International Journal for

Digital Art History

October 2016 - Issue #2

ISSN 2363-5401

Visualizing Big Image Data

Graphentis Verlag

Graphentis
Verlag

International Journal for

Digital Art History

Publisher
Graphentis Verlag e.K., Munich, Germany
www.graphentis.de

Issue 2, October 2016
ISSN: 2363-5401 (online)
ISSN: 2363-5398 (print)
www.dah-journal.org

Partner Insitutions

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

TUM Technische Universität München

Published in cooperation with

arthistoricum.net SPECIALISED INFORMATION SERVICE ART · PHOTOGRAPHY · DESIGN

Title image: Detail from “Gugelmann Galaxy” by Mathias Bernhard, 2016.

Disclaimer
The Publisher and Editors cannot be held responsible for errors or any consequences arising from the use of information contained in this journal. The views and opinions expressed do not necessarily reflect those of the Publisher and Editors, neither does the publication of advertisements constitute any endorsement by the Publisher and Editors of the products advertised.
Contributions are welcome. Please check our website for submission details: www.dah-journal.org. We accept no liability for unsolicited manuscripts and pictures. Authors agree to the terms and conditions and assure that the submission is free of third parties rights. The author grants a royalty-free and irrevocable right to freely publish documents submitted.

ISBN 978-3-942819-11-4

ISSN: 2363-5398 (print version)
ISSN: 2363-5401 (electronic version)

Contents

Contents

Editorial

Big Image Data as new research opportunity in Art History

Harald Klinke, Liska Surkemper

First of all, we would like to thank the DAH-community for the high level of interest and for the overwhelmingly positive responses to the first issue and to the call for contributions. This interest in the DAH-Journal shows that Digital Art History has come of age. After tweeting that the first issue was available and free to download, the tweet was shared exactly 99 times reaching over 30,000 viewers. The journal has been downloaded more than 8,000 times and over 500 readers from all over the word have registered for the newsletter – both numbers are still increasing. What a fantastic start for the DAH-Journal and yet another step forward toward strengthening the community of likeminded scholars.

In the first issue, we discussed "What is Digital Art History?" and learned that it is not new but has perhaps been slowly evolving since the 1980's. That seems to be changing now due to new technology, thus emerging new fields of research in DAH and increasing institutional support. Many have noticed, after Lev Manovich's article, that Digital Art History is largely being driven by data. This ought to be a new paradigm in Art History that spurs new questions: What does working with digital images mean? What possibilities do large data sets present? What are the challenges ahead? And what are the opportunities for art-historical research methods?

While digital methods open up new opportunities, at the same time they question the objectives of our discipline: What does connoisseurship mean in the digital age? What is a master narrative today? What exactly is the *digital* in historical research on art? To address these questions, this current issue focuses on Big Image Data (BID) and the changes in the catalogue of art-historical methodology that come with it.

Big Data and BID open up tantalizing new vistas to the art historian. BID as a sub-category or—better yet—an extension of Big Data affords the possibility of processing and analyzing massive amounts of visual material using computational methods. Among other things the computer creates meta-visualizations comprised of an image data corpus.

Unfortunately, BID is perceived in some quarters as a threat—lack of familiarity with digital technology awakens fears that art historians as traditionally conceived of them will be supplanted by screen-bound technophiles stripped of all aesthetic sensibility. But while some may say that theory and methodological approaches past and present will become wholly obsolete thanks to Big Data, we beg to differ. We believe that as well as expanded access to image banks, BID will also provide art historians with a whole new set of analytic tools, adding new tonal range to our discipline without discarding any of the traditional art historical methods.

This is the school of thought powerfully propagated by Heinrich Klotz—founder of the ZKM (*Center of Art and Media*) and the HfG Karlsruhe (*University of Arts and Design*) in the early 1990s—who encouraged everyone to embrace new technologies in art: "I always say, we do not discard the grand piano, only because we have the synthesizer. We take the synthesizer *and* the grand piano, video *and* painting!"[1] In other words, we do not throw out iconology and iconography just because we have BID. Rather BID should enhance these traditional skills.

Heinrich Klotz saw the need for institutions which would house artists, philosophers, architects and art historians (like Jeffrey Shaw, Boris Groys, and Hans Belting), who would systematically analyze and discuss the opportunities offered by new media in art even—or *because*—this was still unknown territory. In this issue of DAH-Journal, we pursue the same line of thought by encouraging authors to write about BID.

We are living in a world in which we communicate visually like never before and Art History has the experts to analyze and reflect upon such visual artifacts. The digital's predicted impact on art-historical methodology and its repeated questioning of Art History can be viewed positively, because it is helping bring our discipline into the 21st century.

Of course, one could ask if the projects and research approaches presented here, which deal with large amounts of image data, can realistically be labeled BID. Big Data means more than just a large quantity of data; it is the *processing* aspect that differentiates Big Data from mere data per se. One way of defining it is by the *three V's*: "Big data is high *volume*, high *velocity*, and/or high *variety* information assets that require new forms of processing to enable

Liska Surkemper, Harald Klinke

(Photo: Janusch Tschech. Artwork "Nachschub": Li-Wen Kuo) enhanced decision making, insight discovery and process optimization.”[2] The problem of extending the Big Data definition to BID *and* characterizing art-historical projects as such is, that processing image data is a much more complex task than analyzing texts with computational methods.

Furthermore, it remains unclear how art historical data can be operationalized for digital research purposes; issues surrounding access rights, technical problems to do with digitization of images and quality of scans are just some of the obstacles still to be overcome. Computer Vision, i.e. making the computer ‘see’ in a sense that it can differentiate, compare, and thus categorize images, is one of the biggest challenges being faced right now – by the scientific community as well as for software giants, such as Google. As Richard C. Johnson from Cornell stated in the interview in our first issue: “The time commitment to gain access to scientific quality data has proven formidable. It remains a high barrier.”[3]

Nonetheless, there are already a variety of innovative approaches allowing researchers to handle millions of images with unprecedented facility. And although BID is still in its infancy, we see this not as a shortcoming but as an opportunity to critically engage with the exciting developments this emerging field has to offer. The articles and case studies gathered here illustrate the rapid development the field has seen by mapping the progress of various applications of BID to real world praxis. They also suggest ways in which it will and should – or should not – develop in the future.

Harald Klinke’s introductory article “Big Image Data within the Big Picture of Art History” drills into the nature of the digital image to the information layer that lies beneath it. Secondly, it poses the important question of what art historians can extract from BID in information terms. In what ways can the digital code behind the forms, colors, etc., which one sees as a recompiled image, increase an understanding of and ability to manipulate an image?

Would you like to be part of the journey? Welcome. You can contribute in many ways: as author to share your findings, as peer reviewer to enhance the quality of the journal, as copy editor to improve its language, as advertiser or sponsor to support its organization or in many other ways. This is a journal for the community by the community. We can’t wait to hear your thoughts. Contact us at editors@dah-journal.org

Editorial

Maximilian Schich, the featured author of this issue, begins with a bang in his article “Figuring out Art History”, introducing the reader to a new definition of art history. He invites us to take on a different perspective on our discipline by introducing a systematic approach which is now opening up through the possibilities of Big Data and BID.

Next, the contributions of Babak Saleh and Mathias Bernhard both analyze and work with digitized corpora. They focus on different collections and approaches to archiving, retrieving and presenting those large multimedia datasets. While Saleh investigates low and high visual features and explains the methodology of metric learning approaches in order to achieve automated classifications, Bernhard takes the reader on a journey into the Gugelmann Galaxy. This latter project allows users to explore a collection of digitized images and texts in an immersive three-dimensional cloud.

Stefka Hristova describes why and how to combine methods of Art History and Digital Humanities, and the knowledge gained by integrating cultural and visual knowledge as well as data science. She demonstrates this method on a case study of Aby Warburg’s *Mnemosyne Atlas* and emphasizes, that the visualization mode used is always imbricated in a complex network. A network that is not only algorithmic but also historic. Matthew D. Lincoln analyses the development of a professionalized and highly centralized printmaking industry in northern Europe during the mid-sixteenth-century by using network analysis.

We also interviewed George Legrady, Professor of Interactive Media at Santa Barbara, who, since his beginnings as an analog photographer became an early adopter of digital technology, integrating and interrogating the use of computational methods in his own work. He has pioneered in the artistic use of Big Data visualizations at the intersection of Technology and the Arts.

This issue also includes two case studies of projects that deal with large data sets and detailed reviews of workshops and summer schools on digital art history to showcase how institutes worldwide are engaging students and researchers interdisciplinarily in our discipline.

At the end, we would like to share some news concerning the publication itself. What we learned from the website analytics within the last months is: 20,61% of our readers are on mobile devices. While a pdf-file is great for print and on a big screen, it is difficult to read on a display on a smart phone. Since we believe, this journal is also an experiment on what an Open Access journal could be for Art History, we are trying something new with issue #2. We are now also publishing in the epub-format. This open standard is great for small screens. We hope you

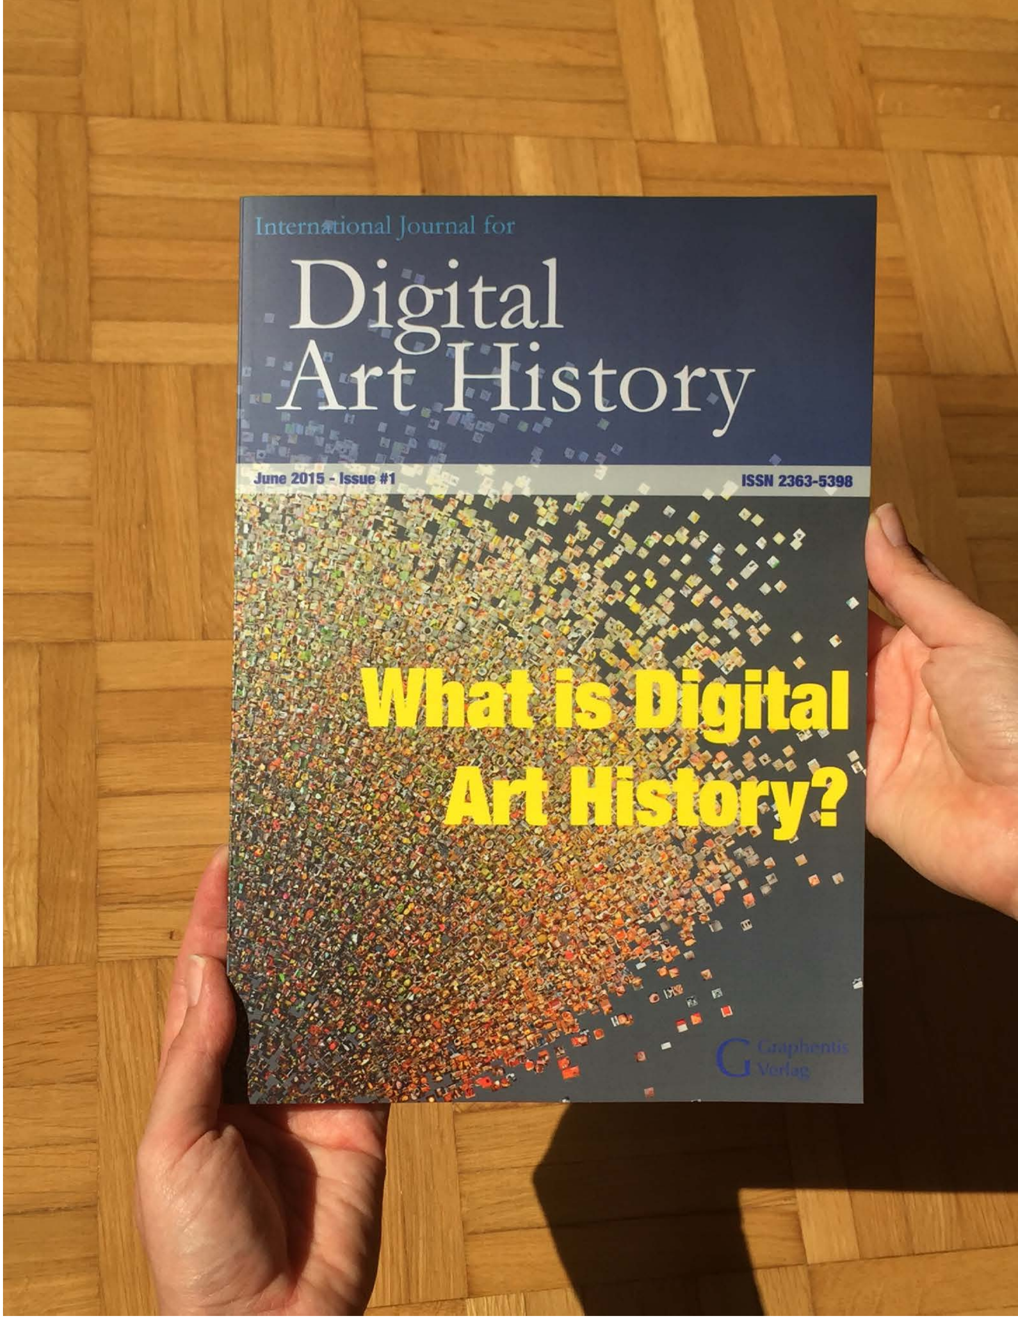

like it. Please, send us your experience; we are looking for the best way to bring you content as comfortable as possible.

Paper is still an essential medium for reading. Thus, we have also published the journal in print. And it looks beautiful. We are proud to say, in the digital age, to weight a book in your hand, flip through the pages and read in high quality print is still an experience.

At the same time, we are excited to see how the digital will change the way we publish, the way we access and analyze research data, and in the end possibly the way we think. On this journey, there are opportunities that can be pursued with heartfelt passion, but it also needs a critical eye to keep track on the direction we follow. What do we want Art History to look like in 10 or 20 years? Now is the time to do the first steps.

Notes

1 Heinrich Klotz, Rektoratsreden (Hamburg: Ausnahme-Verl., 2009), 90 (editor's translation and italics).
2 Douglas Laney, "The Importance of ‚Big Data': A Definition," in: Gartner, 2012 (editors' italics). http://www.gartner.com/resId=2057415.
3 Richard C. Johnson, Park Doing, "On Applying Signal Processing to Computational Art History – an Interview,“ International Journal for Digital Art History, no. 1 (2915), 89.

Harald Klinke has a Ph.D. in art history and a Master of Science in Information Systems. Currently he is Assistant Professor at the Ludwig Maximilian University, Munich, and responsible for the doctoral program "Digital Art History". He conducts research on visual communication, digital media, and Big Image Data in art-historical contexts.
From 2008 to 2009, he worked as a Lecturer of Visual Studies (Bildwissenschaft) at the Art History Department of the University of Göttingen. From 2009 to 2010, he conducted research, supported by a grant from the German Research Foundation (DFG), as a Visiting Scholar at Columbia Uni-versity, New York. He has published books on American history paintings, digital images and art theory as visual epistemology.

Correspondence e-mail: h.klinke@lmu.de

Liska Surkemper is Associate Researcher for architectural and cultural theory at the Technical University Munich. She conducts research on visual epistemology and the interrelationship of pictures, architecture and economy.
From 2010 to 2014, she was a researcher and lecturer at the Department of Art Research and Media Philosophy at the University of Arts and Design Karlsruhe. She was also coordinator for the project "Memory of Scientific Knowledge and Artistic Approaches", which was supported by the German Federal Ministry of Education and Research (BMBF). Together with computer scientists, designers and arts scholars she helped develop the web application "Presenter" (http://presenter.hfg-karlsruhe.de): a tool for visualizing, sharing and archiving scientific and artistic knowledge.

Correspondence e-mail: liska.surkemper@tum.de

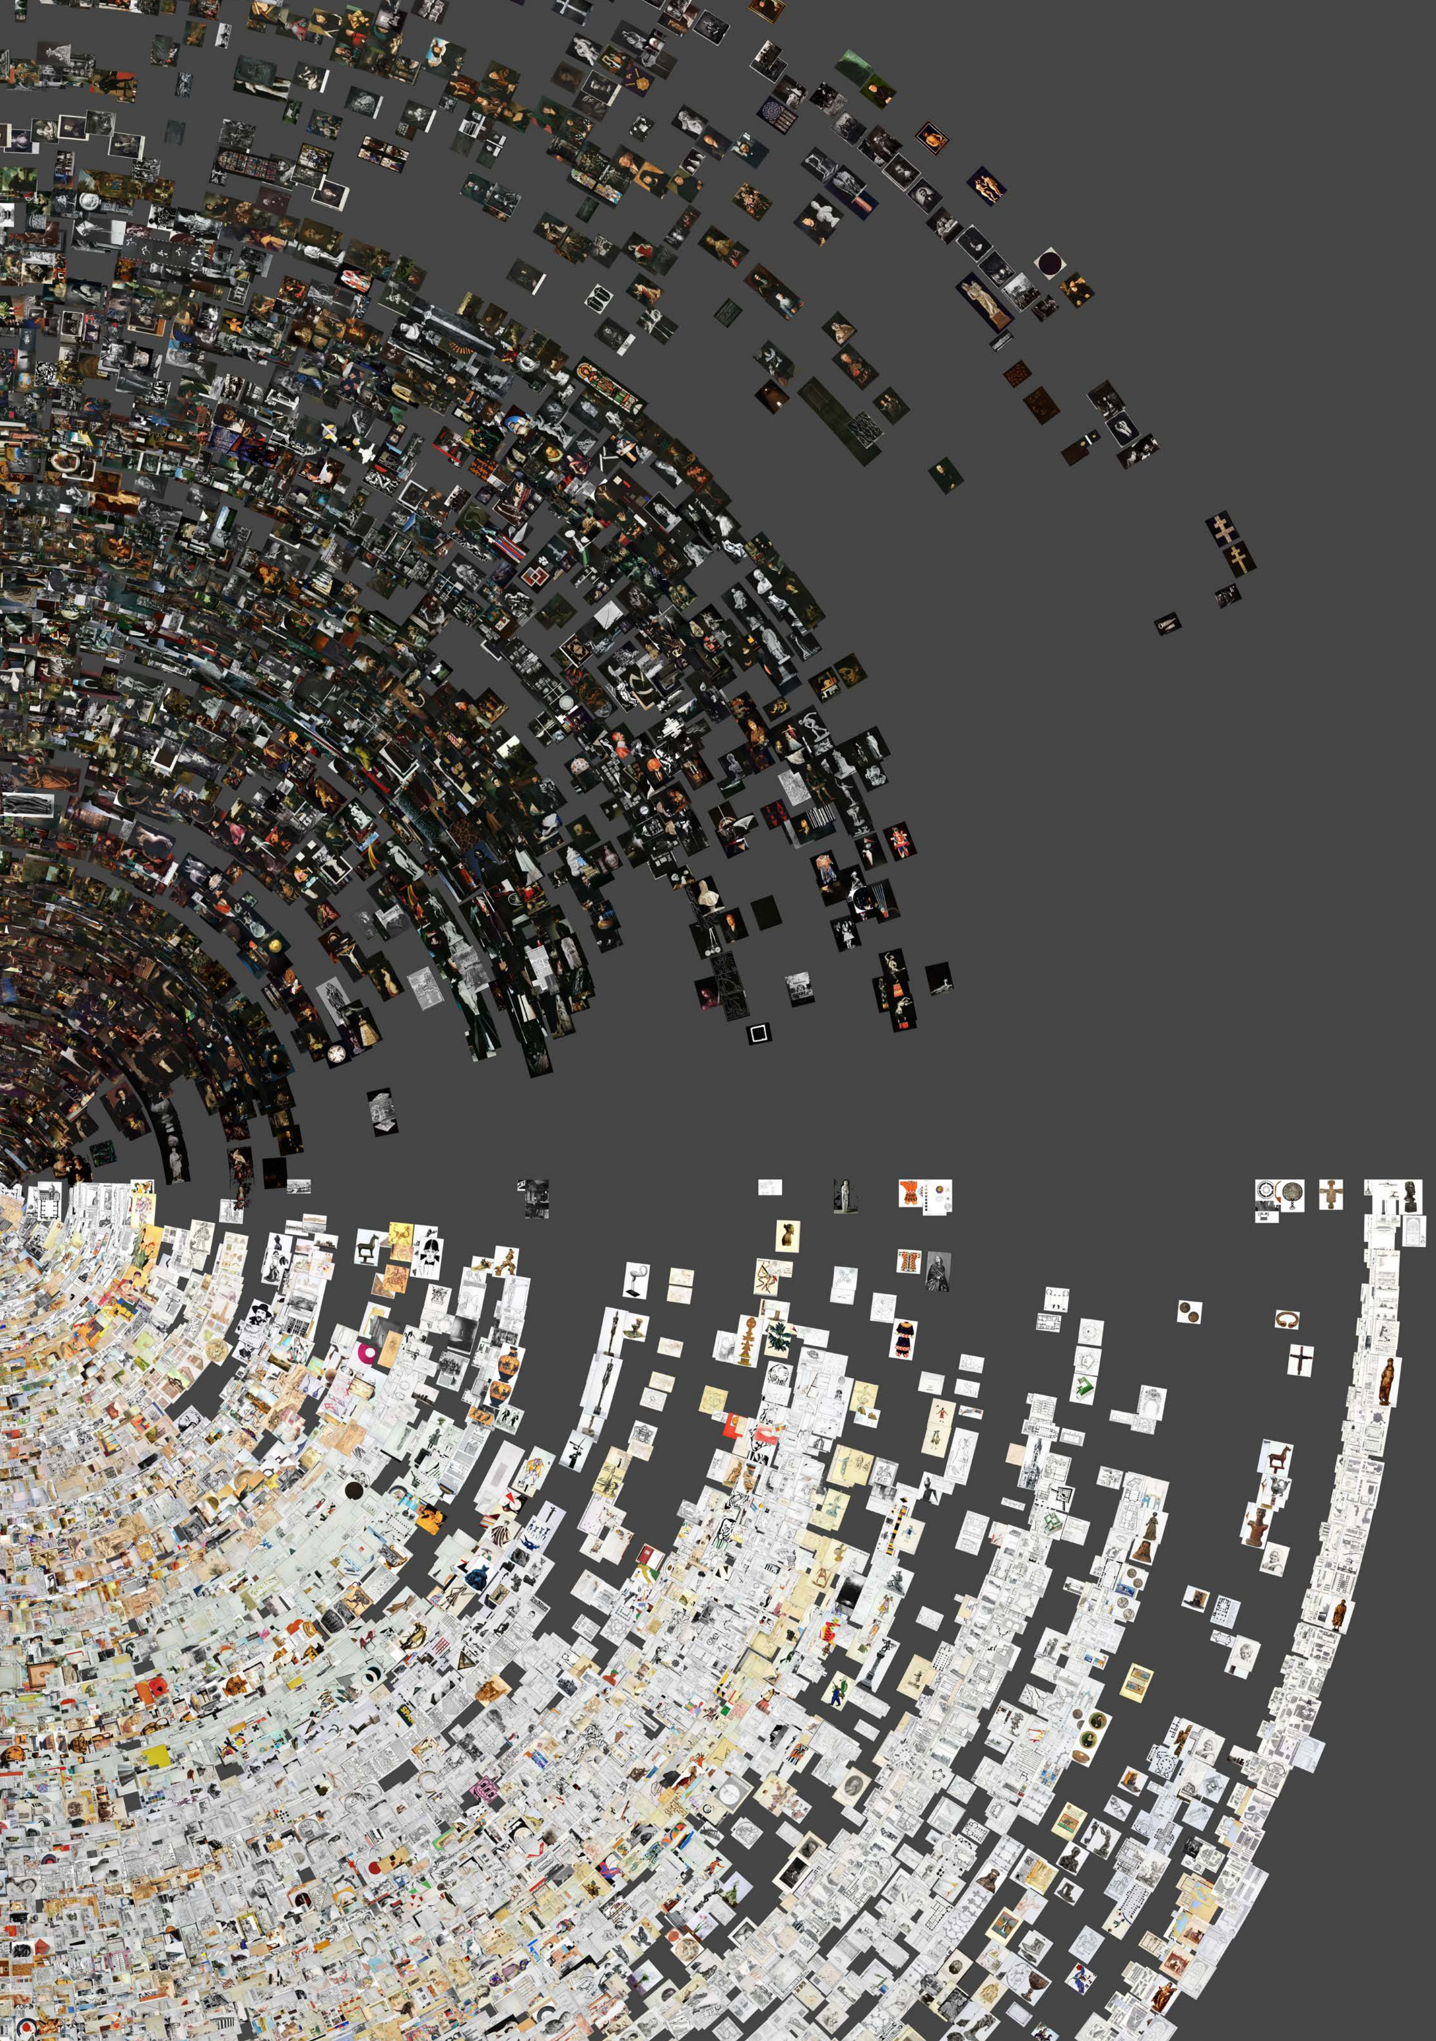

Big Image Data within the Big Picture of Art History

Harald Klinke

Abstract: The use of the computer in Art History is changing the approach towards our objects of research. Now, we are able to compute more images than a human can see in a lifetime. That, in turn, calls for a new definition of the role of the researcher and the tools being used. The access to large amounts of visual data stands in a tradition of conventional methods of Art History, but also augments them with quantity. This article proposes a theoretical model on which to build an understanding of the meta image with which we interactively derive our conclusions.

