Table 2 summarizes key contextual data on each of the museums, these data were obtained from the information on their website, online collections, annual reports or directly provided by the interviewees.

4.1. Harvard Art Museums

Harvard Art Museums' Collection has 250,000 objects but <1% of them are on view. Therefore, the way to access these objects is through their website. The museum has been exploring alternative methods for categorizing, describing and tagging the collection objects, especially for users that are new to art and don't have the knowledge to search for style, periods or art history terms on the website. They have been using the following computer vision services: Microsoft Cognitive Services, Google Vision, Imagga, Clarifai and AWS Rekognition (see the example in Figure 1).

The museum provides direct access to the machine-generated data via the API (Harvard Art Museums 2020) and a website interface called AI Explorer (Harvard Art Museums n.d.a, n.d.b) and has applied these results in a series of playful experiments including, for example, 'Magic Message,' where users receive image fragments based on a sentence, or 'Face Match' that invites the user to add faces to the corresponding bodies (Harvard Art Museums n.d.a, n.d.b).

Table 2. List of the museums that participated in the study. Data accurate as of April 2020.

Museum	Location	Number of visitors	Collection objects	Collection objects online (as of January 2020)
Harvard Art Museums	Cambridge, USA	145,000 (avg. attendance past 3 years)	250,000	233,977
Princeton University Art Museum	Princeton, USA	206,622 (2018–2019)	111,700	54,178
Science Museum Group	London, Manchester, York, Shildon and Bradford, UK	5,210,000 (2018–2019)	7,300,000	325,700
The Metropolitan Museum of Art	New York, USA	7,027,858 (2008–2019)	1,500,000	490,000
The Museum of Modern Art	New York, USA	3,000,000 (2017–2018)	200,000	83,537

4.2. Princeton University Art Museum

Princeton University Art Museum is exploring how computer vision can be implemented in collections to enable object, scene, text and facial recognition. As part of the process, they consider it important to integrate this work into the existing data pipeline and expose these results as data and image annotations. The museum has tested this approach with two projects. The first one, a prototype, consisted of identifying and annotating Mayan glyphs (Figure 2) with the intent of building a model that could be used to surface shared characteristics. The second project, led by a computer science student at the University, is focused on identifying Chinese paintings to find shared visual characteristics across and within dynasties (Kong 2020).

4.3. Science Museum Group

The Science Museum Group is currently engaged in a mass digitization project of their collections. The museum is using machine learning with AWS Rekognition (Amazon) to generate tags of their objects, and although they are not shown to users on the website, these are being considered for use in the backend, in the search index (powered by Elasticsearch) that is used for the Online Collection and for the Public API. The goal of the usage of computer vision to generate tags is to improve the discoverability of collection records that have little metadata, allow discovery through a visual lens, create synonyms and non-academic terms, and create visual relationships between objects in the collection (Figure 3). The tags generated show a long-tail shape in their usage distribution and the museum is evaluating the accuracy of the results (Stack 2020).

4.4. The Metropolitan Museum of Art

The museum currently contains more than 450,000 digitized records and is growing in number with each passing week. Major collections belonging to the museum include

Figure 2. Training a model to identify Mayan glyphs.

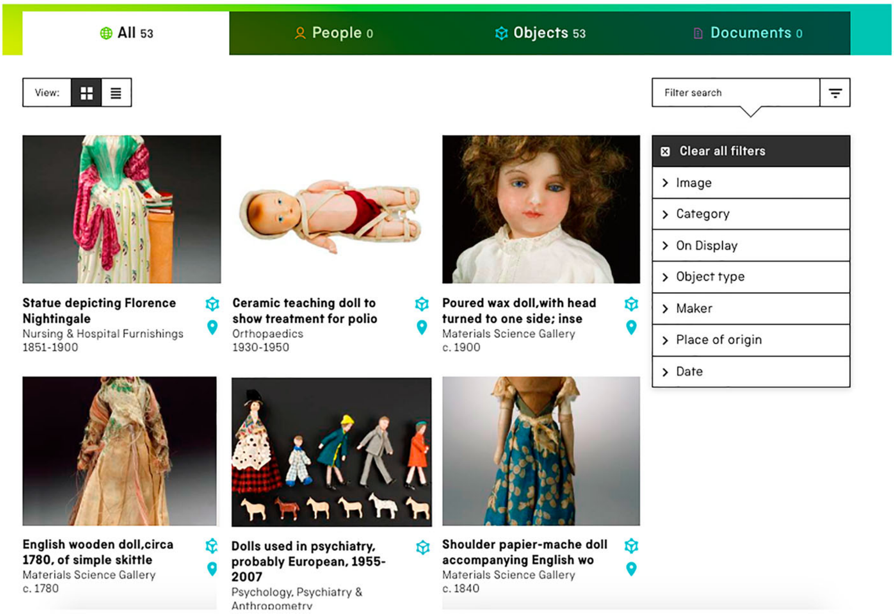

Figure 3. Search results for the computer-generated tag 'doll' on the Science Museum Group's website https://collection.sciencemuseumgroup.org.uk/search/imgtag/doll. Image credit: Science Museum Group. Statue depicting Florence Nightingale. A661037. Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co121484/statue-depicting-florence-nightingale-figurine; Ceramic teaching doll to show treatment for polio. 2002–360. Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co8015261/ceramic-teaching-doll-to-show-treatment-for-polio-teaching-doll; Poured wax doll,with head turned to one side; inse. 1999–1183, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co522862/poured-wax-doll-with-head-turned-to-one-side-inse-wax-doll; English wooden doll,circa 1780, of simple skittle. 1999–1184, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co522863/english-wooden-doll-circa-1780-of-simple-skittle-wooden-doll; Dolls used in psychiatry, probably European, 1955-2007. 2006–124, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co8078338/dolls-used-in-psychiatry-probably-european-1955-2007-toys-recreational-artifacts; Shoulder papier-mache doll accompanying English wo. 1999–1185, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co522864/shoulder-papier-mache-doll-accompanying-english-wo-papier-mache-doll.

American paintings and sculpture, European paintings, Egyptian art, arms and armor, the art of Africa, Oceania, and the Americas, ancient Near Eastern art, Asian art, costume, drawings and prints, European sculpture and decorative arts, Greek and Roman art, Islamic art, medieval art, modern and contemporary art, musical instruments, photographs and the Robert Lehman Collection. With such a large and diverse collection, an ongoing challenge faced by staff at The Met is developing new ways to document, and interpret the museum's collection in a way that will allow it to become searchable and browsable online. Many objects that have been digitized have very little information to support them, which means that whilst a digital image exists, a lack of metadata or keywords signifies that a user is not able to discover these items through search. The museum is working on the generation of tags manually and testing with computer vision. The goals of tagging the museum collection are to increase user engagement, improve search and discovery of the collection, make the collection accessible to the widest possible audience and explore using tags as training data for AI models. The museum has tested the usage of various computer vision technologies including Google Vision and Microsoft Azure to generate tags automatically.

4.5. The Museum of Modern Art

The Museum of Modern Art collaborated with Google Arts & Culture (GA&C) in 2018 on a project that used an object recognition algorithm to identify collection artworks in a large trove of historical exhibition photography. The museum had recently put online over 30,000 photos that documented the exhibition program from the institution's inception in 1929. The photos formed a remarkable visual overview of MoMA's exhibition program; however, the works depicted in these images had not been documented and as such, the photos were not searchable. GA&C's algorithm matched works contained in the exhibition images with images of works from the museum's online collection. The museum was then able to incorporate the resulting data set into the exhibition history and online collection section of MoMA's website. This allowed users to see which objects appear in exhibition images and, conversely, in which exhibition images an object is present (The Museum of Modern Art 2018). Another outcome of the project was the development of tools for the imaging team to manually identify artworks in images. With this in place, imaging workflows have been modified so that virtually all works that are documented in the online collection and also appear in exhibition images have been identified.

5. Results

5.1. Overall results

Analysis of the interview data and project presentations shows that while computer vision presents significant opportunities to enhance collection data, museums are still in an early stage of adoption when it comes to the application of data created through computer vision tools. Within a museum context, the adoption of these technologies has been slower than other sectors, this study identifies a number of underlying tensions that have created sticking points, including a lack of resources, challenges with algorithmic biases and internal resistance to the accuracy of data outputs. The results of this study are presented here thematically.

5.2. Possibilities to offer new ways to analyze, describe and present the collection

The presentation of case studies and interviews manifested a range of usages of computer vision including automated subject tagging, colors and patterns identification, face and object recognition, relationship creation between collection objects, and label production, among others. As Table 1 shows, the number of objects that these museums

Figure 4. Examples of the tags generated for two objects.Image credit: Science Museum Group. ICI fertiliser, 1977, ICI no.8 NPK composite ferti. 1977-270/6, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co8021776/ici-fertiliser-1977-ici-no-8-npk-composite-ferti-chemical-fertilizer; Beyer Peacock Printing Block. Y1966.24.1.717, Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co8595648/beyer-peacock-printing-block-printing-block.

Source: https://johnstack.github.io/what-the-machine-saw.

have in their collections is enormous and the number of objects getting digitized continues to grow. The human effort to document the objects would take a tremendous amount of time; therefore, computers could play a productive role in this gigantic task. Two interviewees provided specifics insights around speed, and data production and a common focus on collections as being dynamic and growing:

> We acquire thousands of objects a year, we digitize objects weekly. So, we have a growing need to identify subjects in our artworks and I think computer vision is a potential solution to this though it is not perfect at this point.

> The collection is absolutely huge and we digitize it at a quite fast pace […] the records are so thin that we started thinking if there is a way of creating additional metadata for these objects that would help their discovery online. So far the experiments have been around keyword

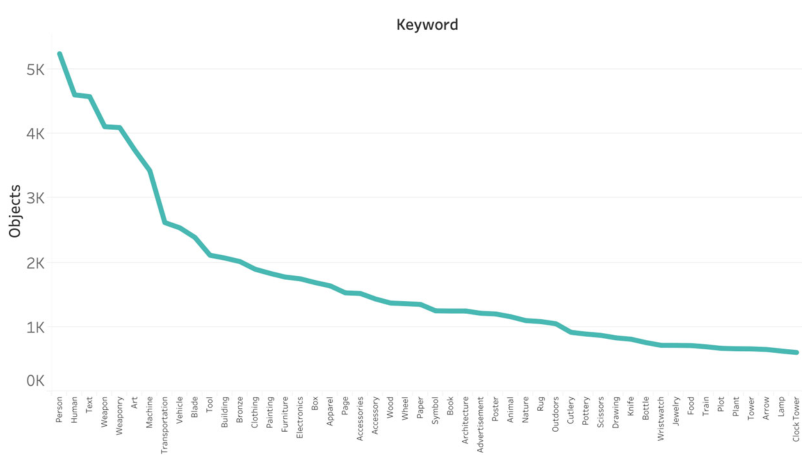

Figure 5. Distribution of usage of the top 50 tags.

tagging of images which are on one level really spookily accurate and on another level, kind of completely out there.