Keywords: Big Data, Digital Humanities, Distant Viewing, Information, Media, Meta Image, Methodology, Qualitative/quantiative Research, Semantic Gap

The digitalization of Art History is creating enormous opportunities. Among the most striking is that we now have the opportunity to not only look at a single work of art—or the comparison of two—but that we are also able to compare more pictures than a human can look at in a lifetime. This concept is called Big Image Data (BID).

Taking a look at our discipline's methodology, the first systematic approach that comes to mind is Erwin Panofsky's iconography and iconology.[1] According to him, and in contrast to all other disciplines in the Humanities, art historians take the visual record seriously and derive knowledge from it by reconstructing an artwork's meaning layer by layer. The goal here is to embed the individual object into its cultural and historic context. Art History is, thus, an empirical science that gathers and interprets visual data. This is one reason why Art History is, in German-speaking countries, also called "Bildwissenschaft," the science of pictures. It should be noted that Panofsky knew that starting with the visual phenomenon does not allow the exclusion of any written records. It is only natural that scholars in front of a work of art will acquire knowledge from books and compare their findings with those of other scholars, and then add to this shared knowledge through future publications.

Figure 1: About 60.000 images taken from the Artemis database (LMU Munich) sorted by average brightness using ImagePlot developed by the Software Studies Initiative

Another art-historical method comes to mind: the side-by-side comparison of pictures as established by Heinrich Wölfflin.[2] What has been done with originals and painted or printed copies before the advent of slide projections turned into a central method in our discipline through the use of double slide projectors. Comparisons make clear the differences as well as the similarities of the studied objects. The findings can then be discussed and converted into insights, such as the question of the object being an original or a copy, questions of dating (*Datierung*), or authorship attribution (*Handscheidung*).[3] A slide library containing hundreds of thousands of high resolution color images is a powerful means of navigation through the history of art.

Both methods show that no other discipline looks at pictures as systematically as Art History. They also show that individual scholars and their visual experiences acquired over the years are at the center of this kind of research. In both cases, researchers need a mental archive of images in order to identify figures by their attributes and to select objects for comparison. Would it not be great to go a step further and use computers to be able to consult a much greater variety of artworks, to use algorithms to recommend further works that might be of interest to the researcher? Would it not be great to let the computer do the work of sifting through a great number of artworks in order to deliver meta information on an enormous corpus of images?

This capability is exactly what the Digital Humanities promise. As the Natural Sciences have used computers to yield unprecedented insights that were not possible before, such as the sequencing and analysis of the human genome, the Humanities are able to make the computer a tool for their research, too. Ever since Roberto Busa created a lemmatization of the works of Thomas Aquinas using IBM computers in 1951,[4] text-based Humanities have developed many software tools to obtain an overview of text corpora employing statistical methods, to uncover semantic structures via text mining, and to develop new insights into language using Corpus Linguistics. The visual sciences, the foremost being Art History, started to integrate computational methods in the early 1980s,[5] but progressed slowly for several reasons.[6] The most important reason is that images are much harder for a computer to process. This might seem to be an odd statement, especially since digital images consist of nothing beyond a string of information units called bits like all other digital media. But this matrix of pixels does not “mean” anything to the computer at first.

While one string of ASCII characters[7] in a text can be compared to another (search), they can be summed up in a corpus (frequency distribution) or become the basis of statistical calculations (text mining). A string of pixels on a display is nothing more than a sequence of brightness values on the visible light spectrum (red, green and blue) at first. However, we are able to calculate and use these color values as the criterion to sort images[8] in a process involving these "low level features" of images (see figure 1). This process, however, is limited in its epistemic potential.[9] In contrast, "high level features" are what those pixels represent; the content of the images. Between both of these lies the "semantic gap," where high level features are perceived by humans with ease, whereas the computer still struggles. Information Science—Computer Vision in particular—is working on this problem via pattern recognition, deep learning, and other approaches, but there is still much research to do.[10] As research is carried out, it will increasingly provide the means to work with images that Art History can use for its own purposes.

At this point, we have to distinguish two concepts. Reading a book of literature and writing about it has been called "close reading." Having all works of an author in a digital format and comparing them to those of another author by means of statistical analysis is an example of "distant reading."[11] It offers the possibility to use computers to work with data in order to analyze a greater number of books in a matter of seconds than an individual could hope to read in a lifetime. These concepts can also be applied as standard methods of study for Art History, where "close viewing" describes the study of individual reproductions of artwork. In the same vein, "distant viewing" describes what algorithms are increasingly able to offer: examining an infinite number of images at once and deriving meta information from that corpus of images.[12] Dealing with a massive amount of data is a tremendous challenge for Art History. The good news is that there is no need to start from scratch. Other disciplines have dealt with images before: medical applications deal with images, such as high quality x-rays. Biology has developed software to handle pictures produced by microscopy in order to quickly discover clusters or count cells. Astronomy, Geology, and other similar disciplines incorporated such tools into their daily practices long ago. In particular, one field of Information Science has advanced quickly of late: Big Data.

Big Data has been a buzzword for a few years now. In 2012, it was identified as the biggest trend by Bitkom—the German IT Trade Association—and in 2014, Big Data was the theme of the CeBit trade fair. The reason for this recognition is that it promises unprecedented insights into social behavior and new possibilities for business models. At the same time, it stokes fears of surveillance and espionage. Big Data seems to be defined as the collecting and analyzing of data, with the ultimate result and goal to make money; Facebook, Apple, and Google come to mind when discussing Big Data. With data, companies not only know what we do, but also things that we *might* do.[13] Some say that data is the new raw material.[14] And the massive amount of data is not the problem, but the solution.

Big Data, first of all, is an Information Science concept. Only in combination with business does it become a concept for profit maximization. It has the potential, in combination with science, for "knowledge maximization." Moreover, it can become a fruitful tool in the Digital Humanities. While "data" is generally used here to denote values, such as sensor results, the challenge for disciplines like Art History that use images is figuring out how to gain knowledge from a mass amount of visual data. How can we understand the hidden connections between images on a macro scale? How do we deal with what could be called Big Image Data?

What is a Digital Image?

Large amounts of images can be a treasure trove. An art-historical image database can represent a relevant section of art history for specific research questions. But before we are able to unearth such a treasure, we first have to understand what kind of data digital images are. If we want the computer to help us generate art-historical knowledge, we have to go a bit further and ask what an image is, generally speaking.

"What is an image?" was the central question posed in Gottfried Boehm's 1994 publication that spurred Visual Studies. His thoughts were based on the assumption that we are living in a time when we use images in a quantity unprecedented in human history (and that is truer than ever today, 22 years later). Since we have sciences for language, there must also be a science for images (*Bildwissenschaft*).[15] So again, what is an image?[16] An initial approach to answering this question could be to name the objects that we include under the term "image." The history of art knows many such objects, and paintings are only one of them. Photography has arguably been the most important revolution for images. One can incorporate not only artistic productions, but also the broad field of amateur photography. In addition, technical images have to be added because they are mainly used in the natural sciences. Effigies, wax figures, and other such items are also undoubtedly understood as images.[17] Finally, so-called internal or mental images are also a part of the Visual Studies, even if they are of a fundamentally different nature than the aforementioned external images. As a matter of fact, in the English language we sometime prefer to call external images "pictures."[18] The two are closely tied together, as pictures cannot be sufficiently explained without images.

In addition, the term "picture" refers to another important concept. The conventional panel painting has, for centuries, dominated European art history. Following its example, art theories have been developing since the Renaissance. But only with the invention of photography—and with approaches to make them theoretically comprehensible—a new concept has become necessary: the concept of the medium. The term "medium" here denotes the carrier of an image that makes the visual phenomenon possible by its physical configuration.[19] The term is necessary in photography to make clear the relationship between originals and copies. A photograph allows for the production of a series of positives from one negative. Unlike reproduction graphics, the first and the hundredth copy do not differ appreciably. This has been made possible by using a specific medium—chemical photography, in this case. The distinction between image and medium makes clear what is new about photography and its demarcation from painting. A picture became not only something that could represent something else (a distinction between the picture and what is pictured, see figure 2), but it could also appear with different characteristics in its production, distribution, and reception. It is noteworthy that after nearly 100 years of image-making practices with photography, these crucial concepts were discussed for the first time in 1936 by Walter Benjamin.[20] This opened up the opportunity to make the new image phenomenon tangible; make us able to conceive a history of visual media and media theory.[21]

Today, we are living in a time where revolutionary changes in the hegemony of visual media are being made again. A new visual medium has joined its peers and, similar to photography, challenges other media in their artistic standards and continues to drive the democratization of image production and lets the flood of images skyrocket. This new visual medium is the digital image.

The advent of the digital image—just think of the spread of digital photography—challenges the older media at least as radically as photography did in the 19th century. It represents a new challenge to conceptual tangibility, just as photography did. At the same time, it is itself the driving force to develop new ideas about what images are. Just as the photograph required the concept of media, the digital image requires new terminology in order to make its unique ontology understood, as well. The difficulty in making digital

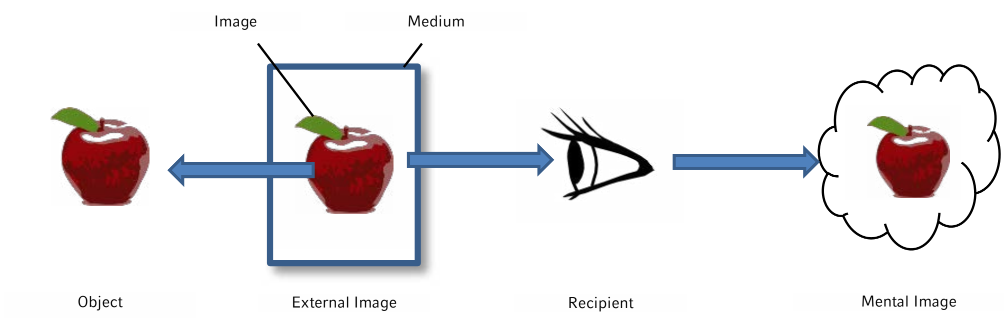

Figure 2: Scheme of image and medium

images understandable comes from the fact that the concepts developed for the older image media must now be challenged: What is the image of a digital image? What is the medium of a digital image? The distinction of image, what is being pictured, and the medium no longer seems sufficient. Therefore, the digital image can only be understood when new terms are introduced.

A few specific examples will make it easier to grasp what has been introduced by the digital image. The simplest example of a digital image is probably the digital photograph. It is noteworthy that even especially high-priced digital cameras today imitate the mostly conventional Single Lens Reflex camera (SLR) in design. The film has been replaced by an image sensor, such as a CCD chip. The advantage of this technique is the immediate consideration of the result on a screen, so much so that in compact cameras, the optical viewfinder has largely been phased out in favor of a visual display. The revolution of photography was, at one time, the chemical fixation of the volatile image of the camera obscura. Today is quite similar, where digital image captures can be displayed on a pixel matrix (see figure 3).

Another example of a digital image is when you work in a graphics program, such as Adobe Photoshop, on a computer. The image is visible on the display and can be changed using a graphical user interface that is controlled by mouse and keyboard peripherals. The possible changes to brightness and contrast, to collage or "paint," have replaced the conventional processes used in a darkroom.[22] The "retouched" image can subsequently be sent or shared via email or social networks and made available to other viewers.

A third common form of the digital image is—as opposed to digitized—the computer generated image (CGI) as used in computer games (first-person shooters, for example), movies, or design software. The on-screen visible

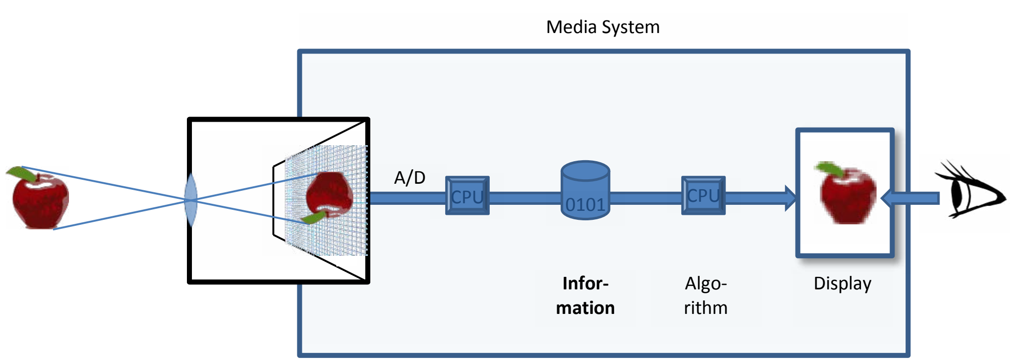

Figure 3: Basics of digital photography

image is now generated based on data and algorithms devised through computational steps. In addition, this allows for manipulation of the image by the recipient (now called the “user”) in real time.

These are all undoubtedly examples of images, and there is something they all share: the visible and the object being pictured. Also, digital images represent something visible by means of the image carrier (the medium). Yet, it is clear that we are dealing with a special picture system here whose unique ontology becomes clear when examining the differences with older picture systems. The “image” on a camera’s memory chip is itself not visible, comparable perhaps to an exposed roll of film in a sealed opaque cartridge, carrying the latent images only. And yet the “latency” of digital images on the chip is different because it is made visible through a different process than chemical photography. In image processing, it is clear that the image can be changed with more ease than performing the same changes in a darkroom. It can be copied without loss and can become visible anywhere in the world almost simultaneously. Finally, computer generated imagery does not need to be in direct proportion to the visible reality; rather, it can be created and modified in real time. This can be done with such quality that the impression of immersion may occur, which was not previously possible with any other visual medium.[23]
If one examines the “image” in the strict sense (what is visible with the help of a medium that has a reference to something), the question arises: what is the medium of the digital image? It is certainly not the display alone. In digital cameras, the memory chip and the files it contains are important elements. When one thinks of Internet applications, it is clear that the display unit and the file do not even need to be in one place. The file is located on a remote server and can be fetched via the network at any time.

Strictly speaking, the CCD chip converts the local brightness values of RGB colors via an analog/digital converter into numbers, which are then compressed using a processor and written to a file (such as the JPEG format) where they can ultimately be presented as on/off states. In reverse, the data is read, decompressed, and outputted as the brightness values of the three primary colors on a display. Any alterations via imaging software takes place only indirectly on the display (the graphical user interface is the interface between users and the data; between man and machine).[24] In fact, the database is changed by a computational process. For example, brightening an image corresponds to the addition of a value to each RGB value of the pixel matrix, darkening an image is a subtraction, and contrast enhancement is actually a stretching of the histogram. The image is finally generated by calculating the numerical values of the individual pixels. In other words, the computer knows no image, only numbers. It does not “see”—it only calculates.[25]

The system, as described, makes it clear that we have to add a new concept to image and medium: the concept of information. The science of digital images must be understood as information theory of images. The individual bits of memory contain the information that must be interpreted by applications to provide meaningful messages. A string of bits, like 01000001, can be interpreted as a decimal (65) or—according to the ASCII table—the letter "A" or a brightness value of a pixel. In this way, not only numbers and letters, but entire pictures, sounds, and myriad pieces of data can be coded in binary, processed, and transported as a signal.

But that is not the image. It is the information that is later turned into an image through hardware and algorithms; in other words, by an overarching system that may be referred to as a medium. This process is called "visualizing." The information can be displayed using a computer screen, two screens, be presented as printout, etc. These are only different forms of visualization of the same information.

The question "where is the image in a digital medium?" can now be answered: it is not in the memory storage device. The information stored by the memory device is required to represent something through the use of the medium. The medium—the system that makes the information visible—can be labeled as a computer, which is designed according to the von-Neumann principles. A von-Neumann-computer consists of controller, arithmetic unit, memory, input unit, and output unit components.[26] Every digital camera, smartphone, and laptop consists of memory, a processor, and a display. The seemingly placeless image, e.g. a picture posted on Facebook, can be retrieved from different places, uploaded, and downloaded. Such images basically consist of a string of bits in a file format, such as JPEG, it is located on an Internet server and is transferred on demand to the local machine. The file is interpreted there by using programs and made visible by the graphics and display hardware as an "image." This entire system can be called the medium, and therefore its task is to visualize the string of bits.

If we have thus defined and distinguished information and media, what then is the "image"? In summation, the distribution of brightness of the red, green, and blue values on a screen's pixel array results in an impression of contiguous areas, which are detected as objects based on one's experience with the real world. In other words, an "image" of something. Consequently, the "image" is the visual phenomenon as perceived by the viewer.[27] More specifically, it is the light projected on the retina, which is then converted into nerve signals and—preprocessed by the metathalamus— perceived in the visual cortex.[28] In short: what is visible is the visual phenomenon. In order to avoid obfuscating the term "image," the following descriptors should be used: information, media, and visual phenomenon.

These terms enable a clear description of the digital images. A digital camera generates information that can be visualized on a display or photo printer. The image processing program changes the information (via a graphical user interface), which is then sent and displayed on devices. To “send a picture” is colloquial and actually means to transmit information signals that are interpreted and displayed on the device as a visual phenomenon recognizable by humans. Computer games generate sequences of bits using data and algorithms. These bits are visualized using the medium and are perceived as visual phenomenon.

These terms can be applied to more than digital images. With conventional panel paintings, there was no need to question the medium because the visible phenomenon was inseparable from the physical medium.[29] Image reproducibility extended the medium’s properties and allowed for the new concept of the new visual medium to form. The term “medium” can then be traced historically back to the paintings: the medium “painting” consists of pigments with a binder (e.g. oil) and applied to a substrate (wood or canvas). Because of this material configuration, a visual phenomenon can be experienced, which creates the impression to refer to objects outside of the picture itself.

Accordingly, in photography, there was no need to question the information because the image information was irrevocably connected with the image carrier. Image data processing systems introduced new complexities to what the information entails. It called for precise descriptions and a more comprehensive understanding of this new imaging system. This concept of information can also be traced back historically. As stated above, the information is hidden behind the visual phenomenon and defines the brightness and color distribution for the output device. Viewed as a whole, these pieces of information coalesce to form a visual impression—specifically, the brightness values of RGB subpixels. The values associated with each cell of the matrix are not permanently fixed, because the brightness on a display can

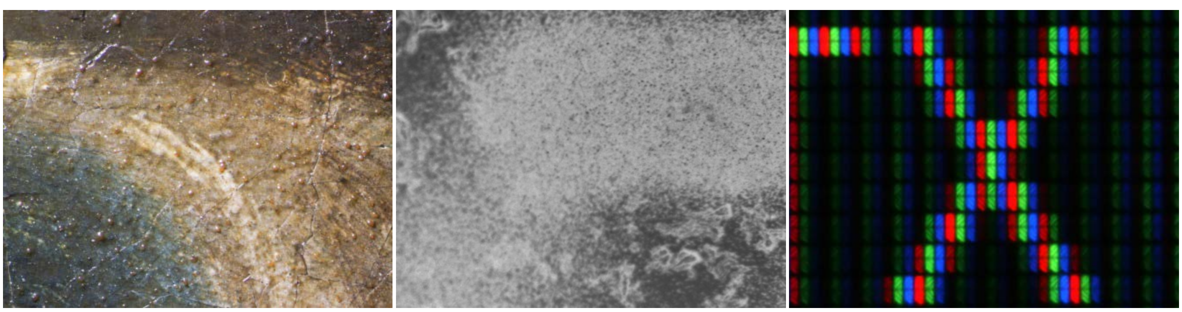

Figure 4: Pigments of a Painting, Silver Nitrate of a Daguerreotype, Pixels of an LCD Display

Source: Micrograph. Light micrograph. Macro shot of an LCD display shift at any point and is controlled by the display's underlying operating system. In a painting, these pieces of color information are also present. However, their values are determined by the artist using a brush in the form of bound pigments and are no longer changeable after the drying process. The information is thus firmly connected to the image carrier. The situation is similar with photography: silver nitrate is changed via light exposure by different degrees. Each silver nitrate molecule at any point is a carrier of brightness information that will convey the overall visual impression (see figure 4). Again, the information is fixed to the medium. Only in the digital imaging medium are these elements separated and act as properties of the new medium that has added new features and options to it. The separation of "information" and "media" makes variability, processing ability, and lossless copying possible—these basic characteristics distinguish the digital image from the conventional image.