For instance, in the case of the Science Museum, AWS Rekognition generated 1811 unique tags with a confidence level of minimum of 70%. These tags were used 163,792 times on the collection objects (See examples in Figure 4). The distribution of the usage of these tags shows that there is a small number of tags used by a high volume of objects and that there is a significant proportion of tags that are used only a few times (Figure 5). Generic terms such as ‘person’, ‘human’, ‘text’ or ‘art’ have been frequently applied to the collection objects. However, the long tail of the tags distribution designates more specific terms. As an example, ‘sports’ is a tag applied to 283 objects, but specific terms related to sports are assigned for particular objects such as baseball glove, soccer ball, tennis racket or football helmet.

The opportunities are not only to create metadata for these records but also to diversify the information that currently exists on the database. Computer vision can find other ways of describing collections that go beyond what museum staff can do. These computer-generated features and data can have a significant impact on providing access and creating more open and friendly paths to navigate through the collection for users

			
	Mary Sylvester 1754 By Joseph Blackburn	*Hermann von Wedigh III (died 1560)* 1532 By Hans Holbein the Younger	*Euphemia White Van Rensselaer* 1842 By George P. A. Healy
Computer generated tags	halloween witch, person, waxwork, kirtle dress, **dress,** weeds, **garment**, celebrity, adult, clothing, fabric, olive green color	clothing, **hat**, apparel, person, human, art, painting, coat, advertisement, poster	apparel, clothing, art, painting, human, person, fashion, cloak, **hat**
Human tags	women, portraits, sheep	men, portraits, latin	portraits, women

Figure 6. Human and computer-generated tags for three artworks from The Met’s collection. Image credit: The Metropolitan Museum of Art. Mary Sylvester. 10174. Accessed January 18, 2021. https://www.metmuseum.org/art/collection/search/10174; Hermann von Wedigh III (died 1560). 436658. Accessed January 18, 2021. https://www.metmuseum.org/art/collection/search/436658; Euphemia White Van Rensselaer. 11055. Accessed January 18, 2021. https://www.metmuseum.org/art/collection/search/11055.

Still Life with Watermelon, 1822

Amazon	Clarifai	Imagga	Google	Microsoft
Plant	fruit	fruit	melon	plant
Fruit	food	edible fruit	still life	dish
Food	no person	food	fruit	fresh
Melon	grow	healthy	watermelon	
Watermelon	sweet	watermelon	painting	
Grapes	melon	strawberry	still life photography	
	still life	fresh	citrullus	
	confection	juicy	artwork	
	table	delicious		
	healthy	melon		
	juicy	diet		
	diet	nutrition		
	leaf	berry		
	nutrition	sweet		
	vegetable	produce		

Figure 7. Top 15 tags generated for *Still Life with Watermelon.* Image credit: Harvard Art Museums. Still Life with Watermelon. 2006.4. Accessed January 18, 2021. https://harvardartmuseums.org/collections/object/4988.

with no expert knowledge in the subject. This view was repeatedly mentioned in the interviews:

> Our staff are rooted in deep scholarly art history and that only covers a small subset of perspective, so we're dealing with highly subjective material and our historical perspective is one – again, it's just one narrow way to perceive – to consume art. So, through AI, we are looking at ways of just finding other systems to describe collections that use completely different terminology than we would by default.

The idea of adding tags had already been explored by The Met but in a more manual way. In 2018, The Met added, working with an outside vendor, subject keyword tags to 300,000 artworks in their online collection (The Metropolitan Museum of Art 2020). The goal of the project was to improve search and discovery of the collection, increase user engagement and provide a new access point around depicted subject matter (Murphy and Villaespesa 2020). The same goals could be potentially achieved using algorithms and the museum has tested the water with various computer vision platforms. In some cases, these tags expand upon and add further details that the human tags do not include. Because of the large size of the collection, details about clothing and accessories were not included in the human tags, but were often identified by the machine tags. For the portrait of *Mary Sylvester*, IBM Watson Visual Recognition included tags for dress and garment, whereas the human tags only included women, portraits, and sheep. Other examples include tags for a portrait by Hans Holbein the Younger and a portrait of *Euphemia White Van Rensselaer*. Amazon Rekognition identified hats in both paintings, which were not included in the human tags (see Figure 6).

Here is another example, this time from the Harvard Art Museums, which was presented during a *Museums + AI network* meeting, and illustrates the diversity of terms that could be generated. For a painting of a still life, most algorithms were able to identify that fruit was presented in the image and specifically which fruit (melon, watermelon, grapes). Moreover, Imagga and Clarifai also generated keywords related to the feelings and experiences of eating them (fresh, delicious, sweet) (see the top 15 tags generated by each system in Figure 7).

A secondary outcome of automatically generating new metadata such as subjects, locations or historical events is the possibility to connect to other platforms such as Wikidata which opens the door to the automated translation of the content to other languages, the generation structured data that machines can easily read and the usage of the wiki tools to examine the data. This opportunity is currently being explored by The Met who is actively collaborating with Wikimedia (Lih 2019).

5.3. Computer vision applications are experimental

A common theme amongst interviewees was that they are in an experimental phase of using these AI technologies. The experimentation model and structure vary per institution with the establishment of either internal or external collaborations. In some cases, this innovative work is generated solely by one or a few members of staff in the digital department. Hackathons, innovation labs and computer science competitions online are some of the formats that have been practiced by the museums in this study.

For the past two years, The Met's open access dataset has been used in a data science competition on Kaggle ("IMet Collection 2019 - FGVC6" 2019). Kaggle is a platform and online community for data scientists to solve challenges around AI and machine learning. The competition focused on fine grained attributes which go beyond basic image recognition and attempt to identify specific subject matter and characteristics in a given image. Participants were provided a training dataset which included data and tags for a subset of records and a test data set for which participants had to create a model that predicted appropriate tags for each image (Zhang et al. 2019; Choi 2020). Leader boards and submitted models were all public and are still accessible on the Kaggle site. Jennie Choi commented on this experience:

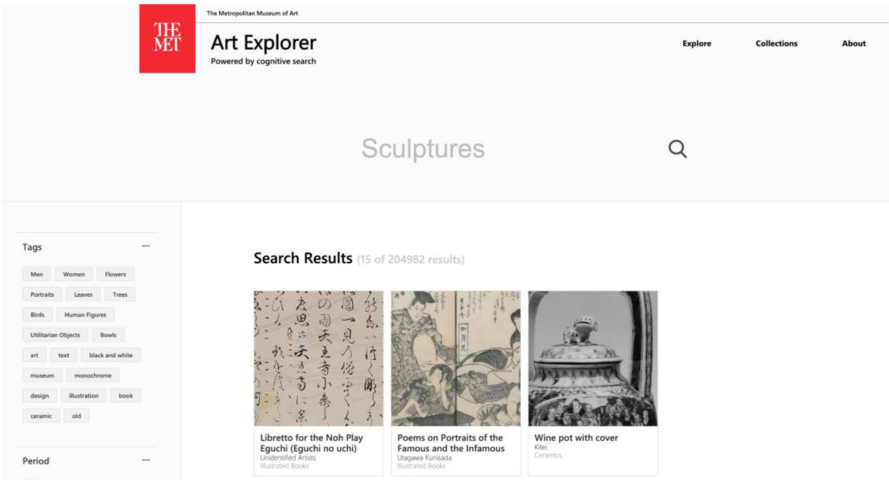

Figure 8. Screenshot of the Art Explorer homepage (tags are presented on the left-hand side and the search box suggests options for users to search based on those tags).

> Though we were unable to view the results or apply the models to other records in our collection due to a lack of AI expertise in-house, we were able to see a new audience of data scientists engage with and be inspired by our collection. We hope participating in these competitions will result in new advances in AI that will benefit other art collections in the future.

The Met also collaborated with Microsoft and MIT in a two-day hackathon around AI in 2018. Five prototypes were created during the event with projects using a range of tools including voice recognition, generative adversarial networks (GAN), tag prediction and the use of AI to discover artworks based on current events ("The Met x Microsoft x MIT" 2018; Kessler 2019). Similar to the Kaggle competition, results have not been

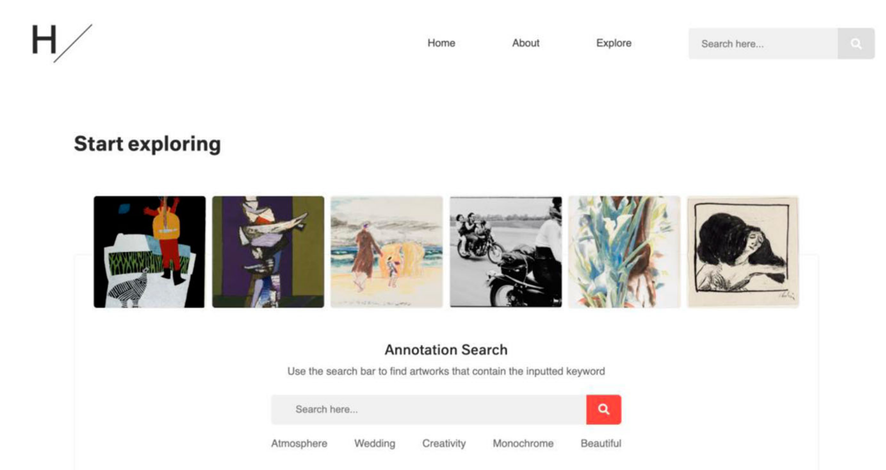

Figure 9. Screenshot of the AI Explorer website where users can enter a keyword or click on a suggested keyword. https://ai.harvardartmuseums.org/explore

implemented into a production environment due to required staff time and the need for ongoing expertise around AI to maintain the projects.

Another format to trigger this work is making these data openly available via a public API or file downloads on GitHub, so users with different backgrounds can play and develop data and interfaces to interact with the collection. For those museums embedded in a University such as Princeton or Harvard, there have been some connections established with individual students or research groups and labs interested in exploring the collection data with these technologies. Another model adopted by these museums is the establishment of a partnership with technology companies such as Microsoft and Google.

One of the most remarkable themes is that besides the opportunities that computer vision brings in terms of creating new user-friendly vocabularies and generating new data to interact and access the collection. Some prototypes of potential digital experiences are The Met's Art Explorer and Harvard Art Museums' AI Explorer and IIIF Explorer.