The introduction of the concept of "information" in image science has further advantages: the above triad of information, media, and visual phenomenon stands in a hierarchical relationship with one another. The media system requires information to create a visual phenomenon. This can be represented as follows:

3. Visual Phenomenon
2. Media System
1. Information

In computer science, this is called a layer model. Among the best-known layer models is the OSI model, which refers to seven perspectives on technical structures in telecommunications. In it, the lowest layer of the physical layer—copper cable or electromagnetic waves—transmits the bits in the form of signals. The higher layers assure the correct transmission between sender and receiver through a series of transmission protocols.[30] A similar layered model can be applied to visual media.

It should be noted that the first layer consists not only of bits, but also of data formats including compression. The second layer consists of the elements of von-Neumann computers, including memory, on which information is stored, processed, and outputted (e.g. displayed). The third layer consists of the observable phenomenon that is purely optical.

On this foundation a system of human perception can be built: an image is only an image when it is viewed by a recipient. Higher layers describe the steps of assigning meaning. For this, a proven model of additional layers that we can make use of already exists: Erwin Panofsky describes a three-stage general epistemology of visual perception when discussing iconology and iconography. Accordingly, first the objects are named (in his example serving Christian iconography: man, knife), and then the overall image is identified (Saint Bartholomew) and finally placed in its cultural and historical context.[31]

What Panofsky fails to describe is that before the cognitive process—which relies on cultural experience—can begin, the sensation must first be pre-processed to distinguish shades and surfaces in order to identify areas as connected objects. This requisite edge detection, contrast enhancement, and reduction of information already happens in part inside layers of the retina and the metathalamus.[32] This pre-processing is necessary to distinguish a human from an object based on the amount of stimuli in a visual phenomenon and can be added to Panofsky's description as layer zero. In this way, we obtain the following layer model:

3. Iconological Interpretation
2. Iconographic Analysis
1. Pre-iconographic Description
0. Object Recognition

Since the structured image description as shown here is only complete with a description of the image medium, these two layer models can stacked upon each other:

7. Iconological Interpretation
6. Iconographic Analysis
5. Pre-iconographic Description
4. Object Recognition
3. Visual Phenomenon
2. Media System
1. Information

This model describes the picture medium and its reception by the viewer in full and is applicable not only to the digital image, but any image phenomena. It is the relationship between the visual phenomenon caused by a physical configuration and the viewer that constitutes what is called an "image." It is the human being who perceives images and fills them with meaning, and it is the human being who creates images in order to visually communicate ideas. With the digital image, it becomes clear: to talk about images, one must connect four terms into a relationship: recipient, visual phenomenon, media system, and information. In short, an "image" consists of multiple interconnected parts.

What is Big Data?

Having defined the term "image," the question remains: what is "Big Data" and how can we use it, especially Big Image Data (BID), for science? Big Data simply sounds like a lot of data. The image database Prometheus (University of Cologne) contains 1.5 million images,[33] Art Store contains 2 million,[34] Google Cultural Institute contains almost 5 million,[35] and the Artigo project (University of Munich) has generated about 8 million tags so far.[36] A heap of data, however, does not yet represent knowledge. We have to distinguish pure signs, such as bits or bytes that can be interpreted; data, or unstructured facts about certain occurrences in the real world; information, or contextualized, relevant meaning; and knowledge, or understanding through experience (see figure 5).[37] The interpretation of results is still exclusive to humans, but computers can help us with the steps leading to that destination.

Since Big Data promises to acquire new relevant knowledge, it refers to data on one hand, and the analysis of that data on the other. For example, the warehouse chain Target employed a marketing strategy aimed at pregnant women as their target market and tried to address their needs earlier on in their pregnancy than competitors. By analyzing the patterns in a customer's purchase history, a woman's pregnancy and her child's birth date could be predicted, sometimes even before the women knew, it seems.[39] Such patterns are derived from a combination of age, place of residence, and purchase history. This analysis discipline is called "predictive analytics" and shows how knowledge about the future can be generated by using the past and present. An important point should be mentioned here: such predictions are not definitive, rather, they come with a probability. That means the result of such an analysis is not a statement of fact, but is instead a fuzzy projection (often with a notably high probability).

In order to generate knowledge from pure data, Big Data is characterized by three Vs: Volume, Variety, and Velocity.[40] Volume allows for the access to a large set of data, and its heterogeneity allows for the drawing from multiple ends in order to find the item searched for. Its real-time analysis allows for a higher probability of the prognosis. Thus, Big Data means to acquire, store, and analyze large amount of data that are generated quickly and are not always structured, as with SQL databases. Open source solutions like the Hadoop Framework, or database systems such as Cassandra and MongoDB, were all developed to help analyze such data quickly.

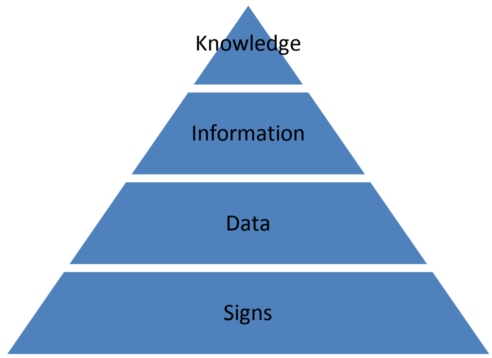

Figure 5: The "Pyramid of Knowledge" as derived from Ackoff (1989)[38]

What is Big Image Data

Big *Image* Data, then, is about a special kind of data: visual data. Analysis of it can lead—just like ordinary Big Data—to statistical key figures, diagrams, overviews, or graphs that show structures such as clusters or outliers.[41] Image data visualization is essentially images about images—meta images, perhaps—since they make a digest of many images understandable. How that digest is calculated determines the epistemic content of the image. Similar to text mining, the process could be called image mining, for which algorithms are currently being developed. It is a projection from many to one and, thus, an information reduction defined by mathematical calculations.[42] In short, they reflect other images. Or, in the words of W.J.T. Mitchell, "Any picture that is used to reflect on the nature of pictures is a metapicture."[43]

A few examples of volume in BID: Apple's iCloud solution has over 782 million users (as of February 2016),[44] and its photo sharing function allows users to upload the last 1,000 photos from a device, such as the iPhone, and sync them with other devices. That means up to 782 billion images can be kept in Apple's data centers, such as the one in Maiden, South Carolina. Shortly after the introduction of this service, the 1,000 image limit was lifted. Starting in 2013, virtually all images can be synchronized. Facebook has over 1 billion users and is said to receive about 350 million photo uploads per day.[45] The sheer number of images is not the treasure here, but the variety: GPS location integration, likes, friends, and other contextual data. The velocity is another treasure and includes the real-time acquisition of images, since they are uploaded almost immediately after they are taken on a mobile device.

Big Data content is often acquired in the "Internet of Things" by means of sensors. Visual data is acquired with visual sensors, such as cameras in smartphones. The number of these images is being generated and accumulated like never before in history. Their analysis can generate enormous insights into societal developments and individual habits. Most of all, visual data speaks to what we see, what we show to others, and what we think the world looks like. Art History is still far away from such repositories, but still has to answer how a *Bildwissenschaft* makes use of image analysis technology. So how can Art History acquire data big enough to derive meaningful knowledge about the history of art? How can we utilize the tools information science is offering for our objectives? Digital Humanities promises that the computer can help us reach our epistemic aims better. The computer can support our conventional methods, and with its help, we are able to develop new digital methods. What does this mean for Art History?

Apple and Facebook are presumably interested in understanding customer behavior in order to match advertisements with customer interests. To achieve such an objective, the collected data requires the proper structure, and the challenge comes in how to analyze the data in order to receive meaningful insights. Art History is also interested in human behavior, but on the other hand, it is also interested in image-making processes (as a *homo pictor*),[46] in specific images like art works, and in networks existing between artwork and artist.[47] This means that not only are the number and structure of data relevant, but also their quality.

Therefore, sometimes there is another important fourth "V" added: veracity. Veracity refers to the accuracy of data in relation to the real world. For Art History, this begs the question: how much can we trust a digital reproduction of an artwork? Fact is, a reproduction always bears a difference to the original. But looking at dozens of variations of depictions of the Sistine Chapel in our database makes us wonder if we are able to comprehend the width of that gap. We need to address the integrity of our data. Hence, we must first ascertain what data needs to be acquired to be used as the base for subsequent analyses. If we build on top of our digitized slide libraries, like our art-historical image databases, we also need to address other quality concerns apart from accurate color depiction. The sufficient richness of visual data remains unanswered: paintings are not flat and are therefore not adequately represented in a JPEG-file, but they have a relief structure. Paintings would be more adequately represented with an additional layer of grazing light or x-rays and a 3D model. Such a database for visual computing could be called *Rich Image Data*.

Once the necessary kind, structure, and amount of data have been acquired, how can it all be used to gain knowledge on a higher level? First of all, art-historical research should not be driven by the promises of information technology, but by its very own epistemic goals. Digital Art History should therefore not ask what to do with all those images, rather start with what we want to find out, and then ask how computers can assist us in this task. As a third step, we should ask what data we need, in what kind of structure, what algorithms we should use in order to analyze the data, and what visualizations should be used to show these results and explore the data.

Data analysis for Art History means the introduction of *quantitative* methods into a field that has used exclusively *qualitative* methods throughout its existence. However, these quantitative methods are not a substitute for conventional methods, but an addition. Results of data analysis can help to answer to our research questions, but may also spark new hypotheses that can be followed using a bouquet of methods that Art History has. For some, Big Data seems to be the end of theory but one could also argue that Data Analytics helps us to see the big picture and create new hypotheses for qualitative research.[48] This way, the computer remains a tool in the hands of the researcher, or as Steve Jobs defined the role of the computer to the human, "a bicycle for the mind."[49]

Figure 6 details how the human system (i.e. the researcher) makes use of the computer (denoted as the media system) in order to access visual data. Such meta images from BID are observed by the viewer just like individual images, except they are a digested representation of the amount of data, can be interactive, and reference on visual data. The referenced visual data relates to the individual physical art-historical object, which the art historian likes to place in a cultural context. The goal of that epistemic process is to form an abstract mental representation of developments in the history of art.[50] In this process, the computer acts as a computational co-processor to the researcher's mind. To quote Mitchell once again: "The metapicture is a piece of moveable cultural apparatus, one which may serve a marginal role as illustrative device or a central role as a kind of summary image, [...] that encapsulates an entire episteme, a theory of knowledge. [...] In their strongest forms, [metapictures] don't merely serve as illustrations to theory; they picture theory."[51]

Today, the computer is able to make low-level features computable and is increasingly able to compute high-level features. That is particularly fruitful for Art History when applied to large amounts of images in order to pre-process them, visually examine them in order to create a distant view, to explore them, and to find them. The processing of low- and high-level features usually takes place in the human retina, metathalamus, and visual cortex. With the ability of the computer to process these plain pieces of data into meaningful information while generating meta images, these processes sink on the hierarchy of levels into the realm of the media system. Augmenting the layered model described above with the computer's visual computation abilities leads to the scheme shown in Table 1.

Thus, the computer becomes a powerful tool in the hands of an image researcher because it can help us to access large amounts of visual data. At the same time, it leaves cultural contextualization to the art-historical researcher, who can dive deeper into the individual images. Art is made by humans for humans and the computer

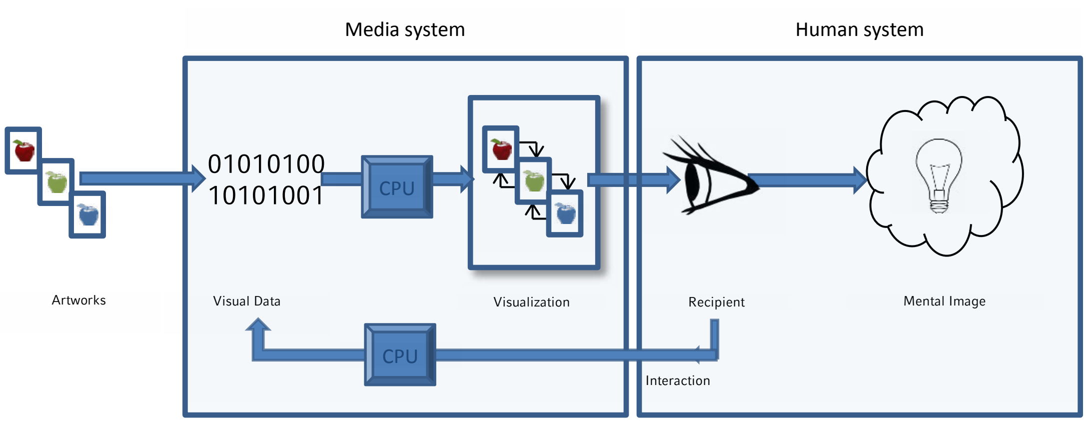

Figure 6: Scheme of the computer as an epistemic tool in the service of image research

can help us cut through the ever-expanding forest of mass visual data. Such digital methods do not replace the human researcher, even if the computer takes over more and more tasks. It helps us to focus on those tasks that we excel in: interpretation, contextualization, and appreciation. The mental archive will remain central and will be extended with computers where applicable. In short, the processing of increasingly larger numbers of images by the computer (*meta* images) allows for deeper interpretation of the individual image by researchers (*mesa* image).

The computer, used as a tool to explore visual data via meta images, puts visualization into a core concept of BID. This concept, visualization, can be added as the fifth "V" in Big Data. The computer and the data it stores are being used via a graphical user interface—an interface between the human and the data. It processes data, reduces information, and transforms it to a dynamic visual phenomenon. This transformation process can apply different algorithms and can be controlled by a number of variables. At the same time, the user interface allows for the manipulation of data via interaction. It allows selection, exploration, and alteration. The human using the system will see, understand, and comprehend the visual phenomenon and then form a mental image of the data, which becomes their structure and meaning. For a long time, Art History has studied the recipient in front of an image. What is being seen is not a passive process, but shaped by previous visual experiences, cultural norms, and expectations.[52] Having the ability to not only contemplate the image, but to also interact with it, makes the viewer an active user and designer of data and allows explorative access to virtually infinite image libraries.

Human system	7. Iconological Interpretation
	6. Iconographic Analysis
	5. Pre-iconographic Description
	4. Object Recognition
Media system	3. Visual Phenomenon of Meta Image
	2. Media System *(pre-processing, including object recognition and pre-iconographic analysis)*
	1. Information

Table 1: Layers in Generation and Understanding of a Meta Image from BID in Computer-aided Art History

BID seems to hold great promise for Art History. The methods that we are about to develop are not only for our own scientific needs, but might also be used to grasp the billions of images that are created outside the artistic realm. The methods can be applied to answer the questions that society has regarding the digital image in general. Who else can make this contribution to society if not the discipline that has for centuries acquired a great historical and visual experience—Art History?

The future form of Art History means more than the analysis of Big Image Data, but at the moment that seems to be the greatest promise *and* challenge for the discipline. More than any other area of the Humanities, Art History deals with images and the analysis of visual data. It is a great opportunity to modernize the discipline. At the same time, we should not forget that, in research, we draw information from multiple sources—not only images. The prototype of an Art History department at a university has historically had two repositories of knowledge: the library and the slide library. While we have digitized the slide library and are discussing how to access its data, the physical book library is yet another source of information. The next step will be to combine and connect both sources of data digitally and make that the basis of our research.

The acquisition of Big Image Data, data analysis, and visualization will open up to new methods in Art History. We need to team up with information scientists, define goals, and develop tools that meet our objectives. Now is the time to develop such new approaches. However, we must remain critical: data analysis does not replace quantitative research; it is an aid to the researcher. Results always need interpretation; they do not replace hypotheses, but are their premise. And statements can only be made based on the data present. Meaning today, we teach our students source criticism, i.e. evaluating a text by asking who has written what, when and why. In the future, we need to add critiquing digital resources to the curriculum—data critique must become a core competence in art-historical studies.

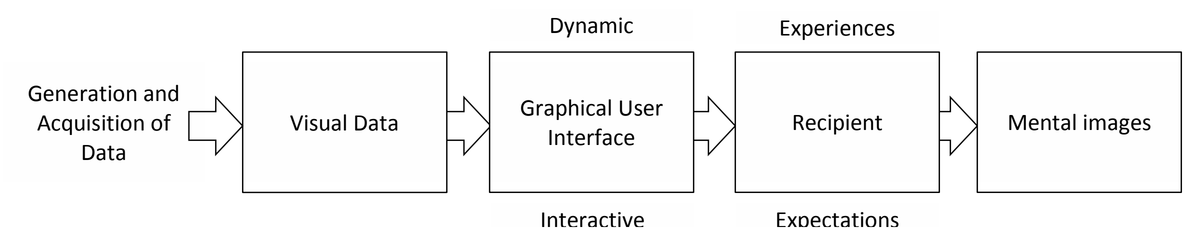

Figure 7: Scheme of a research workflow via a graphical user interface that accesses visual data

We are witnessing a historical shift. Art historians used to be in the position of image recipients. With data visualizations, we have become producers of images ourselves. Which discipline has the toolset to understand what kind of images such visualizations represent, if not Art History? Thus, a critique of the digital image should be included in art-historical curriculums as well—a visual literacy in the digital age, so to speak. And that, too, requires digital competence. If we want to remain a *Bildwissenschaft* in the 21st century, we need to embrace digital methods. Not only to apply them on reproductions of art or popular culture, but also to understand contemporary art that is increasingly digital-born.

Notes

1 Erwin Panofsky, “Iconography and Iconology: An Introduction to the Study of Renaissance Art,” in Erwin Panofsky, Meaning in the Visual Arts: Papers in and on Art History (Garden City, NY: Doubleday, 1955), 26–54.
2 Lena Bader and Martin Gaier/ Falk Wolf (eds.): Vergleichendes Sehen (München 2010).
3 See also Udo Kultermann, ”Der Dresdner Holbeinstreit,” in Udo Kultermann: Geschichte der Kunstgeschichte. Der Weg einer Wissenschaft (München: Prestel Verlag, 1996), 136–141.
4 Thomas Nelson Winter, “Roberto Busa, S. J., and the Invention of the Machine-Generated Concordance,” The Classical Bulletin 75:1 (1999), 3-20 http://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1069&context=classicsfacpub
5 See Anna Bentkowska-Kafel, “Debating Digital Art History,” International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21634>. Date accessed: 19 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21634.
6 See Benjamin Zweig, “Forgotten Genealogies: Brief Reflections on the History of Digital Art History,” International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21633>. Date accessed: 19 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21633.
7 American Standard Code for Information Interchange is a table of characters alongside their numerical repre-sentation. The now widely used UTF-8, which is capable of encoding all possible characters, is fully backward compatible to ASCII.
8 Sorting of pictures according to their median brightness, hue or saturation can easily be done with ImagePlot developed by the team around Lev Manovich: http://lab.softwarestudies.com/p/imageplot.html
9 R. Datta, D. Joshi, J. Li, and J. Z. Wang, “Image retrieval: Ideas, influences, and trends of the new age,” ACM Computing Surveys, 40(2):1-60, April 2008
10 See the article by Babak Saleh in this journal.
11 Moretti, Franco and Alberto Piazza, “Graphs, Maps, Trees: Abstract Models for Literary History” (Verso, 2007).
12 K. Bender, “Distant Viewing in Art History. A Case Study of Artistic Productivity,” International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21639/15412>. Date accessed: 09 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21639.
13 Charles Duhigg, “How Companies Learn Your Secrets,” The New York Times Magazine, http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html, Feb. 16, 2012. The story is based on the presentation “How Target Gets the Most out of Its Guest Data to Improve Marketing ROI” by Andrew Pole at Predictive Analytics World Conference in Washington D.C. (October 19, 2010).
14 Craig Mundie in: “Data, data everywhere“, The Economist, http://www.economist.com/node/15557443 , Feb 25th 2010, see also Clive

Humby, Address at the ANA Senior Marketer's Summit at the Kellogg School (2006), in Michael Palmer, Data is the New Oil (Nov. 3, 2006), http://ana.blogs.com/maestros/2006/11/data_is_the_new.html.

15 Gottfried Boehm, "Die Bilderfrage," in Was ist ein Bild?, ed. Gottfried Boehm (München, Fink, 1994), 325-343.

16 The following paragraphs have already appeared in German in: Harald Klinke, "Bildwissenschaft ohne Bildbegriff," in Bilder der Gegenwart, ed. Harald Klinke, Lars Stamm (Göttingen: Graphentis Verlag, 2013), 11-28.

17 Hans Belting, Bild-Anthropologie. Entwürfe für eine Bildwissenschaft (München: Fink, 2001), 16.

18 Belting, op. cit., 15.

19 "Es gibt keine Daten ohne Datenträger. Es gibt keine Bilder ohne Bildschirme" (Claus Pias, "Das digitale Bild gibt es nicht. Über das (Nicht-) Wissen der Bilder und die informatische Illusion," Zeitenblicke Bd. 2 (2003), §53).

20 Walter Benjamin, "Das Kunstwerk im Zeitalter seiner technischen Reproduzierbarkeit," in Gesammelte Schriften Bd. 1, ed. Rolf Tiedemann und Hermann Schweppenhäuser (Frankfurt/M.: Suhrkamp, 1980), 431-469.

21 Narrations of a general history of media usually start with the invention of the printing press - a text medium. A history of visual media on the other hand focuses on the visual media and therefore begins with the first human image artifacts (e.g. cave paintings). The history of visual media is also not limited to the history of technology. Rather, it is rather concerned with the examination of the cultural consequences of the development of visual media and the definition of the properties of the medium in its historical context.

22 See W.J.T. Mitchell, "Realismus im digitalen Bild," in Bilderfragen. Die Bildwissenschaften im Aufbruch, ed. Hans Belting (München: Fink, 2007), 237-255.

23 Oliver Grau, Virtual Art. From Illusion to Immersion (Cambridge, Mass: MIT Press, 2003), 15.

24 Margarete Pratschke, "Interaktion mit Bildern. Digitale Bildgeschichte am Beispiel grafischer Benutzeroberflächen," in Das Technische Bild, ed. H. Bredekamp and others (Berlin: Akademie Verlag, 2008), 68-81.

25 See Frieder Nake, "Das doppelte Bild," in Bildwelten des Wissens. Kunsthistorisches Jahrbuch für Bildkritik Vol. 3, No. 2, Digitale Form, ed. Margarete Pratschke (Berlin: Akademie Verlag, 2006), 47ff.