After The Met's hackathon, the museum continued its collaboration with Microsoft in the development of the proof-of-concept site *Art Explorer* ("The Met - Art Explorer" n.d.) (see Art Explorer's homepage in Figure 8). This project utilized Microsoft's AI cognitive search technology in providing improved access to the Met's collection. Art Explorer includes additional search options for visually similar artworks, tag prediction for image content, data enrichment by linking to other data sources like artist biographies on Wikipedia, and an interactive visual graph showing connections between artworks based on time period, artists, subject matter and medium. Each of these features provides AI-based search and browse capabilities for the collection. Due to the technical requirements needed to integrate Microsoft's cognitive search tools with our existing site, the project remains a proof-of-concept and has not been implemented on the museum's website.

The experimental work that Harvard Art Museums has done testing five different computer vision systems is presented on the *AI Explorer* website (see AI Explorer's explore page in Figure 9). For each individual artwork, the tags, captions and object, face, and text recognition from each system is displayed on the web page. The site includes a search functionality that allows users to search for keywords and find artworks that include those computer-generated tags. Another site developed by Harvard Art Museums is the *IIIF Explorer* site which includes eight different experiments using data generated from computer vision systems. For example, *Magic Message* composes sentences using pieces of artwork images from the collection as the user types words.

These examples show that there is very limited implementation of these results in the user interfaces such as the official museum website. As one interviewee said:

> We're not surfacing that right now in any real way, because we still have this kind of internal debate about how do we present that this is not something that we think it is something that an algorithm thinks.

The adversity to risk in museums to display information that is not completely accurate or that has not come from academic sources is the main reason for not surfacing the computer vision results to users. This was a common concern expressed by interviewees:

> The bigger challenges internally are convincing staff the value of using services to do descriptions. Especially when a service can be so wrong and those, you know, it's just culture, a museum culture issue with trying to get staff to be okay with allowing terms that they

> might not fully agree with but they can't unequivocally say is wrong, so it's that problem like letting other people have opinions about the art.

In the case of MoMA, they took a conservative approach where they prioritized accuracy over coverage, which involved significant quality assurance work and algorithm refinement until the results were showing minimal mistakes. On the website, they provide an email address in case any user spots an error in the artworks identified.

One challenge expressed regarding the application and sustainability of computer vision and other artificial intelligence algorithms is the rapid obsolescence rate. This field is evolving in such a speedy fashion that the algorithms, as they become more accurate and precise, would need to be applied literally over and over again and the human effort to check, tweak the results and implement the changes in the production of the digital product will also need to be repeated. This requirement to iterate the computer vision results to improve the final digital output overtime is not compatible with most of the funding models in museums which tend to be temporary and project-based.

The early adoption of computer vision means that there is no formal strategy and structured capabilities in the usage of these technologies, including not having a defined evaluation framework to assess the success of these activities. There were some expressions in the interviews of what success would look like which are linked to the opportunities described earlier in this paper about access, discoverability and data enrichment. Improving the user experience with serendipitous journeys and increasing traffic to objects that have small or no web traffic are common measures of success mentioned by the interviewees:

> I consider it a success if we learn something about our collections that either drives in the research or exposes the collection to audiences that have not been exposed to so far, or forces us to reconsider our collections in certain ways, then I will consider that a success. In terms of identifying like, in actual like, quantifiable KPI of any sort, I'm still not sure what that would be.

> Are more people finding our stuff? Are we getting more page views? Are people finding our objects on other platforms? Do we have increased engagement on Wikipedia? I think those would be the main indicators because those were our original goals, to engage audiences and to make our collection more accessible. So, if we get more eyes on our objects whether it's on our own website or third-party website through these tags, I think that would indicate success.

5.4. Bias existence in the whole data life cycle process

Computer vision technologies bring a set of notable algorithm biases. The presentation of the results obtained by the participants running these systems highlighted the challenges that arise in each of the phases of an AI initiative from the data collection to the training, application and evaluation of the results.

Looking at the first step, the data input, there is already a bias in the collection data as each collection has different origins, donations history and new acquisition policies, and therefore, there may be gaps in the collection and what is represented. One of the participants mentioned this initial challenge:

> Our collection is inherently biased. So, even if we get the perfect machine learning model, it's going to be biased which, you know, that's just what we have to live with. That's just how our

The Met	Science Museum Group	Harvard Art Museums
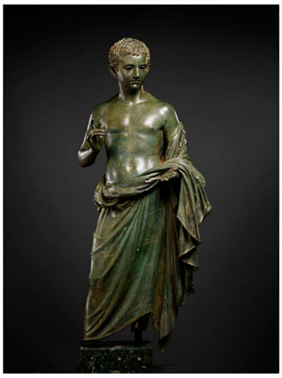		
Bronze statue of an aristocratic boy 27 B.C.–A.D. 14	Sylvester, coal mining jacking device. C.1940	Dagger Blade
lady (Google Cloud Vision)	bracelet, accessories, jewelry, chain (Amazon Rekognition)	echinoderm, invertebrate, sea cucumber, bread (Imagga)

Figure 10. Examples of incorrect tags generated by computer vision algorithms. Image credit: The Metropolitan Museum of Art. Bronze statue of an aristocratic boy. 248891. Accessed January 18, 2021. https://www.metmuseum.org/art/collection/search/248891; Science Museum Group. Sylvester, coal mining jacking. Y2002.19/T879. Science Museum Group Collection Online. Accessed January 18, 2021. https://collection.sciencemuseumgroup.org.uk/objects/co8413000/sylvester-coal-mining-jacking-jack; Harvard Art Museums. Dagger Blade. 1991.43. Accessed January 18, 2021. https://harvardartmuseums.org/collections/object/303702.