26 The Von-Neumann principle describes a computer, which is universal, i.e. independent of the problems to be processed and controlled by binary coded commands.

27 Edmund Husserl, "Phantasie, Bildbewusstsein, Erinnerungen. Zur Phänomenologie der anschaulichen Vergegenwärtigung," in Husserliana Band 23, ed. Eduard Marbach (Den Haag: Martinus Nijhoff, 1980).

28 See Harald Klinke, "The image and the mind," in Art Theory as Visual Epistemology, ed. Harald Klinke, 1-10.

29 See Belting, op.cit., 38.

30 See Peter Stahlknecht and Ulrich Hasenkamp, Einführung in die Wirtschaftsinformatik 11th ed. (Berlin, Heidelberg, New York: Springer, 2005), 94-97.

31 Panofsky, op. cit..

32 Specifically in the Corpus geniculatum laterale (see Hans-Otto Karnath und Peter Thier, Kognitive Neurowissenschaften, 3rd ed. (Berlin: Springer Medizin, 2012), 36-37.

33 http://www.prometheus-bildarchiv.de/

34 http://www.artstor.org/content/artstor-digital-library-features-benefits

35 https://www.google.com/culturalinstitute/browse/

36 http://www.artigo.org/

37 Jay H. Bernstein, "The data-information-knowledge-wisdom hierarchy and its antithesis," in Proceedings North American Symposium on Knowledge Organization, ed. E. K. Jacob, and B. Kwasnik , Vol. 2 (Syracuse, NY, 2009), 68-75.

38 R. L. Ackoff, "From data to wisdom," Journal of Applied Systems Analysis 15 (1989), 3-9.

39 Charles Duhigg, op. cit.

40 Douglas Laney, "The Importance of ‚Big Data': A Definition," in Gartner, 2012. http://www.gartner.com/resId=2057415.

41 Harald Klinke, "Heinrich Wölfflin in Zeiten digitaler Kunstgeschichte," in Kunstgeschichten 1915. 100 Jahre Heinrich Wölfflin: Kunstgeschichtliche Grundbegriffe, ed. Matteo Burioni and others (Passau, 2015) , 415-421.

42 "Dimension reduction is a projection of a space of many dimensions into a fewer dimensions – in the same way as a shadow of a person is a projection of a body in three dimensions into two dimensions." Manovich,

Lev, "Data Science and Digital Art History." International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21631>. Date accessed: 19 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21631., 30

43 William J. Thomas Mitchell: Picture Theory. Essays on verbal and visual representation (Chicago/London: The University of Chicago Press, 1994), 57.

44 "(…) more than 1 billion active devices and 782 million iCloud users" (Juli Clover: Eddy Cue and Craig Federighi Discuss Bloated Software Accusations, Upcoming iTunes Plans, http://www.macrumors.com/2016/02/12/eddy-cue-craig-federighi-bloated-software/ February 12, 2016.

45 Facebook Reports Second Quarter 2016 Results, https://investor.fb.com/investor-news/press-release-details/2016/Facebook-Reports-Second-Quarter-2016-Results/default.aspx.

46 Hans Jonas, "Homo Pictor. Von der Freiheit des Bildens," in Was ist ein Bild?, ed. Gottfried Boehm (München: Fink, 2006), 105-124.

47 Of course one could question if today the differentiation between artistic and general images should be sustained. Lev Manovich recently published a Manifesto in which he argued that there is no need for that anymore. A democratic approach would now be possible in which every image process would be stored through the simple fact that there is enough storage on the servers. And with BID there would be methods at hand that could integrate every one of the billions of images (Lev Manovich: Manifesto for Democratic Art History, http://lab.softwarestudies.com/2016/02/manifesto-for-democratic-art-history.html). But is it just elitism that makes the difference between art and plain images made by laymen? Or shouldn't we think of artists as the ones who are experimenting with new image media and try to test what an image can be? (Thanks to Liska Surkemper for the thoughts on this paragraph.)

48 Liska Surkemper, "Zu einer Theorie vom Ende der Theorie," Archplus, 48 (221), 2015, 66-69.

49 Steve Jobs, in: Memory & Imagination. New Pathways to the Library of Congress, Michael Lawrence Films, 1990

50 A discussion of the concept of the master narrative see: Jean-François Lyotard, The Post-modern Condition: A Report on Knowledge, 1979.

51 Mitchell: Picture Theory, op. cit., 49.

52 Klinke: The image and the mind, op. cit.

Bibliography

Ackoff, R. L. "From data to wisdom." Journal of Applied Systems Analysis 15, 1989, 3–9.

Andrew Pole. "How Target Gets the Most out of Its Guest Data to Improve Marketing ROI", presentation by at Predictive Analytics World Conference in Washington D.C. (October 19, 2010).

Bader, Lena and others (eds.). *Vergleichendes Sehen*, München 2010.

Belting, Hans. *Bild-Anthropologie. Entwürfe für eine Bildwissenschaft* (München: Fink, 2001).

Bender, K.. "Distant Viewing in Art History. A Case Study of Artistic Productivity." *International Journal for Digital Art History*, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21639/15412>. Date accessed: 09 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21639.

Benjamin, Walter: "Das Kunstwerk im Zeitalter seiner technischen Reproduzierbarkeit." In *Gesammelte Schriften*, Bd. 1, edited by Rolf Tiedemann und Hermann Schweppenhäuser (Frankfurt/M.: Suhrkamp, 1980), 431–469.

Bentkowska-Kafel, Anna. "Debating Digital Art History." International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21634>. Date accessed: 19 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21634.

Bernstein, Jay H. (2009). "The data-information-knowledge-wisdom hierarchy and its antithesis." In *Proceedings North American Symposium on Knowledge Organization*, ed. by Jacob, E. K. and Kwasnik, B. . Vol. 2, 2009, Syracuse, NY, 68–75.

Boehm, Gottfried. "Die Bilderfrage." In Was ist ein Bild?, edited by Gottfried Boehm (München, Fink, 1994), 325–343.

Charles Duhigg: "How Companies Learn Your Secrets." In *The New York Times Magazine*, http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.

Clover, Juli. Eddy Cue and Craig Federighi Discuss Bloated Software Accusations, Upcoming iTunes Plans, http://www.macrumors.com/2016/02/12/eddy-cue-craig-federighi-bloated-software/ February 12, 2016.

Datta, R. and others. "Image retrieval: Ideas, influences, and trends of the new age." In *ACM Computing Surveys*, 40(2):1-60, April 2008.

Grau, Oliver. *Virtual Art. From Illusion to Immersion* (Cambridge, Mass: MIT Press, 2003).

Husserl, Edmund. "Phantasie, Bildbewusstsein, Erinnerungen. Zur Phänomenologie der anschaulichen Vergegenwärtigung:" In *Husserliana,* Vol. 23, edited by Eduard Marbach (Den Haag: Martinus Nijhoff, 1980).

Jobs, Steve, in Memory & Imagination. New Pathways to the Library of Congress, Michael Lawrence Films, 1990.

Jonas, Hans. "Homo Pictor. Von der Freiheit des Bildens." In *Was ist ein Bild?,* edited by Boehm, Gottfried (München: Fink, 2006), 105–124.

Karnath, Hans-Otto and Peter Thier (eds). *Kognitive Neurowissenschaften*, 3rd ed. (Berlin: Springer Medizin, 2012), 36–37.

Klinke, Harald. "Bildwissenschaft ohne Bildbegriff." in *Bilder der Gegenwart*, edited by Harald Klinke and Lars Stamm (Göttingen: Graphentis Verlag, 2013), 11–28.

Klinke, Harald. "Heinrich Wölfflin in Zeiten digitaler Kunstgeschichte." In: *Kunstgeschichten 1915. 100 Jahre Heinrich Wölfflin: Kunstgeschichtliche Grundbegriffe*, edited by Matteo Burioni and others (Passau, 2015), 415–421.

Klinke, Harald. "The image and the mind." In *Art Theory as Visual Epistemology*, edited by Harald Klinke, 1–10.

Kultermann, Udo. "Der Dresdner Holbeinstreit." In Udo Kultermann. *Geschichte der Kunstgeschichte. Der Weg einer Wissenschaft* (München: Prestel Verlag, 1996) 136–141.

Laney, Douglas. "The Importance of ‚Big Data': A Definition," in: Gartner, 2012. http://www.gartner.com/resId=2057415.

Lyotard, Jean-François. *The Postmodern Condition: A Report on Knowledge*, 1979

Manovich, Lev. "Data Science and Digital Art History." International Journal for Digital Art History, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21631>. Date accessed: 19 mar.

2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21631., 30.

Manovich, Lev. *Manifesto for Democratic Art History*, http://lab.softwarestudies.com/2016/02/manifesto-for-democratic-art-history.html.

Mitchell, William J. Thomas. "Realismus im digitalen Bild," in *Bilderfragen. Die Bildwissenschaften im Aufbruch*, edited by Hans Belting (München: Fink, 2007), 237–255

Mitchell, William J. Thomas: *Picture Theory. Essays on verbal and visual representation* (Chicago/London: The University of Chicago Press, 1994), 57.

Moretti, Franco, and Alberto Piazza. *Graphs, Maps, Trees: Abstract Models for Literary History* (Verso, 2007).

Mundie, Craig. "Data, data everywhere", *The Economist*, http://www.economist.com/node/15557443 , Feb 25th 2010, see also Clive Humby. "Address at the ANA Senior Marketer's Summit at the Kellogg School." In: Michael Palmer, *Data is the New Oil* (Nov. 3, 2006), http://ana.blogs.com/maestros/2006/11/data_is_the_new.html.

Nake, Frieder. "Das doppelte Bild" In *Bildwelten des Wissens.* Kunsthistorisches Jahrbuch für Bildkritik Band 3, Nr. 2, Digitale Form, edited by Margarete Pratschke (Berlin: Akademie Verlag, 2006).

Panofsky, Erwin. "Iconography and Iconology: An Introduction to the Study of Renaissance Art." In: Panofsky, Erwin: *Meaning in the Visual Arts: Papers in and on Art History* (Garden City, NY: Doubleday, 1955), 26–54.

Pias, Claus. "Das digitale Bild gibt es nicht. Über das (Nicht-) Wissen der Bilder und die informatische Illusion:" *Zeitenblicke* Bd. 2 (2003), §53.

Pratschke, Margarete. "Interaktion mit Bildern. Digitale Bildgeschichte am Beispiel grafischer Benutzeroberflächen," in *Das Technische Bild*, edited by H. Bredekamp and others (Berlin: Akademie Verlag, 2008), 68–81.

Stahlknecht, Peter and Ulrich Hasenkamp. *Einführung in die Wirtschaftsinformatik,* 11th ed. (Berlin, Heidelberg, New York: Springer, 2005), 94-97.

Surkemper, Liska. "Zu einer Theorie vom Ende der Theorie." In: *Archplus*, 48 (Winter 2015), 66–69.

Winter, Thomas Nelson. "Roberto Busa, S.J. , and the Invention of the Machine-Generated Concordance." *The Classical Bulletin* 75:1 (1999), 3-20 http://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1069&context=classicsfacpub.

Zweig, Benjamin. "Forgotten Genealogies: Brief Reflections on the History of Digital Art History." *International Journal for Digital Art History*, [S.l.], n. 1, june 2015. ISSN 2363-5401. Available at: <http://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21633>. Date accessed: 19 mar. 2016. doi:http://dx.doi.org/10.11588/dah.2015.1.21633.

Big Image Data

Harald Klinke has a Ph.D. in art history and a Master of Science in Information Systems. Currently he is Assistant Professor at the Ludwig Maximilian University, Munich, and responsible for the doctoral program “Digital Art History”. He conducts research on visual communication, digital media, and Big Image Data in art-historical contexts.
From 2008 to 2009, he worked as a Lecturer of Visual Studies (Bildwissenschaft) at the Art History Department of the University of Göttingen. From 2009 to 2010, he conducted research, supported by a grant from the German Research Foundation (DFG), as a Visiting Scholar at Columbia Uni-versity, New York. He has published books on American history paintings, digital images and art theory as visual epistemology.

Correspondence e-mail: h.klinke@lmu.de

Featured Article

Featured Article

Figuring out Art History

Maximilian Schich

Abstract: World population and the number of cultural artifacts are growing exponentially, or faster, while cultural interaction approaches the fidelity of a global nervous system. Every day hundreds of millions of images are loaded into social networks by users all over the world. As this multiplicity of new artifacts veils the view of the past, like city lights obscuring the night sky, it is easy to forget that there is more than one Starry Night, the painting by Van Gogh.

Like in ecology, where saving rare species may help us in treating disease, art and architectural history can reveal insights into the past, which may hold keys to our own future. With humanism under threat, facing the challenge of understanding the structure and dynamics of art and culture, both qualitatively and quantitatively, is more crucial now than it ever was. The purpose of this article is to provide perspective in the aim of figuring out the process of art history – not art history as a discipline, but the actual history of all made things, in the spirit of George Kubler and Marcel Duchamp. In other words, this article deals with the grand challenge of developing a systematic science of art and culture, no matter what, and no matter how.

Keywords: art history, culture, data, science, digital, quantitative, process

Introduction[1]

Imagine we knew there was life on Mars and that there were biologists equipped with both the right skills and fearless enough to figure it out. It would be obvious to fund more research into Martian biology, including to build and send a spaceship to observe Martian life up close, at least from orbit. In art history, we are in a similar situation. We know there is cultural complexity emerging from local activity of large numbers of cultural agents. Based on qualitative observation and quantitative measurement we are familiar with intricate large-scale patterns in art and culture that are non-average, non-random, and hard-to-understand. While many traditional departments of art history are stagnant or shrinking, there are rapidly growing numbers of researchers in computer science and physics both curious and well equipped to advance our understanding of cultural complexity. Combining their curiosity and skills with solid expertise in the domain of arts and culture, we are ready to build laboratories that will advance our

Figure 1: Between this first visualization and publication lie three years of doing science. The figure shows a still from an animation tracking individuals from birth to death (from blue to red). Published within "A Network Framework of Cultural History" in Science Magazine and in a poetic transformation as "Charting Culture" in the Nature video channel, the original purpose of the visualization was to help the group of researchers to find and understand quantitative patterns. The final animation eventually accumulated more than one million views, and was featured among "Best Data Visualizations in 2014" in FlowingData, "Best American Infographics 2015", the "NSF/PopSci Vizzies top 10", and "Macroscopes to Interact with Science" at Scimaps.org. Visualization: Maximilian Schich & Mauro Martino, www.cultsci.net.

understanding beyond the anecdotal, theoretic, or what can be achieved by even the most productive researchers using qualitative inquiry alone.

This article provides a perspective towards a systematic science of art and culture where possible advances are driven by explosively growing amounts of data, including images, and by visualizations, or more precisely, scholarly figures that act as the lingua franca in this joint enterprise. Figures, like data and processing power, accelerate research as they allow for communication between communities of practice that internally rely on mutually opaque terminologies, differential equations, algorithms, or distinct workflows (cf. Figure 1).

Like the featured article by Lev Manovich in the last issue, my arguments will be heavily informed by my own expertise, which includes a scholarly background in art history, classical archaeology, psychology, what is now called graph data, and complex network science. My own work addresses questions and challenges of art, architectural, and cultural history, using a multidisciplinary approach that integrates qualitative inquiry and observation, with methods of computation, natural science, and information design. The resulting research processes are mostly characterized by international collaboration and co-authorship. Work procedures are expressed in a distributed, lab-style environment inspired by architectural think tanks, corporate design studios, and labs in physics or systems biology. Products aim at high impact journals, conference proceedings, and occasional monographic publications, all of which ideally also cater to a broad audience. Striving to deal with images and figures in the manner of high-quality artist publications, some results of our work are also increasingly themselves exhibited as artworks – not accidentally. Though inspired by science, my approach is not without precedent in art and architectural history, standing on the shoulders of practitioners such as Geymüller, Barr, Malraux, Kubler, the Eameses, Venturi & Scott-Brown, Doxiadis, Koolhaas, and others.[2]

While I believe that the approach outlined here will become a prevalent model in the ecology of methods aiming to understand art and culture, it is important to mention that what is currently called digital art history is both less and more. Just like systems biology has established itself besides more traditional forms of practice, such as the observation of individual birds, the quantitative multidisciplinary approach will exist beside other forms of art historical practice. Digital art history is less than what I outline below, as many large-scale technological projects in the field are characterized by engineering approaches that aim to build tools to find patterns and facilitate traditional practice, such as the comparison of individual images. Aiming to understand the large-scale patterns we reveal, I will underscore that such engineering needs to be complemented by science, in the sense of physics, i.e. by formalizing quantitative laws.[3] Digital art history is also more than what I outline below, as it includes a large variety of methods that do not require scientific, computational, or aesthetic skills, while being immediately accessible even to non-tech-savvy or science-minded art historians. These immediate aspects include high-bandwidth browsing of source texts, images, and urban environments, ever closer or distant readings, and of course the simple usage of apps, digital libraries, databases, and other exploration tools that continue to be developed over time.[4]

Looking into the future, all these immediate aspects of digital art history will probably revert to simply being called art history, like digital astronomy, after serious debate in the 1980s, reverted back to astronomy, as almost no astronomer could imagine working without digital data or digital methods any longer.[5] The approach outlined below on the other hand may grow into a systematic science of art and culture, with a growth trajectory similar to systematic biology, "Broadly defined".[6] This systematic approach will also shed the denominator digital over time, as methods may include analog, quantum, and other forms of computation, and data may come in forms other than digital zeros and ones. With that in the back of our mind, for now, we can safely locate the outlined approach within the scope of digital art history.

In the following paragraphs I will argue that the process of art history is both transcending and exponential, while the discipline of art history, in principle, has no limits in method. While the term 'big data' is either relative or nonsense, I will show that "more is different" and that understanding the resulting "organized complexity" in art and culture requires an integration of natural science and humanistic inquiry, which is not something to fear, but a positive development to embrace. I will convince the reader that humanistic inquiry and natural science share the same basic research pipeline, and that norm data is simply the clear end of a massive gradient of uncertainty. Concluding, I will point to outstanding examples of art historical research beyond the discipline of art history.

The process of art history is transcending

It is not new to point out that the process of art history transcends the boundaries of specialized disciplines dedicated to more or less arbitrary subsets (cf. Figure 2). Salvatore Settis, for example, reminds us that the radical divorce between classical archaeology and modern art history is obviously made up, as we all know that we are dealing with a single historical process.[7] Nevertheless, students of

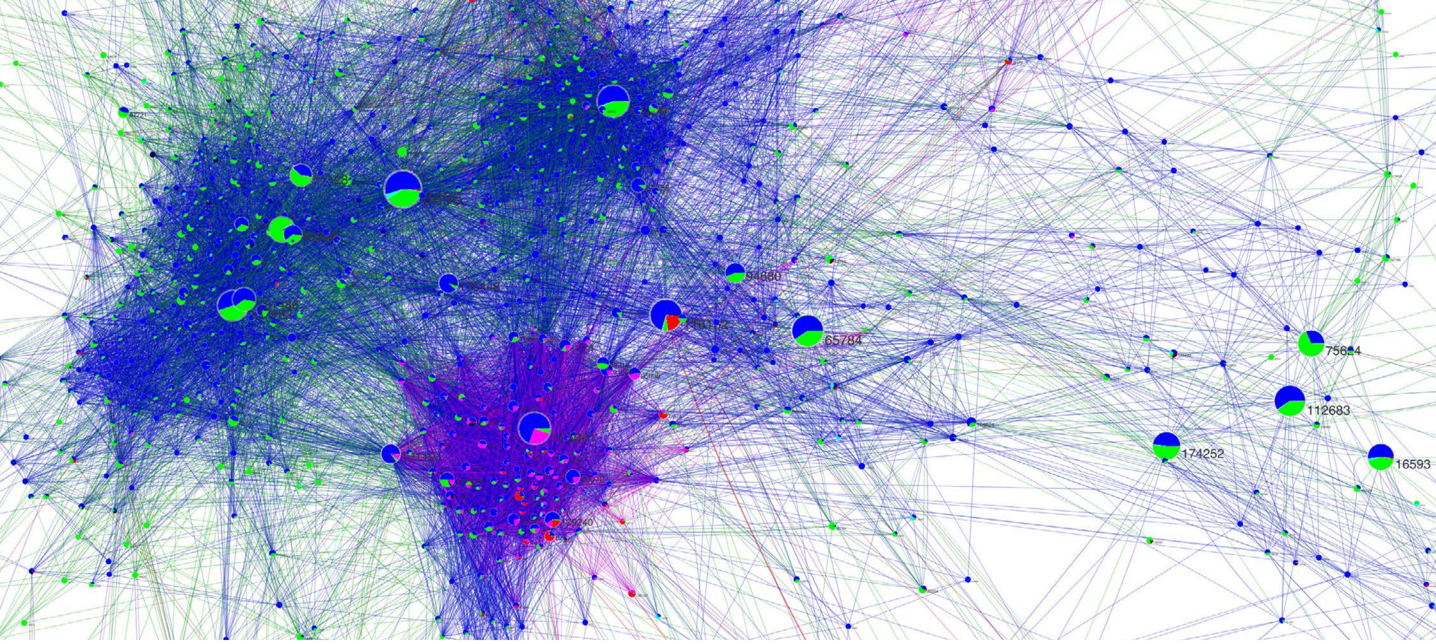

Figure 2: Subject areas in art and culture are highly entangled. This becomes visible by using computation and visualization to filter and map emerging communities of topics (nodes) and their overlap (links), across subjects (blue), locations (green), eras (magenta), and individuals (red). The evolution of subject areas and their overlap, as discussed further in Schich and Coscia "MLG 2011", bears striking stabilities over decades, while also indicating non-intuitive growth that needs to be measured, similar to the evolution of the "PACS" classification in physics. Like in classical archaeology (shown here), we can expect similar organized complexity in general art history, with a stronger focus on known individuals, in addition to objects, locations, and eras. Visualization: Maximilian Schich & Michele Coscia.

art history are constantly exposed to supposedly necessary limitations. Art history pioneer Heinrich Wölfflin famously said that even though "it is hard to answer the man who regards history as an endless flow," on the other hand "intellectual self-preservation demands that we should classify the infinity of events with reference to a few results,"[8] obviously implying the use of his famous pairs of terms, as well as nations, or stylistic periods as necessary categories. In Wölfflin's tradition, a popular introduction to art history in the German language dedicates several sub-chapters to delineate subject areas of art history as necessary specializations to get a job.[9] The most curious products of such definitions are tenure-track positions for aspiring young faculty that are sometimes limited to extremely narrow topics, such as profane buildings of a particular family, covering a couple of decades in a particular Italian city, analyzed with a particular method. On the other hand, brilliant art historians that have "not specialized enough" are often limited to teaching general survey courses in non-tenured adjunct positions or acting as guides in the tourist industry.