> collection is formed. Most of our artists are men. We have a much smaller percentage of our collection from Africa, South America, and Native America. So, we're sort of perpetuating this.

The second step in the process, which is related to the first challenge encountered, is model development and training, which raises a lot of questions about whether museum collections are valid training datasets, due to the representativeness of the data, or as one of the participants pointed, due to not having enough training data:

> We only have 600 cats and typically you need tens of thousands of training records to properly train machine learning algorithms. So, even though our collection is large, when you look at the individual numbers of tags and objects per tag, most of our tags, I think, the statistic is more than 50% of our tags have less than a thousand occurrences and a lot of our tags had less than a hundred but those happen to be some of the most interesting tags.

For example, in the case of MoMA, the algorithm worked well to identify paintings in particular and other two-dimensional objects in the exhibition photos. However, it was not very good at identifying three-dimensional objects or more complex types of art such as moving image, performance or installation works. Anecdotally, the team noticed that when it came to these types of works the image in the collection database had often been taken from a different perspective or point in time, than that of the installation photograph, and the technology was not able to cross-tabulate these differences. Challenges with the computer vision algorithm's outputs also happened in other museums. Several examples of inaccurate results were also presented during the Museums + AI Network project (see examples in Figure 10).

All of the case studies examined use third-party algorithms from various technology companies to process and extract information from images, some of which have already got problematic results. The so-called black box of these tools was repetitively mentioned by the interviewees as a problematic and ethical consideration in the process:

> Just working with these third-party, huge multinational corporations like Google, Microsoft, Amazon, they have our data. What does that mean long term? What can they do with it? So, there are ethical issues around that.

> This artwork is depicting a challenging subject and we know by extension that like the way it can be used in an AI system by an external user could go against our own ethics.

The challenges around inaccuracy and algorithm biases raise important ethical questions about how the implementation of these tools could affect the reputation of the museum. While ethics is definitely a concern among the interviewees, due to the exploratory nature of the projects, museum ethics policies do not currently cover these technologies. Interestingly, one of the museums in this study has an ethics group that recently started to discuss the ethics of data usage, which could be a potential model to be adopted by other museums. Another museum from this study has a Digital Advisory Group which consists of three trustees and three external experts that is interested in having an agenda item in their meetings to look at AI and data and the associated ethics questions.

6. Discussion and practical implications

6.1. Operationalizing AI: from experimentation to practical application and integration in museum systems and processes

Most of the cases examined and presented through the *Museums + AI Network* workshops and discussed in detail in this paper show that these are pilot projects. However, AI technologies are here to stay and the prediction from the interviewees is that the advantages of using these tools will surpass the challenges around inaccuracy and bias. One interviewee argued that internet users are becoming more aware of the usage of AI:

> I think a lot of these things that people are doing as experiments now with their collections in AI will kind of become the norm at a certain point. And five years, that's probably a good window for that to happen, I would say where all of these kinds of concern about how do we identify computer-generated data versus human-generated data and what are the implications of that, I think to our audiences which are becoming like, increasingly smart and adept to identifying those things, it will just kind of cease to be a question.

The first step for having computer vision as a continuous method to enhance and complement the collection data would be to strategize its usage and make it part of the digital plan of the museum. Due to the nature of these algorithms and their rapid evolution, these initiatives would need to be undertaken on an ongoing basis and not as one-off projects. The normalization of the usage of computer vision would require a continuous assessment of the tools and quality assurance through internal processes or with the help of online communities via crowdsourcing initiatives.

An immediate consequence of implementing the results of computer vision applications is the impact on the user experience. Some options discussed with the interviewees are the inclusion of these data outputs in the backend of the website search engine, the incorporation of filters to browse the online collection. The potential inaccuracy and subjectivity of the results in, for instance, tag creation or entity identification, call for openness and transparency about the data sources and their limitations when those are presented in the user interface. For instance, the Science Museum Group has added a comment in their API to highlight that the computer-generated tags are experimental.

The ethical component of the process will require critical decisions about the approval or not of tags, labels and other metadata generated by machines. Museums would need to develop an 'ethics playbook' to specify the criteria for approval, the accepted taxonomy and other actions to mitigate the bias. As one of the participants raised:

> How can we, as a museum, craft AI-based collection data initiatives that extend our approach to teaching with collections while not perpetuating existing biases in our datasets?

6.2. Collaborations across the museum sector

Some of the common themes and challenges in these case studies included the ethical implications of working with third-party algorithms and the lack of sufficient training data in the collection. A potential solution proposed during the project was the collaboration across museums in different aspects and stages of the application of AI tools. Firstly, potentially a group of large museums with the necessary resources could collaborate to develop an open-source algorithm with the documentation and guidelines to be reused by other museums. Secondly, related to the previous point, museums could work together to train an algorithm with multiple collections adding then more records to the training dataset. This approach could be undertaken using large collections or focusing on specific variables such as technique, medium, historic period or subject to increase the accuracy of the results. Moreover, the group discussed during the workshops the need to more generally make heritage content available to creators of image recognition systems along with associated data to help improve their systems. Finally, the collaboration in the museum sector could involve the production of guidelines on computer vision in the existing ethics codes and policies provided by museums associations and other professional bodies.