As a systematic science of art and culture transcends such limitations, the exploration and summary of the process of art history, as a whole, becomes a research priority that is again as important and justified as the inquiry into local specifics. Widening the scope it aims to advance our understanding of the history of all made things, in the spirit of George Kubler, as well as following Marcel Duchamp in asking if there can be any works that are not "of art."[10] As Ernst

Gombrich pointed out, in an obvious allusion to Charles Darwin, "the coral reef of culture was built by short-lived human beings, but it's growth is a fact not a myth."[11] As with coral reefs in the sea, we can and should study individual subsections of art history, while not forgetting that we might get novel insights by looking down at the whole structure and dynamics from space. Due to the organized complexity involved, as detailed below, we will be rewarded with breathtaking beauty and radically new insights that cannot be achieved by local inquiry.

The process of art history is exponential

The process of art history as a whole is intimidating. World population currently grows at a faster than exponential rate, which means the so-called Malthusian explosion is indeed exploding itself.[12] As technological innovation extends the carrying capacity of the planet and raises the amount of artifacts that can be produced by a single individual, the dynamics involved are approaching the fidelity of a global nervous system, whose understanding becomes crucial for our future survival.[13] As the subjects of art history, both past and present, grow with the system and determine parts of it, their study can provide a segue towards a better understanding of the system as a whole.

While the entire history of the Paris salon, spanning over more than two centuries, comprises less than 160,000 artworks,[14] in 2013 more than 350 million pictures were uploaded by Facebook users every single day.[15] Judging from the fraction of Instagram images identified as self-portraits in the Selfiecity project,[16] this likely means that our daily output eclipses the whole prior history of portraits, not only as covered by the discipline of Western art history, but from the moment our species started to produce images more than 40,000 years ago,[17] to at least well into the 20th century.

The number of known artists, as noted in Allgemeines Künstlerlexikon (AKL) and the Getty Union List of Artist Names, grows exponentially for about 800 years, on a trajectory that is faster (still) than world population growth.[18] Today, as a result of this explosion, the creative industry is $4.29 trillion dollars in size. If it was a country, this Orange Economy, as the Inter-American Development Bank calls it,[19] would be larger than the German economy, only surpassed by Japan, China, and the United States. With $646 billion dollars it would be the ninth-largest exporter of goods, and have the fourth largest labor force, with 144 million workers. This means the current labor force in the creative industries eclipses the documented creatives and artists in AKL by about two orders of magnitude. There are over 100 times more creatives making a living today than noted in history since 1200 CE.

Looking at the inventories of well-funded museums, similar growth trajectories would become evident, which either means the number of objects grows more or less exponentially over time, or we tend to forget material in an exponential way. No matter how much the actual growth of production, the documentation bias, and the decay of preservation contribute to this situation, there can be no doubt that the exponential nature of our record presents a serious challenge in working towards a better understanding of the process of art history.

There is no established means of qualitative inquiry that can deal with this form of dynamics. And there is no trivial way to dissect the exponential growth trajectories into meaningful, non-overlapping periods. As the exponential growth of cultural output, like world population, contributes to a large number of indicators that characterize the sustainability of our species, understanding the process of art and cultural history advances from a harmless hobby horse to a mission critical application of research and education. As such, exponential art history feeds into a deep history that ties all disciplines into a single narrative of phase transitions from the big bang to our own daily experience.[20]

The discipline of art history has no limits in method

As we strive to understand the process of art history, we are using established methods and developing new approaches, both qualitative and quantitative, and communicating our results to emerging communities of interested scholars, as well as a broader audience. Ironically, much time is spent defining the in-crowd, to rewrite the creation myth of our practice, to debate on a purely theoretical level, or to reframe the field from individual perspectives. All this is necessary, and this journal consciously provides a forum for such discussion, but we should not forget that our mission is first and foremost to understand the process of art history. Isn't it ironic that a cited search for Warburg's Bilderatlas returns a wealth of literature theorizing the approach,[21] while the majority of practitioners that deal with large amounts of images have never heard about Warburg, even though his idea of Mnemosyne may be as important to our visual cognition and practice as the ideas of Planck are to quantum mechanics? As Vitruvius recommends for good architects, we must combine theory and practice to avoid hunting shadows, while reaching authority and getting to the substance.[22] Just as the architect's goal is to build, our own goal is to understand the historical process. What we call the procedure of reaching this goal is secondary.

Similar to the menu in a Vietnamese restaurant, we are currently confronted with a large variety of concepts, many of which share similar ingredients, while only the initiated are familiar with the subtle and sometimes radical differences. Digital art history, digital humanities, humanities computing, computational art history, culturomics, cultural analytics, and data science in art history are only some of the pertinent concepts on offer in the naming game that leads up to a major tipping point or phase shift in the system.[24] To achieve relevance towards our aim of understanding the process of art history, and make an impact on the audience, it makes no sense to build walls and hide behind one name or the other. It also makes no sense for self-identified "traditional" art historians to avoid, exclude, be afraid of, or look down on those engaged in new perspectives and approaches. Scientists of art and culture will not take over traditional art history. They are not computer-people that provide researchers with a visualization or an automatic tool. They are not a service. Scientists of art and culture are researchers sharing the same goal, namely to understand the subjects and processes of art history. The only difference is that they do not stick with Pad Thai, but opt for the entire menu to reach the goal, and if necessary they change restaurant, or learn to cook. Or, even, invent their own cuisine.[25]

Since I did my PhD in art history, with "too much" classical archaeology, and joined a physics lab as a post-doc, I have been asked very frequently to define what I consider myself. My usual answer is an anecdote about artist Anish Kapoor, who is often asked if he considers himself British, Indian, or Jewish. His smart reply is to point out that we have to stop compartmentalizing people. Instead of providing one of the concepts above, I point out that my aim is to understand the nature of culture by integrating art history with computation, physics, and information design. I am a Professor in Arts & Technology and a founding member of the Edith O'Donnell Institute of Art History. As such, I am teaching courses in art history as well as courses engaging in cultural data science and information design. My research combines both strains and refuses to limit itself to a particular discipline.

Big data is relative or nonsense, but more is different

Both my own work and the work of Lev Manovich has been described as dealing with 'big data', which reflects the size difference of our projects in comparison to other work in art history. We have to admit, however, that we are not overwhelmed by data in the same way as data scientists that deal with real-time streams that are gigabytes per second in size. We do not have to remove or cloak potentially useful data as it comes in. And most of what we do even runs on a single machine, such as the one on your desk. Our data is large, but it would be much larger in an ideal world. In a system of 120,000 individuals moving from birth to death, we have a mere couple of thousand data points over two thousand years, even for the largest centers, such as Paris; in Selfiecity, out of 120,000 Instagram

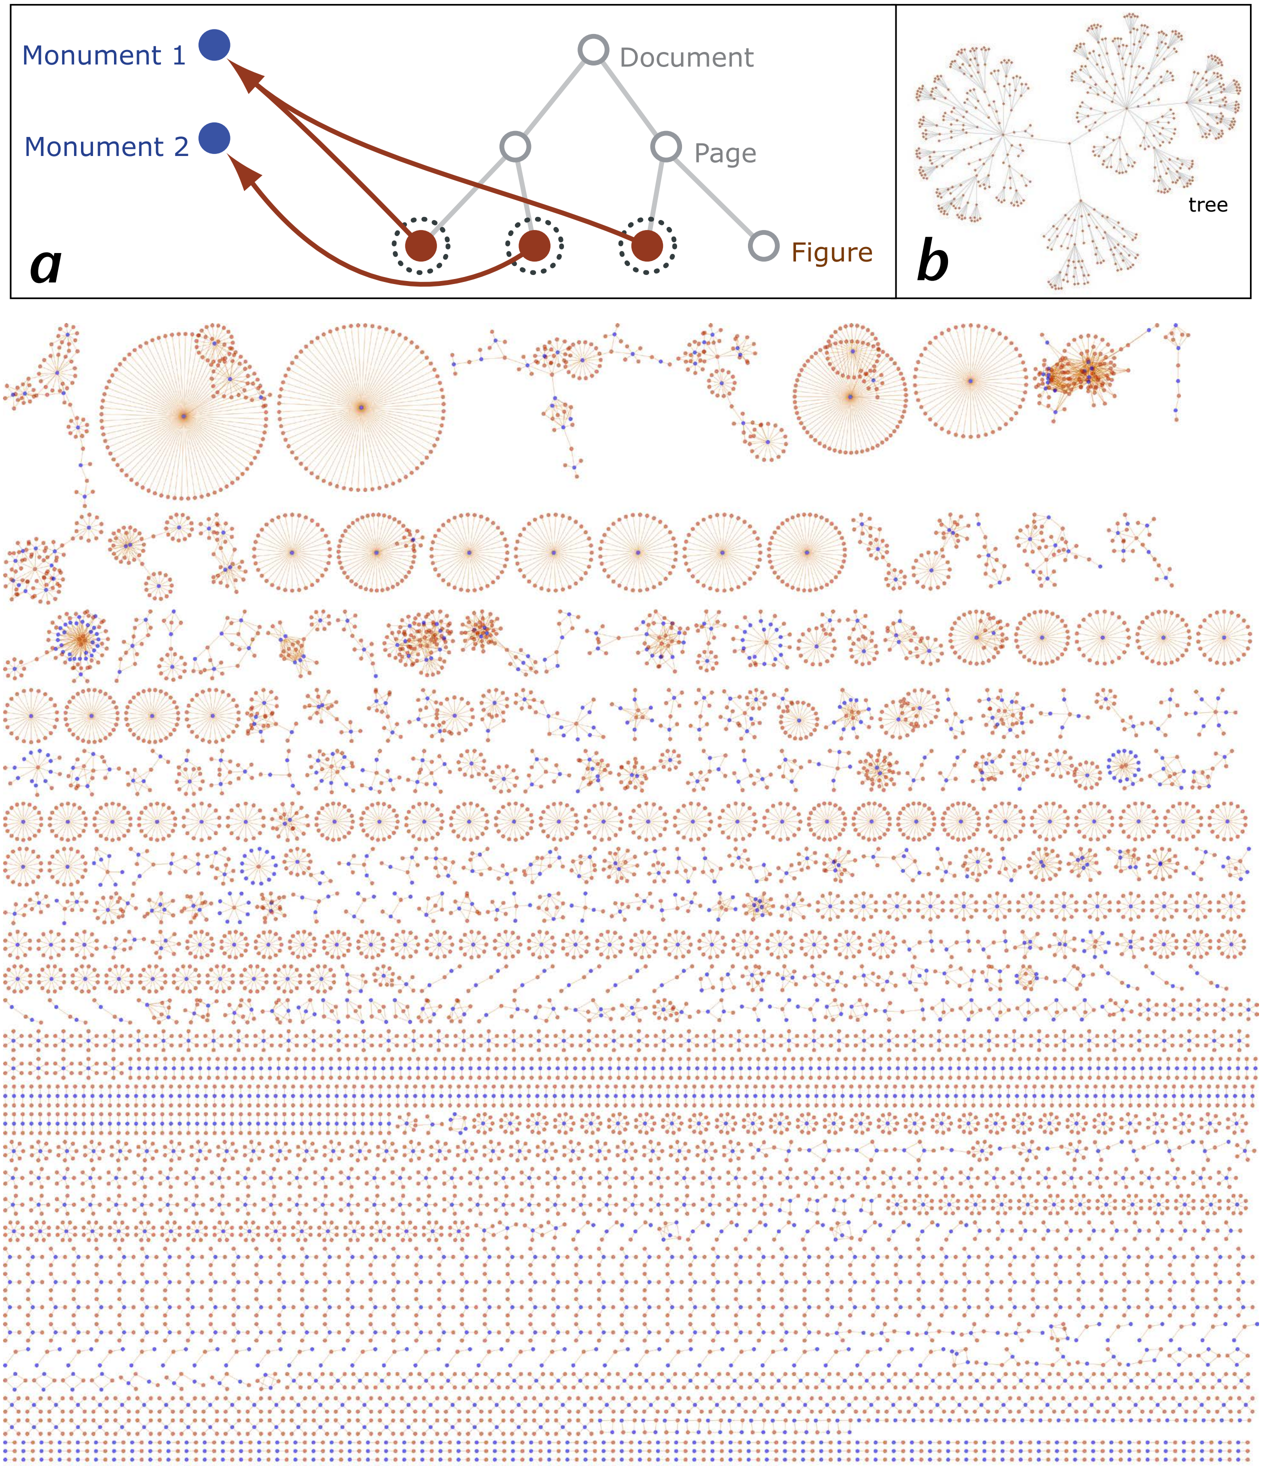

Figure 3: Organized complexity emerges from aggregates of local specifics. To the left, modern documents (brown) are connected with ancient monuments (blue) in the "Winckelmann Corpus". To the right, nodes summarize whole documents, integrating individual drawings and text occurrences into books, etc. As a consequence, the system undergoes a so-called phase transition, forming a single connected cluster. Similar to other complex networks, such behavior

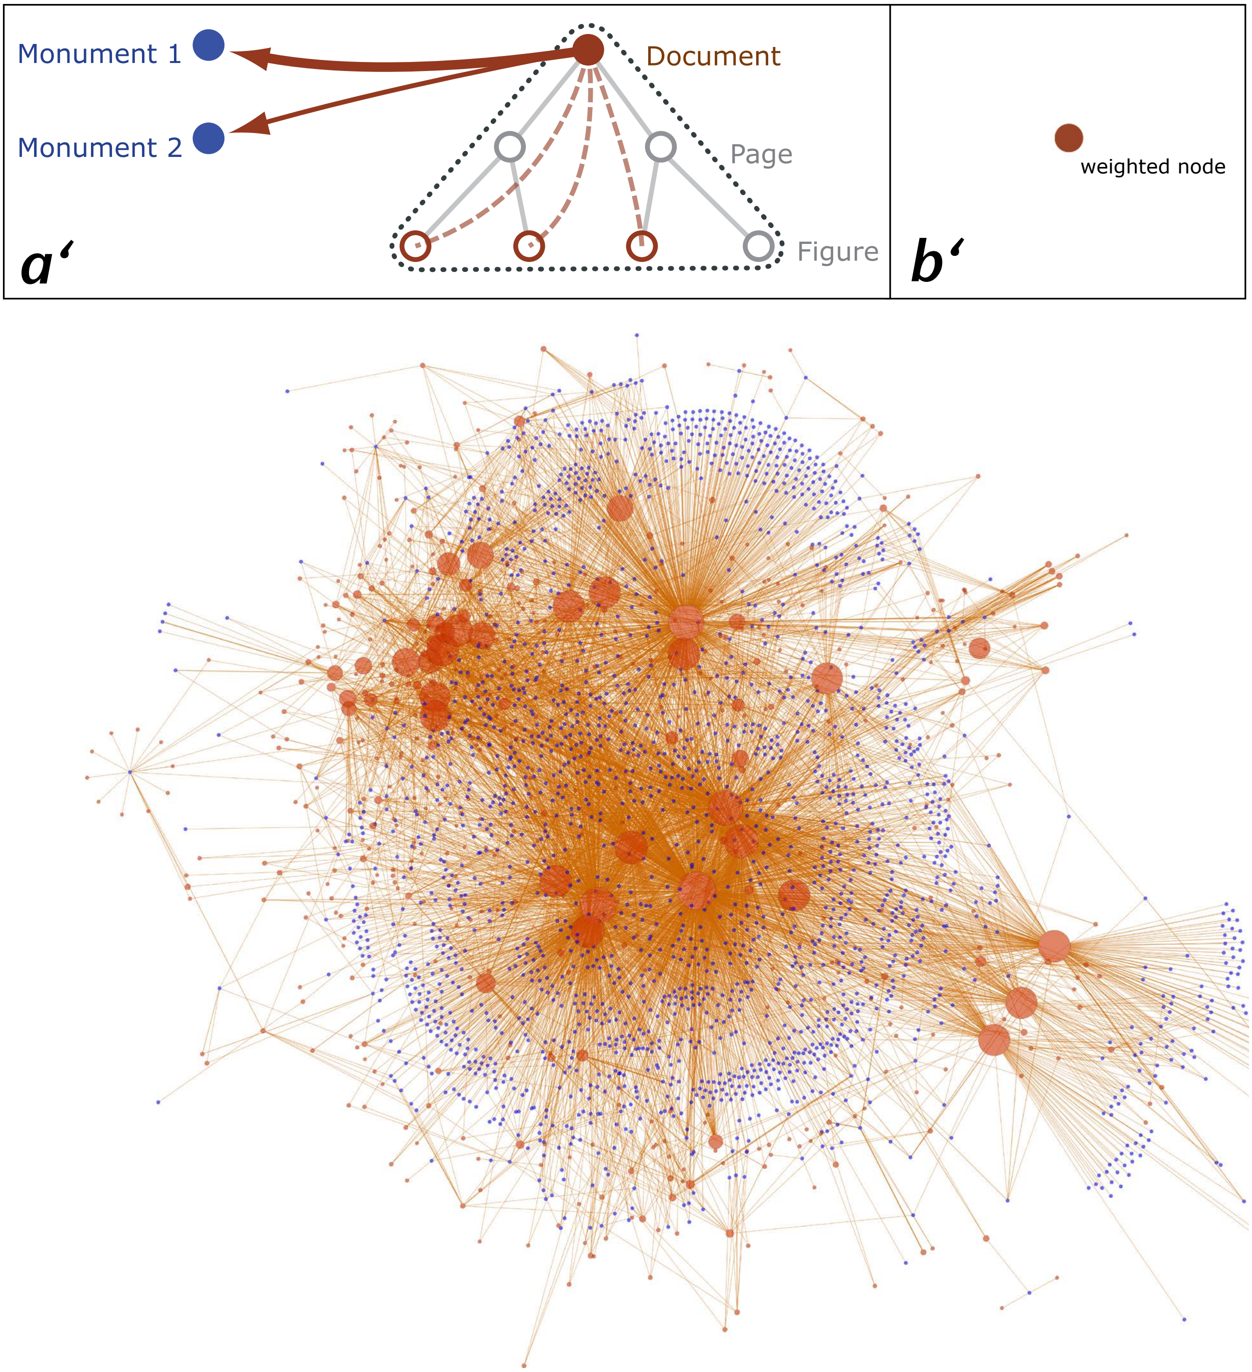

is the subject of mathematical graph theory and physics. Its observation also has immediate consequences for further funding and research. With the largest cluster spanning 100%, as opposed to an expected 90% (cf. Schich “Revealing Matrices”), the Corpus obviously contains monuments “known by Winckelmann”, excluding those “known by his time but not himself”. Data: Kunze & Betthausen “Corpus”, Visualization: Maximilian Schich.

images, only 3600 selfies are above the threshold of quality to make it into the visualization.[26] As a consequence, like most quantitative scientists, when working towards publication, we are worried about issues of undersampling and bias; in short, we worry about having enough data, not about being overwhelmed by too much. Even with the current explosion of data availability, these issues will remain, and, like in economics, social science, and biology, the discussion of bias will occupy a significant amount of time and effort. The extensive discussion of bias in the supporting online material of our recent paper in Science Magazine is a striking example.[27] On the other hand, the discussion of bias is not a weakness, but a strength of quantification. It is easy of course to observe that minority artists are underdocumented, while it is an actionable insight for future funding and research to say by how much compared to the population as a whole.

At its best, the term 'big data' is not an absolute, but relative term that should be avoided in practice, even though it may (still) help when journalists use it. 'Big data' is similar to the colossal order in architecture. Standing in front of Palazzo del Capitanio in Vicenza, the columns indeed seem colossal and intimidating relative to the facade as a whole, while in fact the building is not exactly the size of New St. Peter's in Rome. In a similar way, relatively small amounts of data may look intimidating in relation to qualitative methods of inquiry. For an art historian doing the catalogue raisonné of a very prolific artist, 1 million AKL artists or 1 million Manga pages may seem big.[28] But for data scientists big is when considerable infrastructure is needed to store data, such as 10,000 Tweets per second as they come in, or when they run into the necessity of throwing away data unseen, as in case of the Large Hadron Collider (LHC), where too much image data is generated to even store, let alone to fully analyze, while using the best technology available. For those curious, the CMS (Compact Muon Solenoid) detector of the LHC produces 40 million images at 1 gigapixel resolution per second, which is more than 25 times the number of images in the Prometheus Bildarchiv, at the largest resolution available in some select cases within Google Art Project.[29] From that perspective the available amount of digital data in art history is almost ridiculously small.