7. Conclusion and next steps

As Trevor Panglen's work shows us, computer vision technologies are not without their flaws, many of these flaws are echoed in museum collections, through decisions made by humans. Museum collections are full of subjective decisions that determine how art and objects are collected, categorized and displayed, and as such, in many ways, museums are well placed to develop strategies and processes that can quality assure the data produced by computer vision technologies. Documentation of process and peer review of projects that engage with computer vision technologies will be key to developing this new mode of museum practice. In many ways, it could be argued that as computer vision moves from the pilot project to standard practice, it will become as ubiquitous as, or indeed invisible within industry-standard collection management databases.

This paper shows that we are at a critical juncture in the history of collection management within museums, and museum professionals are faced with a choice. Blindly adopt what is possible, and what is affordable from third-party vendors, or engage in a public, critical, yet constructive process of development that can support ethically robust technology development. If museums push back, and ask big questions, around accountability, authenticity, representation, diversity and unintended consequences, they can fulfill their wider mission as social purpose institutions. Museums are sites of cultural development, they are active spaces, and as a result, they can empower and activate a wider public dialogue about these technologies, and help to harness the potential benefits they could offer. As such, whilst this paper provides an important foundation from which to develop the conversation, and practices around machine vision, in many ways, it poses more questions, than it provides answers. We propose a need for further research that focuses on user testing to provide insights into how visitors engage with computer-generated interpretation of museum collections.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Acknowledgments

The authors would like to thank the interviewees for their time and contributions to this research. We would also like to acknowledge the other participants in the Museums + AI network whose discussions during the project helped to examine the topics presented in this paper.

Funding

This work was supported by Arts and Humanities Research Council: [Grant Number AH/S012516/1].

Notes on contributors

Dr. Elena Villaespesa works as an Assistant Professor at the School of Information, Pratt Institute. Her research and teaching areas of interest include digital strategy, data analytics and user experience research and evaluation applied to the museum sector. She is International Co-Investigator on the Museums and AI Network (funded by the AHRC). She completed a Ph.D. in Digital Heritage at the School of Museum Studies, University of Leicester (United Kingdom) and an MA in Arts Management, Universidad Carlos III, Madrid (Spain). She has previously worked as a Digital Analyst at the Metropolitan Museum of Art and Tate.

Dr. Oonagh Murphy is a Lecturer in Arts Management at Goldsmiths, University of London. As an arts manager, writer and lecturer her research has taken her around the world to explore international best practice on the scalability of emerging technologies for cultural organisations. She is Principal Investigator on the Museums and AI Network (funded by the AHRC). She completed a Ph.D. in Museum Studies and Digital Culture, and a Masters of Research in Heritage and Conflict at Ulster University, Northern Ireland. A recipient of a Winston Churchill Memorial Trust Fellowship, she is a Fellow of the Higher Education Academy.

ORCID

Elena Villaespesa http://orcid.org/0000-0002-7790-5495
Oonagh Murphy http://orcid.org/0000-0002-5095-8861

References

"Auckland Art Gallery's New Chatbot Demonstrates Art-Ificial Intelligence to Give New Access to 17,000 Artworks | Auckland Art Gallery." n.d. Accessed 10 March 2020. https://www.aucklandartgallery.com/page/auckland-art-gallerys-new-chatbot-demonstrates-art-ificial-intelligence-to-give-new-access-to-17000-artworks.

Bernstein, Shelley. 2017. "Computer Vision so Good." *Medium: The Barnes Foundation (blog)*, August 23. https://medium.com/barnes-foundation/computer-vision-so-good-3f4162a35d3f.

Choi, Jennie. 2019. "Exploring Art with Open Access and AI: What's Next?" *The Metropolitan Museum of Art* (blog), June 11. https://www.metmuseum.org/blogs/now-at-the-met/2019/met-microsoft-mit-exploring-art-open-access-ai-whats-next.

Choi, Jennie. 2020. "Engaging the Data Science Community with Met Open Access API." *The Metropolitan Museum of Art* (blog), February 13. https://www.metmuseum.org/blogs/collection-insights/2020/met-api-computer-learning.

Cook, Sarah, Alona Pardo, Trevor Paglen, and Barbican Art Gallery. 2019. *Trevor Paglen: From Apple to Anomaly*.

Corbin, Juliet, and Anselm Strauss. 2014. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. 4th ed. Thousand Oaks: Sage Publications, Inc. https://us.sagepub.com/en-us/nam/basics-of-qualitative-research/book235578.

Crawford, Kate, and Trevor Paglen. 2019. "Excavating AI: The Politics of Images in Machine Learning Training Sets." *Excavating AI*. Accessed 14 March 2020. https://www.excavating.ai/

Davies, Roy E. 2012. *Computer and Machine Vision: Theory, Algorithms, Practicalities*. 4th ed. Oxford: Elsevier.

Davis, Hannah. 2020. "A Dataset Is a Worldview - Towards Data Science." *Medium*, March 5. https://towardsdatascience.com/a-dataset-is-a-worldview-5328216dd44d.

French, Arianna, and Elena Villaespesa. 2019. "AI, Visitor Experience, and Museum Operations: A Closer Look at the Possible." In *Humanizing the Digital: Upproceedings from the MCN 2018 Conference*, edited by Susan Anderson, Isabella Bruno, Hannah Hethmon, Seema D Rao, Ed Rodley, and Rachel Ropeik, 101–113. Museums Computer Network.

"Google Arts & Culture Experiments – X Degrees of Separation." *Google Arts & Culture*. 2018. https://artsexperiments.withgoogle.com/xdegrees.

Harvard Art Museums. 2020. "Harvard Art Museums API Documentation." February 21. https://github.com/harvardartmuseums/api-docs.

Harvard Art Museums. n.d.a. "AI Explorer." Accessed 17 February 2020a. https://ai.harvardartmuseums.org/.

Harvard Art Museums. n.d.b. "AI Explorer." Accessed 17 February n.d. "IIIF Explorer." Accessed 17 February 2020b. http://iiif-explorer.herokuapp.com/.

"Hirshhorn Eye - Hirshhorn Museum and Sculpture Garden | Smithsonian." n.d. Accessed 10 March 2020. https://hirshhorn.si.edu/hi/.

"IMet Collection 2019 - FGVC6." 2019. *Kaggle*. Accessed 8 September 2020. https://kaggle.com/c/imet-2019-fgvc6.

Kessler, Maria. 2019. "The Met x Microsoft x MIT: A Closer Look at the Collaboration." *The Metropolitan Museum of Art* (blog), February 21. https://www.metmuseum.org/blogs/now-at-the-met/2019/met-microsoft-mit-reveal-event-video.

Kong, Cathleen. 2020. *Computerized Visual Analysis for Classifying Chinese Paintings by Dynasty*. Princeton University. http://arks.princeton.edu/ark:/88435/dsp010z709048z.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–444. https://doi.org/10.1038/nature14539.

Lih, Andrew. 2019. "Combining AI and Human Judgment to Build Knowledge about Art on a Global Scale." *The Metropolitan Museum of Art* (blog), April 3. https://www.metmuseum.org/blogs/now-at-the-met/2019/wikipedia-art-and-ai.

Majd, Mahshid, and Reza Safabakhsh. 2017. "*Impact of Machine Learning on Improvement of User Experience in Museums*." 2017 *Artificial Intelligence and Signal Processing Conference (AISP)*, Shiraz: IEEE: 195–200. https://doi.org/10.1109/AISP.2017.8324080.

McKim, Joel. 2019. "Trevor Paglen Trains His Sights on the Rise of Machine Vision." *Apollo Magazine*, October 17. https://www.apollo-magazine.com/trevor-paglen-machine-vision/.

Mensink, Thomas, and Jan van Gemert. 2014. "*The Rijksmuseum Challenge: Museum-Centered Visual Recognition*." *Proceedings of International Conference on Multimedia Retrieval - ICMR '14*, Glasgow, UK: ACM Press: 451–454. https://doi.org/10.1145/2578726.2578791.

"The Met - Art Explorer." n.d. Accessed 8 September 2020. https://art-explorer.azurewebsites.net/search.

"The Met x Microsoft x MIT." 2018. *The Metropolitan Museum of Art*. https://www.metmuseum.org/about-the-met/policies-and-documents/open-access/met-microsoft-mit.

The Metropolitan Museum of Art. 2020. "The Tagging Initiative [online]." *The Metropolitan Museum of Art*. Accessed 2 April 2020. https://www.metmuseum.org/about-the-met/policies-and-documents/open-access/tagging-initiative.

Murphy, Oonagh, and Elena Villaespesa. 2020. *AI: A Museum Planning Toolkit*. London: Goldsmiths, University of London. https://themuseumsainetwork.files.wordpress.com/2020/02/20190317_museums-and-ai-toolkit_rl_web.pdf.

The Museum of Modern Art (MoMA). 2018. "Identifying Art through Machine Learning." *The Museum of Modern Art*. https://www.moma.org/calendar/exhibitions/history/identifying-art.

"The Museums + AI Network." 2019. Accessed 17 February 2020. https://themuseumsai.network/.

Pim, Harrison. 2018. "A New Kind of Image Search." *Medium: Wellcome Collection* (blog), October 18. https://stacks.wellcomecollection.org/a-new-kind-of-image-search-5870c2cdadcb.

Saleh, Babak, Kanako Abe, Ravneet Singh Arora, and Ahmed Elgammal. 2014. "Toward Automated Discovery of Artistic Influence." *ArXiv:1408.3218 [Cs]*, August. http://arxiv.org/abs/1408.3218.

Stack, John. 2020. "Computer Vision and the Science Museum Group Collection." *Science Museum Group Digital Lab Medium* (blog), April 4. https://lab.sciencemuseum.org.uk/computer-vision-and-the-science-museum-group-collection-a6c20efb0ac9.

Szeliski, Richard. 2010. *Computer Vision: Algorithms and Applications*. London: Springer.

Tate. 2016a. "IK Prize 2016: Recognition – Exhibition at Tate Britain | Tate." n.d. Tate. Accessed 10 March 2020. https://www.tate.org.uk/whats-on/tate-britain/exhibition/ik-prize-2016-recognition.

Tate. 2016b. "Recognition." Accessed 10 March 2020. http://recognition.tate.org.uk/#info.

Veirum, Niels Einar, and Mogens Fiil Christensen. 2011. "If It's Not on the Net It Doesn't Exist." *Museum Management and Curatorship* 26 (1): 3–9. doi:10.1080/09647775.2011.540123.

Zhang, Chenyang, Christine Kaeser-Chen, Grace Vesom, Jennie Choi, Maria Kessler, and Serge Belongie. 2019. "The IMet Collection 2019 Challenge Dataset." *Cornell University - ArXiv.Org*, June. http://arxiv.org/abs/1906.00901.