At its worst, the term 'big data' is nonsense. Looking for great literature on the topic, it is useful to compare a Google Books search for "big data" with one for "large data." The first returns a mass-market book as the top result, while a search for the latter returns the practical textbook on data science recommended by Lev Manovich in the last issue of this journal.[30]

Be that as it may, on a practical level, 'large' and eventually 'really big' data is relevant to understanding the process of art history as a whole because "more is different."[31] Just as we cannot imagine the full structure and dynamics of the great barrier reef solely by looking at a couple of fish or a bunch of polyps, we cannot understand the large-scale structure and dynamics of the process of art history only by studying a selection of paintings, artists, or archival records. Like a coral reef, the process of art history is a product of "local activity"[32] done by a large number and variety of actors, forming a highly entangled complex system that is literally more than the sum of its parts (cf. Figure 3).[33] The coral reef of culture, like biology, includes large networks of complex networks whose structure and dynamics we can only understand given large amounts of data. [34] The networks involved contain emerging information that is not a property of individual actors, objects, locations, periods, or events, but a property of hard to define aggregates or of the system as a whole. As a consequence, to advance our understanding, we have to combine our traditional domain expertise in art history with methods of complexity science, such as matrix algebra, and advanced graph theory.[35]

Understanding complexity needs science as well as humanities

Nurturing natural science methods to understand the process of art history promises to overcome the long-standing separation of "nomothetic" law disciplines, such as physics, and "ideographic" event disciplines, such as history, as postulated by Wilhelm Windelband in 1894 and famously lamented by C.P. Snow in the 1950s.[36] Warren Weaver in 1948 and Jane Jacobs, implicitly, in 1961 have already argued that such an integration is possible and indeed necessary to address abundant problems of "organized complexity," in both economic and urban systems.[37]

In A Network Framework of Cultural History,[38] we implicitly provide a rigorous mutual justification for such an integration of quantitative and qualitative research. The article shows that quantification in the humanities does indeed work by bringing evidence for the physical "laws of migration" spanning over 800 years, based on simple birth and death records of large numbers of artists and other individuals. On the other hand, the article also shows that quantification cannot replace qualitative inquiry, as the system of cultural history is characterized by massive fluctuations on a local level (cf. Figure 4). Both methods of inquiry bring essential ingredients to the table. Delineating examples, the article further promotes the integration of qualification and quantification by revealing sense-making cultural meta-narratives as they emerge from large amounts of granular information, and by helping to cross-fertilize qualitative domain expertise within the context of a big picture. Finally, as mentioned above, the rigorous quantification of bias on

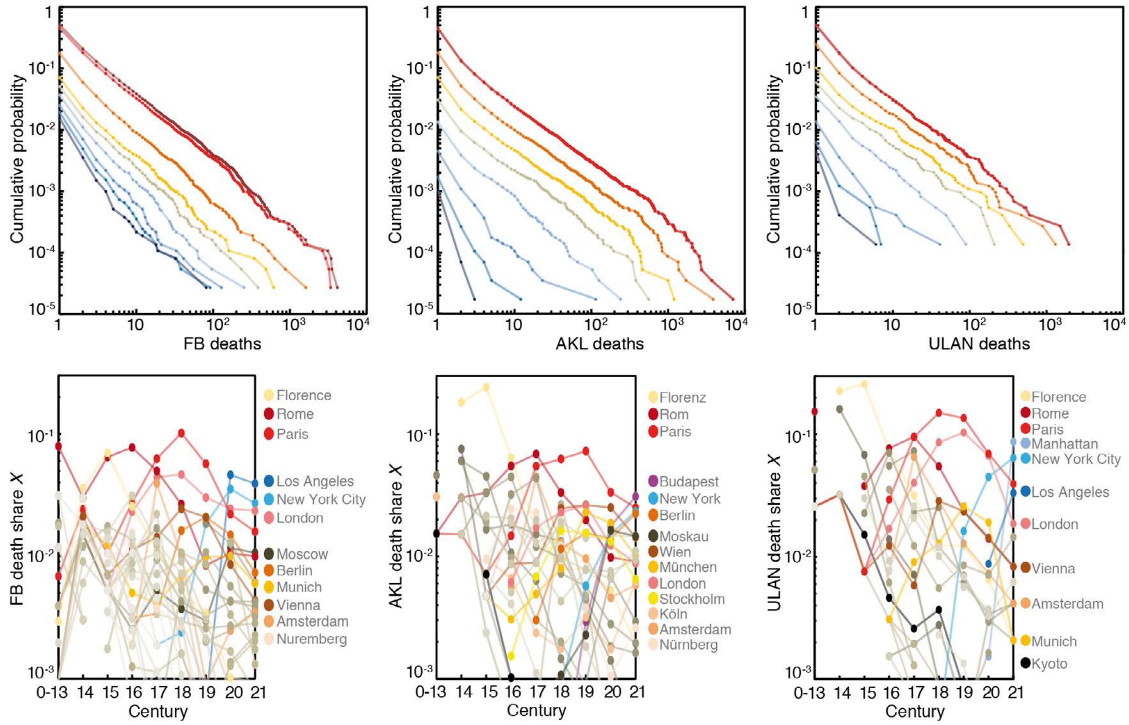

Figure 4: Quantitative science and qualitative inquiry are both necessary and complement each other. The three plots above highlight the necessity to quantify physical laws in cultural history, indicating a heterogeneous size distribution of cultural centers that grows more or less exponentially over time while being stable in slope throughout history. The three plots below make a case for qualitative inquiry, by exposing massive fluctuations in the relative share of notable deaths in cultural centers. Both phenomena are consistent across datasets, even though Freebase.com (FB) has very little overlap with Allgemeines Künstlerlexikon (AKL) and the Getty Union List of Artist Names (ULAN). All plots see Schich et al., "A Network Framework of Cultural History," including the Supporting Online Material.

mesoscopic and global levels in the supporting online material adds to the general usefulness of the combined approach.

While our Science paper was a major breakthrough, the proposal of integrating humanistic inquiry, computation, natural science, and information design to understand the process of art and cultural history, still often invokes a manifest disbelief in quantification and sometimes the almost insulting conviction that such a proposal can't be much more than "data management." Such reactions are not surprising, as the process of understanding art and culture is still dominated by qualitative humanistic inquiry, and the necessary foundations are not taught within the standard curriculum. Technology within the discipline of art history, including quantitative science, is mostly perceived and treated as a complement or service, where qualitative researchers call in computer experts and designers to support their qualitative inquiry. An example of this phenomenon is the social network diagram published recently at MoMA, also cited by Lev Manovich in the last issue. Intended to improve over the famous original Barr chart, the new MoMA diagram has been marketed as the result of a high profile collaboration between a curator at the museum and network analysts in the business school of Columbia University. While the new diagram could have been done by a reasonably talented undergraduate in Digital Humanities within a few minutes, the original chart's irony, using Picasso's bull as a layout algorithm, could only have been produced by an art historian, such as Alfred Barr, who mastered the production of images just as much as their curation.[39]

While some technological applications in art history have achieved flagship status within university departments and research institutes, starting in the 1980s, the differences of perceived authority are still expressed in salary differences between professors and institute leadership recruited from those doing qualitative inquiry, versus lower-paid adjunct or well-paid but temporary employed computer experts and designers on the other side, notwithstanding their pertinent expertise, often underscored by a PhD in art history. It is a step forward to underline that "humanists must work side-by-side with technical experts [...] to get tools, portals, access, etc.," as Thomas Gaethgens, head of the Getty Research Institute, recently said, quoting Johanna Drucker.[40] Such acknowledgment breaks with the implicit pattern of subordination but is not enough. Drucker's statement, and indeed the whole definition of digital humanities according to leading practitioners still implicitly assume that the application of technology in art history is an engineering problem, [41] with the final goal being the production of means that help the actual researchers doing their inquiry. To achieve a deeper understanding of the process of art history we cannot employ such a procedure, akin to civil engineering, where engineers build the street while working side-by-side with future drivers. Instead, deeper understanding is like the honey in a natural beehive. It can only be reached by those who are able to master and adapt the twig or whatever tool will take them there.[42] Everybody involved in the process must have enough expertise in both arts and technology to collaborate towards achieving the ultimate goal of a deeper understanding. It serves to immediately point out that such a proposal is not the suggestion of "white male science" to take over the arts and humanities.[43] Indeed, being modeled on established practices in multidisciplinary network science and systems biology, the proposed science of art and culture promises to attract enthusiasm from a large diversity of researchers, coming from all continents and with much better gender balance than discrete communities of practice.[44]

Uri Alon, author of a popular Introduction to Systems Biology, has introduced a striking model called the "cloud of uncertainty," which can help us to clarify the difference between engineering problems and problems of science.[45] Projects that aim to build tools, portals, and access are engineering problems as they aim to go from problem A to an imagined future solution B. Examples include the digitization of all books ever published, or a database of all paintings in public collections. Both applications require highly skilled researchers, masterful coordination, sophisticated technology, and efficient workflows to be successful. The results may be highly useful to traditional practitioners, but in themselves do not necessarily contribute to a deeper understanding of the subject matter. Projects that aim towards such an understanding, on the other hand, may include some engineering, but are very different in nature, no matter if they choose to employ qualitative or quantitative methods. They need to go where nobody has gone before, even in imagination. The difference is that starting with situation A, we may find out that the imagined solution B is unachievable, putting us into the "cloud of uncertainty," from which we can only escape by mastering whatever method is necessary to reach an unknown and maybe surprising solution C. In addition to scientific skills, this may involve overcoming negative emotions and depression to reach the happiness of insight.[46] As Paul Feyerabend pointed out, this enterprise is essentially anarchic, and we have to act like undercover agents, who play the game of reason, to undercut the authority of reason.[47] There is no fixed workflow pipeline or service that we can call in like a construction firm in civil engineering. Instead, we are required to learn, master and adapt our methods and tools as we go along.

Aiming towards the unknown to eventually find surprising insight is a common trait of basic science and research (Grundlagenforschung), no matter if qualitative or quantitative. Of course, while there is no fixed workflow in this enterprise, there are general recipes and procedures that help to formalize the process of inquiry and raise the chances for new insights. Hedging our resources like an angel investor or venture capitalist, with collaborators being involved in multiple projects, we can minimize the risk and ensure the overall success of a given group of researchers.

Humanistic inquiry and science share the same basic pipeline

A systematic science of art and culture will align with traditional qualitative scholarship in art history, not only in terms of questions, but also in terms of workflow, fixing a major shortcoming in established digital practice within the arts and humanities. Over decades we have spent a large amount of energy developing data models and standards based on formal logic and anecdotal evidence.[48] This was important to get digitization and digital workflows off the ground, but violates a basic principle of scholarship, as it is impossible to arrange material without prior collection and observation of its actual structure.

Preparing an individual piece of scholarship, such as a catalogue entry, a journal article, or a book-sized monograph, we would more or less intuitively follow Cicero's sequence of inventing a speech: First we would collect material; then we would arrange the material, formalize the story, if there is one; and finally, deliver it to our audience.[49] In decades of large-scale database projects we have essentially violated this sequence by arranging the material based on our expectations or sometimes ideology, as opposed to taking a deep look into all the material once it is collected, either by using our own eyes or if necessary more sophisticated measurement instruments. Presupposing large-scale structure to be average, random, or intuitive, many database projects were content to create search platforms or browsing tools, whose aim was to facilitate traditional qualitative inquiry. As the emerging organized complexity in the collected material often was out of sync with presupposed expectations and intuitions, it is not a surprise that many large database projects failed to attract a wider and more persistent audience of users.

Quantitative inquiry that aims to map, understand, and explain organized complexity in large collections of data provides a remedy to this situation. Like in the human genome project, where the successful collection of data did not bring an immediate cure for cancer, but did start a whole new field of inquiry, decades of digital data collection in art and culture provide a highly promising point of departure. In fact, due to systematic quantitative inquiry, the decade long effort will finally stand up to its promise. Once we know and understand the emerging complexity, we will be able to pose and address new qualitative and quantitative questions, which we can't even imagine today. Like in other areas of data-driven science, the resulting pipeline will resemble Cicero's basic sequence of invention, enriched by infinite feedback, as formalized by leading data scientists and designers.[50] Indeed, one could imagine the resulting pipeline with feedback as an auto-catalytic cycle of research breeding more research, with quantification as an accelerating enzyme (cf. Figure 5).

As such, qualitative and quantitative practice will feed into a common cognitive process that will advance our understanding of art and cultural history. As there will be variations in procedure, many papers will start with a figure explaining the pipeline.

Norm data is just the clear end of a massive gradient of uncertainty

A good example for the need of quantitative measurement is our obsession with norm data, authority files, and data model standards. Almost nothing in art history is normal, in the

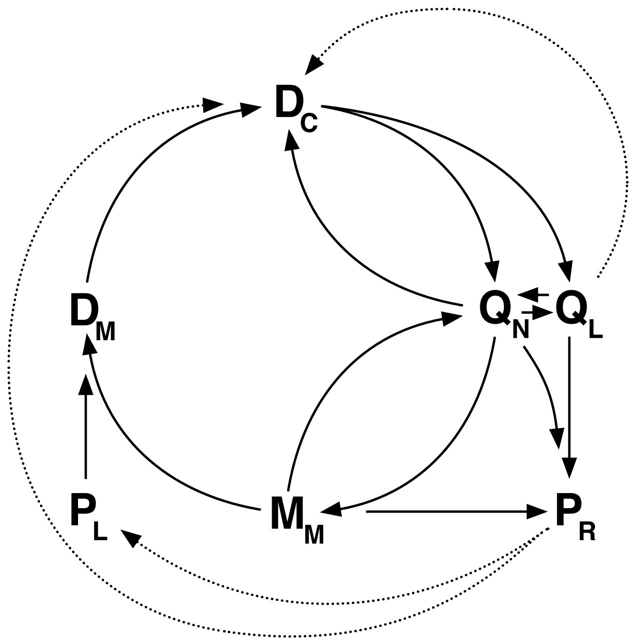

Figure 5: Quantitative science catalyzes the established sequence of digital scholarship. Initially, data models (D_M) are mostly defined using philosophy (P_L), in particular formal logic based on anecdotal evidence. Traditionally, this leads to efficient data collections (D_C), new qualitative observations (Q_L), and eventually the publication of results (P_R), which in turn may lead to better data collection, but usually leaves the original data model intact. Quantitative measurement (Q_N) of organized complexity closes the loop as it leads to creation of mathematical models (M_M), which lead to accelerated change of data models, data collections, and more novel insight. Domain expertise, computation, and visualization are necessary throughout the process. Of course, the figure, inspired by the Eigen and Schuster "Hypercycle", is a cartoon crying for measurement itself. Image: Maximilian Schich.

sense of a normal distribution with a sense-making average, as in the case of a Gaussian bell-curve containing more or less average examples around it. There is no average artist, no typical triumphal arch, no regular Roman sculpture, and no normal nativity scene. Wherever we look we usually find one or a handful of exceptional examples, and a more or less long or "fat tail" of irregular or hybrid examples that are not-so-well-documented, not-so-typical-looking, not-so-well-preserved, or not-so-easy-to-attribute.[51]

I am not saying that there are no well defined groups of objects. What I am saying is that, based on existing evidence, we have to deal with massive gradients of uncertainty. If we want to understand the art market, beyond some well-identified paintings, collectors, auction houses, etc., we have to deal with a vast majority of uncertain attributions, a majority of rare and unknown actors, and of course the unknown amount of "dark data." In other words, art history has to deal with probability distributions and potential sources of bias, just like social science, biology, and other quantitative fields. It is an illusion to think an editorial process can combine normed classification and addressability. If you run an archive: Do assign identifiers to your records, optionally do crude classification, and leave "figuring out" to the whole community of researchers.

To give an example: In the last two years the incredibly talented computer scientist John Resig, who gave us jQuery and Processing.js, is essentially touring prominent visual resource collections in art history to apply computer vision algorithms in order to find duplicate photographs of artworks. Called in like a service, the premise is "to change how photographs and images are managed in archives, libraries, and museums,"[52] working towards a unified or normalized collection of photos that will facilitate traditional scholarship. Such improvement of management by engineering is important, but loses an important chance to accelerate the science of art and culture. As every art historian knows from their own specific practice, highly similar images that can be matched like copies are just the most simple case of visual family resemblence.[53] So one must ask, if we really should split computational management of visual resource collections from scholarly inquiry in art history. Wouldn't it make much more sense to publish the photo archives, like the human genome, to facilitate an explosion of quantitative research not only into duplicates and near-duplicates, but into the entire gradient of similarity? To clarify the potential: A groundbreaking and highly relevant paper, published in the area of computer vision as recent as 2012, already has more than 2500 citations,[54] which means there are likely hundreds of groups that would be more than happy to work with image data that spans more than the last 20 years.

There is outstanding art history beyond art history

It is easy to cite more such examples that are currently beyond the radar of the discipline of art history. Just as social network analysis was beyond the radar of physicists in complex network science 15 years ago,[55] there is a vast amount of work that either precedes or runs parallel to current efforts in digital art history.

In 1967, French geographer Jacques Bertin published his Semiology of Graphics,[56] which, if I had to choose, would be the one book on data visualization that I would take to Mars, if I had to leave everything else behind. Introducing matrix permutation, he claims algorithmic analysis has to go hand in hand with manual sorting. He demonstrates this by using a classification of Merovingian artifacts, i.e. an example taken from the realm of archaeology and art history, likely from a stream of research that discussed the pros and cons of dimensionality reduction, such as principal component analysis (PCA), since more than 50 years ago.[57]

Lev Manovich's image plot software is preceded by a contribution in computer science, published in 1996,[58] just like my own frequency distributions of ancient monuments in Renaissance documents are preceded by Heinrich Dilly, who counted the frequency of artists in the titles of art historical literature over several decades, publishing the result almost ironically in a volume on art history and the Frankfurt school of philosophy.[59] Stanley Milgram did word clouds with a sense-making layout 30 years before they took off, and computer linguists are jealous of James Joyce for having implicitly outlined almost any possible question.[60] We should appreciate and cite such colleagues and giants on whose shoulders we stand. But of course we should also be aware that we can go much further than we could ever before thanks to the unprecedented amount and quality of data today, as well as advances in computational power and scientific method.

Only since recently is it possible to acquire and deal with millions of tourist photos, as well as imagery taken from Google Street view, to extend theories of reption aesthetics (Rezeptionsästhetik) by mapping the density of tourist attention and even calculate the density of viewing cones of individual tourists, as a side effect of reconstructing buildings in 3D without human interaction.[61] Only since recently can we use algorithms to convincingly date architectural details, in order to map the evolution of palace facades in Paris, strikingly mimicking the perception of a well-trained art historian strolling through the city.[62] Only with services the size of Facebook, has it become possible to study the spreading of visual memes on a large scale, revealing cascades that resemble a mathematical theory of biological evolution.[63] Trend analysis in fashion, which traditionally bears striking resemblance with scholarship in art history, is increasingly driven by larger sets of data and quantification.[64] Finally there is an increasing amount of analysis into paintings and artworks, done and published by natural scientists in multidisciplinary environments.[65]

Conclusion

In this article I have outlined a perspective for a systematic science of art and culture that integrates qualitative inquiry with computation, natural science, and information design. As such, cultural science shares the aim of understanding the process of art history with so-called traditional practice. It explores unknown complex emerging structures and dynamics by analyzing large data sets, using both quantitative measurement and qualitative inquiry. Similar to systems biology, the procedure is characterized by multidisciplinary co-authorship and publications that make extensive use of scholarly figures.

The Journal of Digital Art History has the potential to fill an important gap in this enterprise. Positioning itself in a disciplinary niche within an emerging journal hierarchy,[66] similar to Nature Physics, the Journal of Digital Art History complements existing journals that mediate between art and science, such as Leonardo, and multidisciplinary journals, such as

Palgrave Communications, the new social science and humanities equivalent of Nature Communications. The emergence of such a publication infrastructure provides important opportunities for students and researchers engaging in an art and cultural history without limits. With an estimated market demand of more than 140,000 data scientists,[67] and a growing abundance of cultural data, there can be no doubt that the laboratories engaged in the science of art and culture will have an important function in society and are bound to thrive.

Notes

1 Acknowledgments: The author wishes to thank the anonymous donors of UT Dallas ATEC Fellowship #1 and Dirk Helbing at the Professorship for Computational Social Science at ETH Zurich for their generosity and hospitality. As this article is a perspective, not a map, it fails to mention all the excellent work to be summarized in a hopefully forthcoming review.

2 Regarding pioneers see Cortjaens and Heck, Stil-Linien Diagrammatischer Kunstgeschichte on Geymüller; Schmidt-Burkhardt, Stammbäume der Kunst on Barr; and Schich, Rezeption und Tradierung als komplexes Netzwerk pp. 156-160 on Kubler; In addition, compare the workflows leading to Malraux, Le Musée imaginaire de la sculpture mondiale; Eames and Eames, Powers of Ten; Venturi, Scott Brown, and Izenour, Learning from Las Vegas; Koolhaas and Office for Metropolitan Architecture, Content; See also Wigley, "Network Fever," on Doxiadis etc.

3 Compare Richard Feynman's definition of "understanding" in Feynman, Six Easy Pieces pp. 24/25.

4 High-bandwidth browsing cf. "Google Images;" Google Earth;" and "Google Books." Ever closer and distant readings (Moretti and Piazza, Graphs, Maps, Trees.) see for e.g. gigapixel images in "Art Project - Google Cultural Institute;" versus Crandall et al., "Mapping the World's Photos."

5 Roger Malina, "Yes Again to the End of the Digital Humanities ! Please !"

6 "Our Approach | Broad Institute of MIT and Harvard."

7 Settis, "L'opera Di Paul Zanker E Il Futuro Dell'archaeologia Classica."

8 Wölfflin, Principles Of Art History.

9 Belting et al., Kunstgeschichte.

10 Kubler, The Shape of Time; Duchamp et al., À L' Infinitif = In the Infinitive.

11 Gombrich, The Sense of Order, p. 209; Bredekamp, Darwins Korallen.

12 Johansen and Sornette, "Finite-Time Singularity in the Dynamics of the World Population, Economic and Financial Indices."

13 Helbing, "Globally Networked Risks and How to Respond."

14 Whiteley, Index to the Paris Salon Catalogues.

15 "Facebook Has a Quarter of a Trillion User Photos."

16 Manovich et al., "Selfiecity."

17 Aubert et al., "Pleistocene Cave Art from Sulawesi, Indonesia."

18 Schich et al., "A Network Framework of Cultural History." Figs. 1A.

19 Restrepo and Márquez, "The Orange Economy."

20 Exponential indicators see Steffen et al., Global Change and the Earth System; Deep history see Christian, Maps of Time; Exponential patterns in art history see Schich et al., "A Network Framework of Cultural History," and Rosa and Suárez. "A Quantitative Approach to Beauty."

21 Warburg et al., Der Bilderatlas Mnemosyne.

22 Vitruvius, On Architecture. 1,2.

23 Burdick et al., Digital_Humanities; Doing and C. Richard Johnson, "On Applying Signal Processing to Computational Art History;" Michel et al., "Quantitative Analysis of Culture Using Millions of Digitized Books;" "Culture Analytics;" Manovich, "Data Science and Digital Art History."

24 Young, "The Evolution of Social Norms;" and Baronchelli, "Modeling the Emergence of

Universality in Color Naming Patterns."
[25] Compare McGee, On Food and Cooking.
[26] Schich et al., "A Network Framework of Cultural History;" Manovich et al., "Selfiecity."
[27] Multidisciplinary journal articles often include essential information in the Supporting Online Material (SOM). For free access to the Schich et al. SOM enter the Science Magazine website via www.cultsci.net.
[28] The AKL data in Schich et al. was extracted from a dump of 1.1 million XML files; For 1 million Manga pages see Manovich, Douglass, and Huber, "Understanding Scanlation."
[29] "Is the LHC Throwing Away Too Much Data?"; "The Prometheus Image Archive: High-Quality Images from the Fields of Arts, Culture and History;" "Art Project - Google Cultural Institute."
[30] Mayer-Schönberger and Cukier, Big Data; Leskovec, Rajaraman, and Ullman, Mining of Massive Datasets.
[31] Anderson, "More Is Different."
[32] Chua, "Local Activity Is the Origin of Complexity."
[33] Mitchell, Complexity.
[34] Networks of complex networks are abundant in art history: cf. Schich, "Revealing Matrices," and Schich, "Netzwerke von Komplexen Netzwerken in Der (Kunst)Wissenschaft;" The knowledge graph community has a long tradition working with multiple node and link-types: cf. Bollacker et al., "Freebase." The topic now also gains traction in general network science: See the reviews by Kivelä et al., "Multilayer Networks;" and Boccaletti et al., "The Structure and Dynamics of Multilayer Networks."
[35] See the excellent introductions by Barabási, "Network Science Book;" Estrada and Knight, A First Course in Network Theory; and Peter Dodds "Course Home | Matrixology, Season 8."
[36] Windelband, Geschichte und Naturwissenschaft; Snow, The Two Cultures.
[37] Weaver, "Science and Complexity;" Jacobs, The Death and Life of Great American Cities.
[38] Schich et al., "A Network Framework of Cultural History."
[39] "MoMA Makes a Facebook for Abstractionists | ARTnews;" The bull was recently revealed by Adamo and Ortiz. "Leonardo Journal cover page," while in the discipline of art history, we have all been blind.
[40] Gaehtgens, "Thoughts on the Digital Future of the Humanities and Art History."
[41] Burdick et al., Digital_Humanities.
[42] Adams, Parrots, the Universe and Everything.
[43] Chomsky, "Rationality/Science."
[44] Cf. the heterogeneity of practitioners in Schich, Meirelles, and Malina, Arts, Humanities, and Complex Networks.
[45] Alon, An Introduction to Systems Biology; Alon, "How To Choose a Good Scientific Problem."
[46] Alon, "How to Build a Motivated Research Group."
[47] Feyerabend, Against Method.
[48] Cf. Brown University, KODM 2012 Day 1 Panel Discussion.
[49] Cicero, "On Invention." 1,7.
[50] Meyer, "Designing Visualizations For Biological Data."
[51] Newman, "Power Laws, Pareto Distributions and Zipf's Law;" Clauset, Shalizi, and Newman, "Power-Law Distributions in Empirical Data;" see Schich, "Revealing Matrices," and other publications by the author for the prevalence of tailed distributions in art and culture.
[52] Resig, "Using Computer Vision to Increase the Research Potential of Photo Archives."
[53] Wittgenstein, Philosophical Investigations. §67; Rosch and Mervis, "Family Resemblances."
[54] Krizhevsky, Sutskever, and Hinton, "ImageNet Classification with Deep Convolutional Neural Networks."
[55] Freeman, "The Development of Social Network Analysis."
[56] Bertin, Semiology of Graphics; Brilliantly summarized in Bertin, "Matrix Theory of Graphics."
[57] Djindjian, "Fifteen Years of Contributions of the French School of Data Analysis to Quantitative Archaeology," (published 35 years ago).
[58] Mukherjea, Hirata, and Hara, "Visualizing the Results of Multimedia Web Search Engines;" Manovich, "Software Studies: ImagePlot Visualization Software: Explore Patterns in Large Image Collections."
[59] Dilly, "Der kunsthistorische Nachthimmel;" Schich, Rezeption und Tradierung als

komplexes Netzwerk.
[60] Milgram and Jodelet, "Psychological Maps of Paris;" Searls, "With a Wild Surmise: Intimations of Computational Biology in Keats, Carroll, and Joyce."
[61] Agarwal et al., "Building Rome in a Day;" Crandall et al., "Mapping the World's Photos."
[62] Lee et al., "Linking Past to Present."
[63] Cheng et al., "Can Cascades Be Predicted?"
[64] "Fashion Trends for Spring 2015 as Told by Google Data."
[65] Kim, Son, and Jeong, "Large-Scale Quantitative Analysis of Painting Arts;" compares well to Rosa and Suárez, "A Quantitative Approach to Beauty. Perceived Attractiveness of Human Faces in World Painting."
[66] Palla et al., "Hierarchical Networks of Scientific Journals."
[67] Manyika et al., "Big Data."

Bibliography

Adamo, Bélen and Santiago Ortiz: "Leonardo Journal cover Page." Leonardo Journal 47, no. 3 (2014). http://www.mitpressjournals.org/toc/leon/47/3.

Adams, Douglas. Parrots, the Universe and Everything. Vol. 5779. UCTV - University of California Television, 2001. http://www.uctv.tv/shows/Douglas-Adams-Parrots-the-Universe-and-Everything-5779.

Agarwal, Sameer, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M. Seitz, and Richard Szeliski. "Building Rome in a Day." Commun. ACM 54, no. 10 (October 2011): 105–12. doi:10.1145/2001269.2001293.

Alon, Uri. An Introduction to Systems Biology: Design Principles of Biological Circuits. Boca Raton, FL: Chapman & Hall/CRC, 2007.

———. "How to Build a Motivated Research Group." Molecular Cell 37, no. 2 (January 29, 2010): 151–52. doi:10.1016/j.molcel.2010.01.011.

———. "How To Choose a Good Scientific Problem." Molecular Cell 35, no. 6 (September 24, 2009): 726–28. doi:10.1016/j.molcel.2009.09.013.

Anderson, P. W. "More Is Different." Science 177, no. 4047 (August 4, 1972): 393–96. doi:10.1126/science.177.4047.393.

"Art Project - Google Cultural Institute." Accessed October 5, 2015. https://www.google.com/culturalinstitute/project/art-project.

Aubert, M., A. Brumm, M. Ramli, T. Sutikna, E. W. Saptomo, B. Hakim, M. J. Morwood, G. D. van den Bergh, L. Kinsley, and A. Dosseto. "Pleistocene Cave Art from Sulawesi, Indonesia." Nature 514, no. 7521 (October 9, 2014): 223–27. doi:10.1038/nature13422.

Barabási, Albert-László. Network Science Book, Cambridge University Press, 2016. http://barabasi.com/networksciencebook/.

Baronchelli, Andrea, Tao Gong, Andrea Pugliesi, Vittorio Loreto. "Modeling the Emergence of Universality in Color Naming Patterns." PNAS 107, no. 6 (February 9, 2010): 2403-2407. doi:10.1073/pnas.0908533107.

Belting, Hans, Heinrich Dilly, Wolfgang Kemp, Willibald Sauerländer, and Martin Warnke. Kunstgeschichte: Eine Einführung. 7th ed. Reimer, Dietrich, 2008.

Bertin, Jacques. “Matrix Theory of Graphics.” Information Design Journal 10, no. 1 (2001): 5–19.
Bertin, Jacques, and William J Berg. Semiology of Graphics: Diagrams, Networks, Maps. Redlands, Calif.: ESRI Press : Distributed by Ingram Publisher Services, 2011.
Boccaletti, Stefano, G. Bianconi, R. Criado, Charo I. Del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. “The Structure and Dynamics of Multilayer Networks.” Physics Reports 544, no. 1 (2014): 1–122.
Bollacker, Kurt, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. “Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge.” In Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data, 1247–50. ACM, 2008. http://dl.acm.org/citation.cfm?id=1376746.
Bredekamp, Horst. Darwins Korallen. Berlin: Wagenbach Klaus GmbH, 2005.
Burdick, Anne, Johanna Drucker, Peter Lunenfeld, Todd Presner, and Jeffrey Schnapp. Digital_Humanities. Cambridge, MA: The MIT Press, 2012.
Cheng, Justin, Lada Adamic, P. Alex Dow, Jon Michael Kleinberg, and Jure Leskovec. “Can Cascades Be Predicted?” In Proceedings of the 23rd International Conference on World Wide Web, 925–36. WWW ’14. New York, NY, USA: ACM, 2014. doi:10.1145/2566486.2567997.
Chomsky, Noam. “Rationality/Science.” Z Papers Special Issue, 1995. http://www.chomsky.info/articles/1995----02.htm.
Christian, David. Maps of Time. An Introduction to Big History, Berkeley, Calif.: University of California Press, 2011.
Chua, Leon O. “Local Activity Is the Origin of Complexity.” International Journal of Bifurcation and Chaos 15, no. 11 (November 1, 2005): 3435–56. doi:10.1142/S0218127405014337.
Cicero, Marcus Tullius. “On Invention.” Wikisource. Accessed October 7, 2015. https://en.wikisource.org/wiki/On_invention.
Clauset, A., C. Shalizi, and M. Newman. “Power-Law Distributions in Empirical Data.” SIAM Review 51, no. 4 (November 4, 2009): 661–703. doi:10.1137/070710111.
Cortjaens, Wolfgang, and Karsten Heck, eds. Stil-Linien Diagrammatischer Kunstgeschichte. Berlin: Deutscher Kunstverlag, 2014.
Crandall, David J., Lars Backstrom, Daniel Huttenlocher, and Jon Kleinberg. “Mapping the World’s Photos.” In Proceedings of the 18th International Conference on World Wide Web, 761–70. WWW ’09. New York, NY, USA: ACM, 2009. doi:10.1145/1526709.1526812.
“Culture Analytics.” IPAM. Accessed October 7, 2015. http://www.ipam.ucla.edu/programs/long-programs/culture-analytics/.
Dilly, Heinrich. “Der kunsthistorische Nachthimmel.” In Frankfurter Schule und Kunstgeschichte, edited by Andreas Berndt, 69–84. Berlin: Reimer, 1992.
Djindjian, F. “Fifteen Years of Contributions of the French School of Data Analysis to Quantitative Archaeology,” 193–204. Oxford: B.A.R., 1989. http://proceedings.caaconference.org/files/1989/19_Djindjian_CAA_1989.pdf.
Dodds, Peter. “Course Home | Matrixology, Season 8.” Accessed October 7, 2015. http://

www.uvm.edu/~pdodds/teaching/courses/2015-01UVM-124/.
Doing, Park, and Jr C. Richard Johnson. "On Applying Signal Processing to Computational Art History: An Interview." International Journal for Digital Art History 0, no. 1 (June 26, 2015). https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21637.
Duchamp, Marcel, Ecke Bonk, Richard Hamilton, Jackie Matisse, and Typosophic Society. À L' Infinitif = In the Infinitive. Köln: König, 1999.
Eames, Charles, and Ray Eames. 1977. Powers of ten. Santa Monica, CA: Pyramid Films.
Eigen, Manfred, and P. Schuster. 1979. The hypercycle, a principle of natural self-organization. Berlin: Springer-Verlag.
Estrada, Ernesto, and Philip Knight. A First Course in Network Theory. Oxford University Press, 2015.
"Facebook Has a Quarter of a Trillion User Photos." Accessed October 6, 2015. http://mashable.com/2013/09/16/facebook-photo-uploads/#OzS7o.Hu4EqO.
"Fashion Trends for Spring 2015 as Told by Google Data." Think with Google. Accessed October 9, 2015. https://www.thinkwithgoogle.com/articles/spring-2015-fashion-trends-google-data.html.
Feyerabend, Paul. Against Method. London; New York: Verso, 1993.
Feynman, Richard P., Robert B. Leighton, and Matthew Sands. Six Easy Pieces: Essentials of Physics Explained by Its Most Brilliant Teacher. Fourth Edition. Basic Books, 2011.
Freeman, Linton C. 2004. The development of social network analysis: a study in the sociology of science. Vancouver, BC: Empirical Press.
Gaehtgens, Thomas W. "Thoughts on the Digital Future of the Humanities and Art History." Visual Resources 29, no. 1–2 (June 1, 2013): 22–25. doi:10.1080/01973762.2013.761110.
Gombrich, E. H. 1979. The sense of order: a study in the psychology of decorative art. Ithaca, N.Y.: Cornell University Press.
"Google Books." Accessed October 5, 2015. https://books.google.com/.
"Google Earth." Accessed October 5, 2015. http://www.google.com/earth/.
"Google Images." Accessed October 5, 2015. https://images.google.com/?gws_rd=ssl.
Helbing, Dirk. "Globally Networked Risks and How to Respond." Nature 497, no. 7447 (May 2, 2013): 51–59. doi:10.1038/nature12047.
"Is the LHC Throwing Away Too Much Data? | New Scientist." Accessed October 7, 2015. https://www.newscientist.com/article/mg21328564-700-is-the-lhc-throwing-away-too-much-data/.
Jacobs, Jane. The Death and Life of Great American Cities. Vintage Books, 1961.
Johansen, Anders, and Didier Sornette. "Finite-Time Singularity in the Dynamics of the World Population, Economic and Financial Indices." Physica A: Statistical Mechanics and Its Applications 294, no. 3–4 (May 15, 2001): 465–502. doi:10.1016/S0378-4371(01)00105-4.
Kim, Daniel, Seung-Woo Son, and Hawoong Jeong. "Large-Scale Quantitative Analysis

of Painting Arts." Scientific Reports 4 (December 11, 2014): 7370. doi:10.1038/srep07370.

Kivelä, Mikko, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. "Multilayer Networks." Journal of Complex Networks 2, no. 3 (September 1, 2014): 203–71. doi:10.1093/comnet/cnu016.

Koolhaas, Rem, and Office for Metropolitan Architecture. Content. Köln: Taschen, 2004.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks." In Advances in Neural Information Processing Systems 25, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, 1097–1105. Curran Associates, Inc., 2012. http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

Kubler, George. The Shape of Time: Remarks on the History of Things. Rev ed. New Haven Conn.: Yale University Press, 2008.

Kunze, Max, and Peter Betthausen. 2002. Corpus der antiken Denkmäler, die J.J. Winckelmann und seine Zeit kannten. München: Projekt Dyabola.

Lee, Stefan, Nicolas Maisonneuve, David Crandall, Alexei A. Efros, and Josef Sivic. "Linking Past to Present: Discovering Style in Two Centuries of Architecture," 2015. https://hal.inria.fr/hal-01152482/document.

Leskovec, Jure, Anand Rajaraman, and Jeffrey David Ullman. Mining of Massive Datasets. Cambridge University Press, 2014.

Malina, Roger "Yes Again to the End of the Digital Humanities! Please!" Accessed October 5, 2015. http://malina.diatrope.com/2015/08/06/yes-again-to-the-end-of-the-digital-humanities-please/.

Malraux, André. Le Musée imaginaire de la sculpture mondiale. Paris: N.R.F., 1952.

Manovich, Lev. "Data Science and Digital Art History." International Journal for Digital Art History 0, no. 1 (June 26, 2015). https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21631.

———. "Software Studies: ImagePlot Visualization Software: Explore Patterns in Large Image Collections." Accessed October 9, 2015. http://lab.softwarestudies.com/p/imageplot.html#features1.

Manovich, Lev, Moritz Stefaner, Mehrdad Yazdani, Dominius Baur, Daniel Goddemeyer, Alise Tifentale, Nadav Hochman, Jay Chow, "Selfiecity." Accessed October 6, 2015. http://selfiecity.net/#dataset.

Manovich, Lev, Jeremy Douglass, and William Huber. "Understanding Scanlation: How to Read One Million Fan-Translated Manga Pages." Image & Narrative 12, no. 1 (2011): 206–28.

Manyika, James, Michael Chui, Brad Brown, Jacques Bughin, Richard Dobbs, Charles Roxburgh, and Angela Hung Byers. "Big Data: The next Frontier for Innovation, Competition, and Productivity | McKinsey & Company." McKinsey Global Institute, May 2011. http://www.mckinsey.com/insights/business_technology/big_data_the_next_frontier_for_innovation.

Mayer-Schönberger, Viktor, and Kenneth Cukier. Big Data: A Revolution That Will

Transform How We Live, Work, and Think. Houghton Mifflin Harcourt, 2013.
McGee, Harold. On Food and Cooking: The Science and Lore of the Kitchen. Rev Upd edition. New York: Scribner, 2004.
Meyer, Miriah. "Designing Visualizations For Biological Data." Leonardo 46, no. 3 (March 20, 2013): 270–71. doi:10.1162/LEON_a_00568.
Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, et al. "Quantitative Analysis of Culture Using Millions of Digitized Books." Science 331, no. 6014 (January 14, 2011): 176–82. doi:10.1126/science.1199644.
Milgram, Stanley, and Denise Jodelet. "Psychological Maps of Paris." In Environmental Psychology: People and Their Physical Settings, edited by Harold M. Proshansky, William H. Ittelson, and Leanne G. Rivlin, 104–24. New York: Holt, Rinehart and Winston, 1976.
Mitchell, Melanie. Complexity: A Guided Tour. 1 edition. Oxford University Press, 2011.
"MoMA Makes a Facebook for Abstractionists | ARTnews." Accessed October 7, 2015. http://www.artnews.com/2012/10/02/momaabstractionfaceboo/.
Moretti, Franco, and Alberto Piazza. Graphs, Maps, Trees: Abstract Models for Literary History. Verso, 2007.
Mukherjea, Sougata, Kyoji Hirata, and Yoshinori Hara. "Visualizing the Results of Multimedia Web Search Engines." In In Proceedings IEEE Symposium on Information Visualization, 64–65, 1996.
Newman, M. E. J. "Power Laws, Pareto Distributions and Zipf's Law." Contemporary Physics 46, no. 5 (September 1, 2005): 323–51. doi:10.1080/00107510500052444.
"Our Approach | Broad Institute of MIT and Harvard." Accessed October 5, 2015. http://www.broadinstitute.org/what-broad/our-approach/our-approach.
Palla, Gergely, Gergely Tibély, Enys Mones, Péter Pollner, and Tamás Vicsek. "Hierarchical Networks of Scientific Journals." Palgrave Communications 1 (July 14, 2015): 15016. doi:10.1057/palcomms.2015.16.
"Physics and Astronomy Classification Scheme® (PACS)." Accessed October 5, 2015. http://journals.aps.org/PACS.
Ramsay, Steven, Laurent Romary, Kari Kraus, Maximilian Schich, Desmond Schmidt, Andrew Ashton, Julia Flanders,, and Fotis Jannidis, Panel Discussion, Brown University. KODM 2012 Day 1 Panel Discussion: Data Models in Humanities Theory and Practice, 2012. https://www.youtube.com/watch?v=lHJmPT-VjPE.
Resig, John. "Using Computer Vision to Increase the Research Potential of Photo Archives." Journal of Digital Humanities 3, no. 2 (2014). http://journalofdigitalhumanities.org/3-2/using-computer-vision-to-increase-the-research-potential-of-photo-archives-by-john-resig/.
Restrepo, Buitrago, Pedro Felipe, and Iván Duque Márquez. "The Orange Economy: An Infinite Opportunity," October 1, 2013. http://publications.iadb.org/handle/11319/3659.
Rosa, Javier de la, and Juan-Luis Suárez. "A Quantitative Approach to Beauty. Perceived

Attractiveness of Human Faces in World Painting." International Journal for Digital Art History 0, no. 1 (June 26, 2015). https://journals.ub.uni-heidelberg.de/index.php/dah/article/view/21640.

Rosch, Eleanor, and Carolyn B Mervis. "Family Resemblances: Studies in the Internal Structure of Categories." Cognitive Psychology 7, no. 4 (October 1, 1975): 573–605. doi:10.1016/0010-0285(75)90024-9.

Schich, Maximilian. "Netzwerke von Komplexen Netzwerken in Der (Kunst) Wissenschaft." In Die Dynamik Sozialer Und Sprachlicher Netzwerke, 161–78. Springer Fachmedien Wiesbaden, 2013. http://link.springer.com/chapter/10.1007/978-3-531-93336-8_9.

———. "Revealing Matrices." In Beautiful Visualization: [looking at Data through the Eyes of Experts], edited by Noah P. N Iliinsky and Julie Steele. Sebastopol, CA: O'Reilly, 2010. https://library.oreilly.com/book/0636920000617/beautiful-visualization/104.xhtml?ref=toc.

———. Rezeption und Tradierung als komplexes Netzwerk: der CENSUS und visuelle Dokumente zu den Thermen in Rom. Munich: Biering & Brinkmann Verlag, 2009.

Schich, Maximilian, Chaoming Song, Yong-Yeol Ahn, Alexander Mirsky, Mauro Martino, Albert-László Barabási, and Dirk Helbing. "A Network Framework of Cultural History." Science 345, no. 6196 (August 1, 2014): 558–62. doi:10.1126/science.1240064.

Schich, Maximilian, and Mauro Martino (animation), Kerri Smith (narration), and Tristan Perich (music): Charting Culture. Nature video (July 31, 2014). https://www.youtube.com/watch?v=4gIhRkCcD4U.

Schich, Maximilian, Isabel Meirelles, and Malina, Roger, eds. Arts, Humanities, and Complex Networks. 4th edition. Leonardo/ISAST and MIT Press, 2014.

Schich, Maximilian, and Michele Coscia. "Exploring Co-Occurrence on a Meso and Global Level Using Network Analysis and Rule Mining." In MLG'11 Proceedings of the Ninth Workshop on Mining and Learning with Graphs at ACM KDD 2011. http://www.cs.purdue.edu/mlg2011/papers/paper_22.pdf.

Schmidt-Burkhardt, Astrit. Stammbäume der Kunst: Zur Genealogie der Avantgarde. Berlin: Oldenbourg Akademieverlag, 2005.

Searls, David B. "With a Wild Surmise: Intimations of Computational Biology in Keats, Carroll, and Joyce." presented at Shared Horizons: Data, Biomedicine, and the Digital Humanities, University of Maryland, College Park, MD, April 12, 2013. http://mith.umd.edu/sharedhorizons/about/keynote-lecture/attachment/sharedhorizons_searls/.

Settis, Salvatore. "L'opera Di Paul Zanker E Il Futuro Dell'archaeologia Classica." Pegasus 4 (2003). https://edoc.bbaw.de/frontdoor/index/index/docId/917.

Snow, C. P., and Stefan Collini. The Two Cultures. Cambridge University Press, 2012.

Steffen, Will, Regina Angelina Sanderson, Peter D. Tyson, Jill Jäger, Pamela A. Matson, Berrien Moore III, Frank Oldfield, et al. Global Change and the Earth System: A Planet Under Pressure. Berlin ; New York: Springer, 2005.

"The Prometheus Image Archive: High-Quality Images from the Fields of Arts, Culture and History." Accessed October 7, 2015. http://www.prometheus-bildarchiv.de/.

Venturi, Robert, Denise Scott Brown, and Steven Izenour. Learning from Las Vegas: The Forgotten Symbolism of Architectural Form. Cambridge, Mass.: MIT Press, 1977.

Vitruvius. On Architecture, Volume I, Books 1-5. Translated by Frank Granger. Harvard University Press, 1931.

Warburg, Aby, Martin Warnke, Claudia Birnk, and Aby Warburg. Der Bilderatlas Mnemosyne. Berlin: Akademie Verlag, 2008.

Weaver, Warren. "Science and Complexity." American Scientist 36 (1948): 536–67.

Whiteley, J. J. L. Index to the Paris Salon Catalogues. New York: Taylor & Francis, 1988.

Wigley, Mark. "Network Fever." Grey Room - (June 1, 2001): 82–122. doi:10.1162/152638101750420825.

Windelband, Wilhelm. Geschichte und Naturwissenschaft: Rede zum Antritt des Rectorats der Kaiser-Wilhelms-Universität Strassburg, geh. am 1. Mai 1894. Strassburg: Heitz, 1894. http://digi.ub.uni-heidelberg.de/diglit/windelband1894.

Wittgenstein, Ludwig. Philosophical Investigations: The German Text, with a Revised English Translation 50th Anniversary Commemorative Edition. Wiley, 1991.

Wölfflin, Heinrich. Principles of Art History: The Problem Of The Development Of Style In Later Art. Translated by Marie Donald Mackie Hottinger. Literary Licensing, LLC, 2012.

Young, H. Peyton. "The Evolution of Social Norms." Annual Review of Economics 7, no. 1 (2015): 359–87. doi:10.1146/annurev-economics-080614-115322.

Maximilian Schich is an Associate Professor in Arts & Technology and a founding member of EODIAH, the Edith O'Donnell Institute of Art History at the University of Texas at Dallas. In summer 2015, he also was a Visiting Scientist at ETH Zurich in Dirk Helbing's Computational Social Science group, where he wrote parts of this article.

He is the first author of A Network Framework of Cultural History (Science Magazine, 2014) and a lead co-author of the animation Charting Culture (Nature video, 2014). He has visualized networks of complex networks in art research (O'Reilly 2010), and analyzed antique reception and visual citation as complex networks (Biering & Brinkmann, 2009). He is an Editorial Advisor for Arts, Humanities, and Complex Networks at Leonardo Journal, and an Editorial Board member at DAH-Journal and Palgrave Communications, the new humanities and social science equivalent of Nature Communications. He has been invited to SciFoo, DLD*, Edge.org, and the Lincoln Center Global Exchange. His most recent work received global press coverage in 28 languages.

Correspondence e-mail: maximilian@schich.info

Showing Digitized Corpora

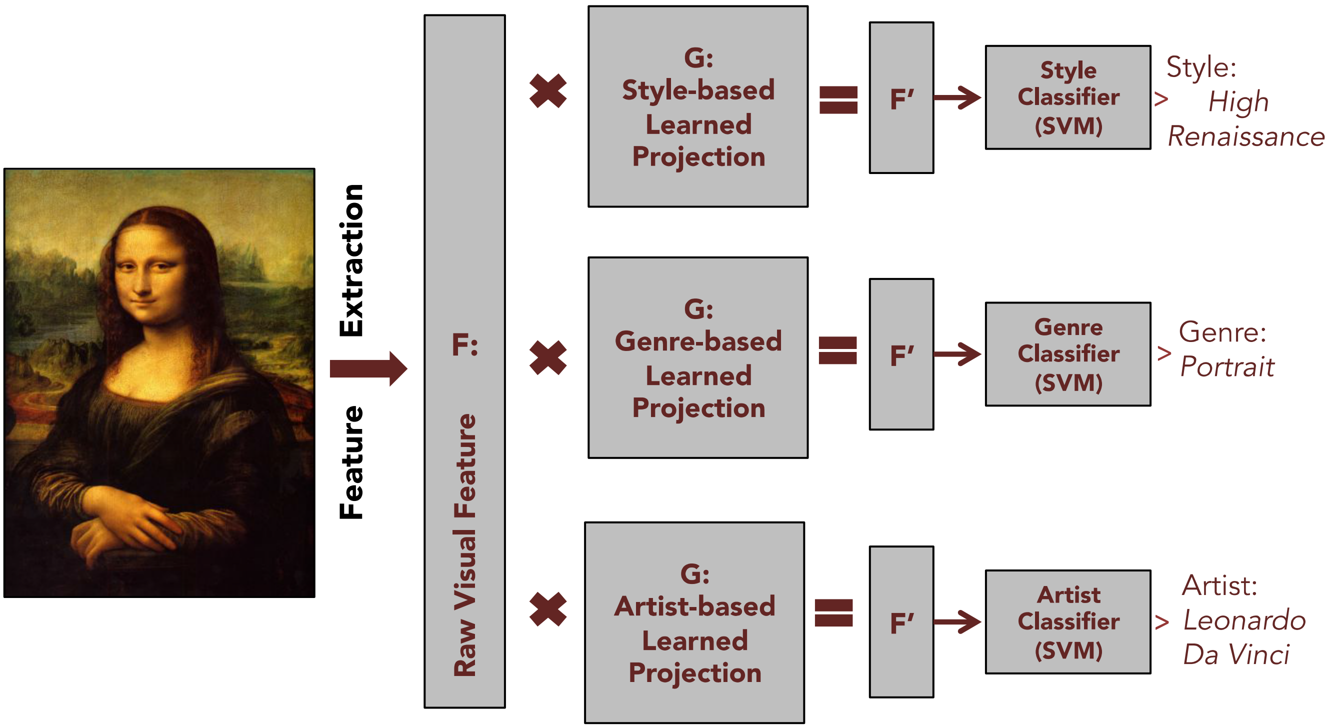

Figure 1: Illustration of our system for classification of fine-art paintings. We investigated variety of visual features and metric learning approaches to recognize Style, Genre and Artist of a painting.

Peer-Reviewed

Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature

Babak Saleh, Ahmed Elgammal

Abstract: In the past few years, the number of fine-art collections that are digitized and publicly available has been growing rapidly. With the availability of such large collections of digitized artworks comes the need to develop multimedia systems to archive and retrieve this pool of data. Measuring the visual similarity between artistic items is an essential step for such multimedia systems, which can benefit more high-level multimedia tasks. In order to model this similarity between paintings, we should extract the appropriate visual features for paintings and find out the best approach to learn the similarity metric based on these features. We investigate a comprehensive list of visual features and metric learning approaches to learn an optimized similarity measure between paintings. We develop a machine that is able to make aesthetic-related semantic-level judgments, such as predicting a painting's style, genre, and artist, as well as providing similarity measures optimized based on the knowledge available in the domain of art historical interpretation. Our experiments show the value of using this similarity measure for the aforementioned prediction tasks.

Keywords: similarity metric, visual features, metric learning, convolutional neural networks, style, genre, artist

1 Introduction

In the past few years, the number of fine-art collections that are digitized and publicly available has been growing rapidly. Such collections span classical and modern and contemporary artworks. With the availability of such large collections of digitized artworks comes the need to develop multimedia systems to archive and retrieve this pool of data. Typically, these collections, in particular early modern ones, come with metadata in the form of annotations by art historians and curators, including information about each painting's artist, style, date, genre, etc. For online galleries displaying contemporary artwork, there is a need to develop automated recommendation systems that can retrieve "similar" paintings that the user might like to buy. This highlights the need to investigate metrics of visual similarity among digitized paintings that are optimized for the domain of painting.

The field of computer vision has made significant leaps in getting digital systems to recognize and categorize objects and scenes in images and videos. These advances have been driven by a wide spread need for the technology, since cameras are everywhere now. However, a person looking at a painting can make sophisticated inferences beyond just recognizing a tree, a chair, or the figure of Christ. Even individuals without specific art historical training can make assumptions about a painting's genre (portrait or landscape), its style (impressionist or abstract), what century it was created, the artists who likely created the work and so on. Obviously, the accuracy of such assumptions depends on the viewer's level of knowledge and exposure to art history. Learning and judging such complex visual concepts is an impressive ability of human perception [2].

The ultimate goal of our research is to develop a machine that is able to make aesthetic-related semantic-level judgments, such as predicting a painting's style, genre, and artist, as well as providing similarity measures optimized based on the knowledge available in the domain of art historical interpretation. Immediate questions that arise include, but are not limited to: What visual features should be used to encode information in images of paintings? How does one weigh different visual features to achieve a useful similarity measure? What type of art historical knowledge should be used to optimize such similarity measures? In this paper we address these questions and aim to provide answers that can benefit researchers in the area of computer-based analysis of art. Our work is based on a systematic methodology and a comprehensive evaluation on one of the largest available digitized art datasets.

Artists use different concepts to describe paintings. In particular, stylistic elements, such as space, texture, form, shape, color, tone and line are used. Other principles include movement, unity, harmony, variety, balance, contrast, proportion, and pattern. To this might be added physical attributes, like brush strokes as well as subject matter and other descriptive concepts [13].

For the task of computer analyses of art, researchers have engineered and investigated various visual features that encode some of these artistic concepts, in particular brush strokes and color, which are encoded as low-level features such as texture statistics and color histograms (e.g. [19, 20]). Color and texture are highly prone to variations

during the digitization of paintings; color is also affected by a painting's age. The effect of digitization on the computational analysis of paintings is investigated in great depth by Polatkan et al. [24]. This highlights the need to carefully design visual features that are suitable for the analysis of paintings.

Clearly, it would be a cumbersome process to engineer visual features that encode all the aforementioned artistic concepts. Recent advances in computer vision, using deep neural networks, showed the advantage of "learning" the features from data instead of engineering such features. However, it would also be impractical to learn visual features that encode such artistic concepts, since that would require extensive annotation of these concepts in each image within a large training and testing dataset. Obtaining such annotations require expertise in the field of art history that can not be achieved with typical crowed-sourcing annotators.

Given the aforementioned challenges to engineering or learning suitable visual features for painting, in this paper we follow an alternative strategy. We mainly investigate different state-of-the-art visual elements, ranging from low-level elements to semantic-level elements. We then use metric learning to achieve optimal similarity metrics between paintings that are optimized for specific prediction tasks, namely style, genre, and artist classification. We chose these tasks to optimize and evaluate the metrics since, ultimately, the goal of any art recommendation system would be to retrieve artworks that are similar along the directions of these high-level semantic concepts. Moreover, annotations for these tasks are widely available and more often agreed-upon by art historians and critics, which facilitates training and testing the metrics.

In this paper we investigate a large space of visual features and learning methodologies for the aforementioned prediction tasks. We propose and compare three learning methodologies to optimize such tasks. We present results of a comprehensive comparative study that spans four state-of-the-art visual features, five metric learning approaches and the proposed three learning methodologies, evaluated on the aforementioned three artistic prediction tasks.

2 Related Work

On the subject of painting, computers have been used for a diverse set of tasks. Traditionally, image processing techniques have been used to provide art historians with quantification tools, such as pigmentation analysis, statistical quantification of brush strokes, etc. We refer the reader to [28, 5] for comprehensive surveys on this subject.

Several studies have addressed the question of which features should be used to encode information in paintings. Most of the research concerning the classification of paintings utilizes low-level features encoding color, shadow, texture, and edges. For example Lombardi [20] has presented a study of the performance of these types of features for the task of artist classification among a small set of artists using several supervised and unsupervised learning methodologies. In that paper the style of the painting was identified as a result of recognizing the artist.

Since brushstrokes provide a signature that can help identify the artist, designing visual features that encode brushstrokes has been widely adapted. (e.g. [25, 18, 22, 15, 6, 19]). Typically, texture statistics are used for that purpose. However, as mentioned earlier, texture features are highly affected by the digitization resolution. Researchers also investigated the use of features based on local edge orientation histograms, such as SIFT [21] and HOG [10]. For example, [12] used SIFT features within a Bag-of-words pipeline to discriminate among a set of eight artists.

Arora et al. [3] presented a comparative study for the task of style classification, which evaluated low-level features, such as SIFT and Color SIFT [1], versus semantic-level features, namely Classemes [29], which encodes object presence in the image. It was found that semantic-level features significantly outperform low-level features for this task. However, the evaluation was conducted on a small dataset of 7 styles, with 70 paintings in each style. Carneiro et al [9] also concluded that low-level texture and color features are not effective because of inconsistent color and texture patterns that describe the visual classes in paintings.

More recently, Saleh et al [26] used metric learning approaches for finding influence paths between painters based on their paintings. They evaluated three metric learning approaches to optimize a metric over low-level HOG features. In contrast to that work, the evaluation presented in this paper is much wider in scope since we address three tasks (style, genre and artist prediction), we cover features spanning from low-level to semantic-level and we evaluate five metric learning approaches. Moreover, the dataset of [26] has only 1710 images from 66 artists, while we conducted our experiments on 81,449 images painted by 1119 artists. Bar et al [4] proposed an approach for style classification based on features obtained from a convolution neural network pre-trained on an image categorization task. In contrast we show that we can achieve better results with much lower dimensional features that are directly optimized for style and genre classification. Lower dimensionality of the features is preferred for indexing large image collections.

3 Methodology

In this section we explain the methodology that we follow to find the most appropriate combination of visual features and metrics that produce accurate similarity measurements. We acquire these measurements to mimic the art historian's ability to categorize paintings based on their style, genre and the artist who made it. In the first step, we extract visual features from the image. These visual features range from low-level (e.g. edges) to high-level (e.g. objects in the painting). More importantly, in the next step we learn how to adjust these features for different classification tasks by learning the appropriate metrics. Given the learned metric we are able to project paintings from a high dimensional space of raw visual information to a meaningful space with much lower dimensionality. Additionally, learning a classifier in this low-dimensional space can be easily scaled up for large collections.

In the rest of this section: First, we introduce our collection of fine-art paintings and explain what are the tasks that we target in this work. Later, we explore methodologies that we consider in this work to find the most accurate system for aforementioned tasks. Finally, we explain different types of visual features that we use to represent images of paintings and discuss metric learning approaches that we applied to find the proper notion of similarity between paintings.

3.1 Dataset and Proposed Tasks

In order to gather our collection of fine-art paintings, we used the publicly available dataset of "Wikiart paintings"; which, to the best of our knowledge, is the largest online public collection of digitized artworks. This collection has images of 81,449 fine-art paintings from 1,119 artists ranging from fifteen centuries to contemporary artists. These paintings are from 27 different styles (Abstract, Byzantine, Baroque, etc.) and 45 different genres (Interior, Landscape, etc.) Previous work [26, 9] used different resources and made smaller collections with limited variability in terms of style, genre and artists. The work of [4] is the closest to our work in terms of data collection procedure, but the number of images in their collection is half of ours.

We target automatic classification of paintings based on their style, genre and artist using visual features that are automatically extracted using computer vision algorithms. Each of these tasks has its own challenges and limitations. For example, there are large variations in terms of visual appearances in paintings from one specific style. However, this variation is much more limited for paintings by one artist. These larger intra-class variations suggest that style classification based on visual features is more challenging than artist classification. For each of the tasks we selected a subset of the data that ensure enough samples for training and testing. In particular, for style classification we use a subset of the date with 27 styles where each style has at least 1500 paintings with no restriction on genre or artists, with a total of 78,449 images. For genre classification we use a subset with 10 genre classes, where each genre has at least 1500 paintings with no restriction of style or genre, with a total of 63,691 images. Similarly, for artist classification we use a subset of 23 artists, where each of them has at least 500 paintings, with a total of 18,599 images. Table 1 lists the set of style, genre, and artist labels.

3.2 Classification Methodology

In order to classify paintings based on their style, genre or artist we followed three methodologies.

Metric Learning: First, as depicted in figure 1, we extract visual features from images of paintings. For each of these prediction tasks, we learn a similarity metric optimized for it, i.e. style-optimized metric, genre-optimized metric and artist-optimized metric. Each metric induces a projector to a corresponding feature space optimized for the corresponding task. Having the metric learned, we project the raw visual features into the new optimized feature space and learn classifiers for the corresponding prediction task. For that purpose, we learn a set of one-vs-all SVM classifiers for each of the labels in table 1 for each of the tasks.

While our first strategy focuses on classification based on combinations of a metric and a visual feature, the next two methodologies that we followed fuse different features or different metrics.

Feature fusion: The second methodology that we used for classification is depicted in figure 2. In this case, we extract different types of visual

Figure 2: illustration of our second methodology – Feature Fusion.

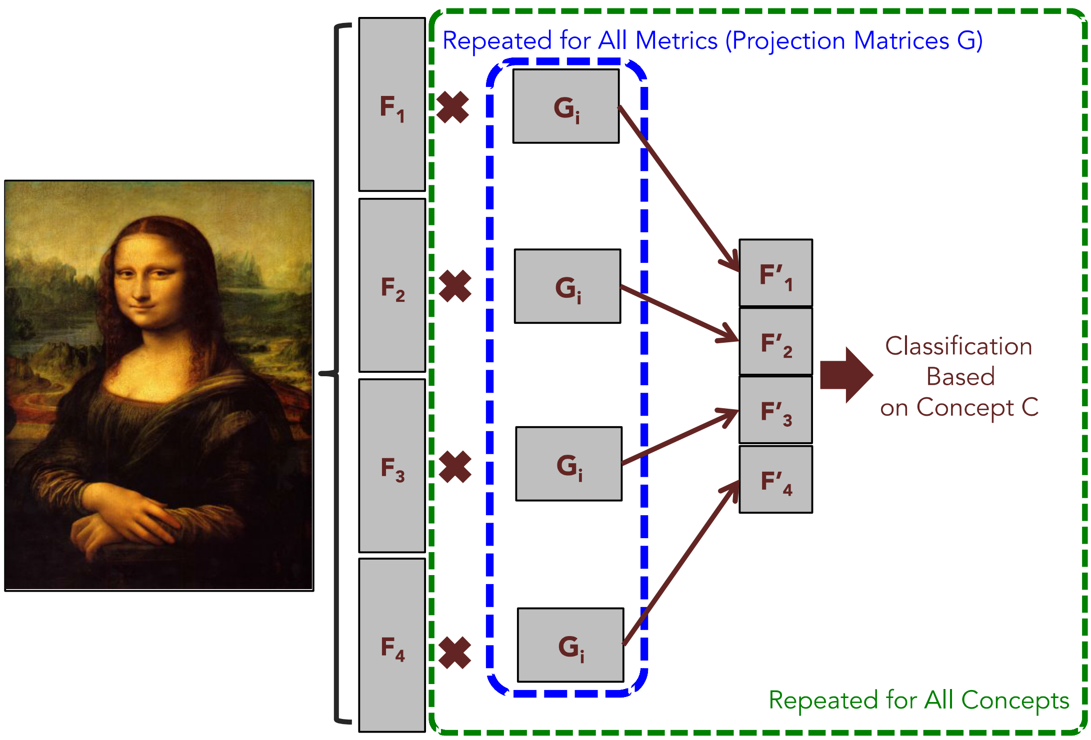

features (four types of features as will be explained next). Based on the prediction task (e.g. style) we learn the metric for each type of feature as before. After projecting these features separately, we concatenate them to make the final feature vector. The classification will be based on training classifiers using these final features. This feature fusion is important as we want to capture different types of visual information by using different types of features. Also concatenating all features together and learn a metric on top of this huge feature vector is computationally intractable. Because of this issue, we learn metrics on feature separately and after projecting features by these metrics, we can concatenate them for classification purposes.

Metric-fusion: The third methodology (figure 3) projects each visual features using multiple metrics (in our experiment we used five metrics as will be explained next) and then fuses the resulting optimized feature spaces to obtain a final feature vector for classification. This is an important strategy, because each one of the metric learning approaches use a different criterion to learn the similarity measurement. By learning all metrics individually (on the same type of feature), we make sure that we took into account all criteria (e.g. information theory along with neighborhood analysis).

Figure 3: Illustration of our third methodology – Metric Fusion.

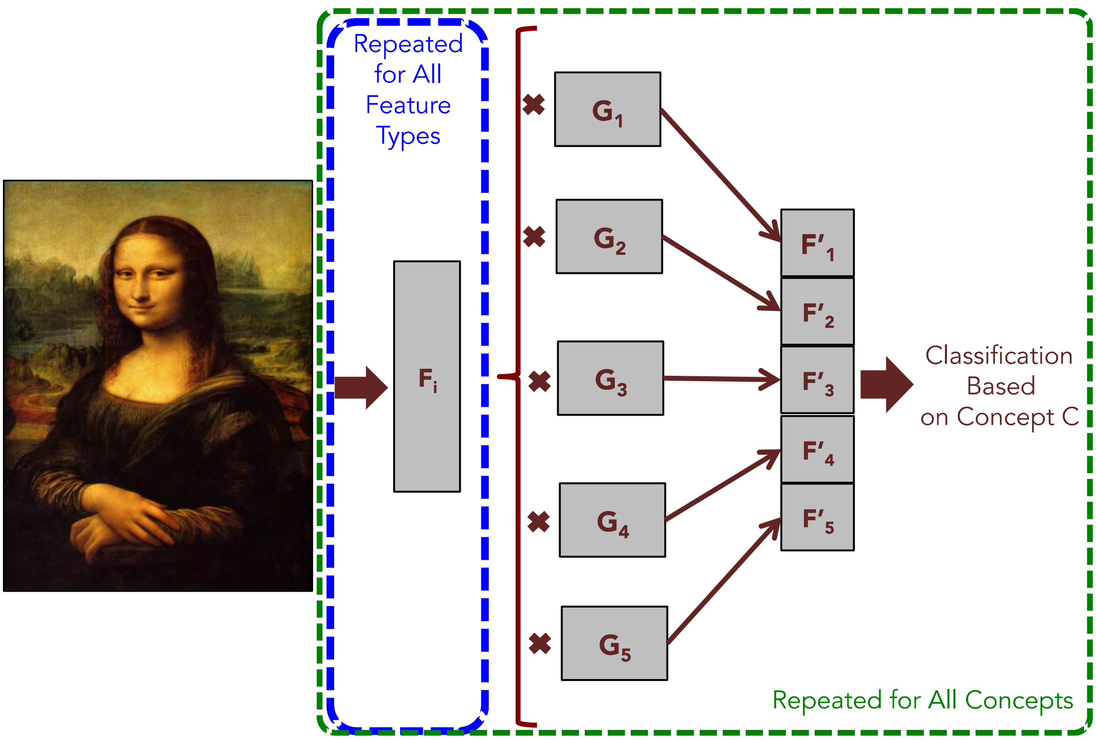

3.3 Visual Features

Visual features in computer vision literature are either engineered and extracted in an unsupervised way (e.g. HOG, GIST) or learned based on optimizing a specific task, typically categorization of objects or scenes (e.g. CNN-based features). This results in high-dimensional feature vectors that might not necessary correspond to nameable (semantic-level) characteristics of an image. Based on the ability to find a meaning, visual features can be categorized into low-level and high-level. Low-level features are visual descriptors that there is no explicit meaning for each dimension of them, while high-level visual features are designed to capture some notions (usually objects). For this work, we investigated some state-of-the-art representatives of these two categories:

Low-level Features, GIST: Human observers can rapidly capture the "gist" of a scene in a quick feed-forward sweep. Therefore, a computational model for "gist" seems a reasonably essential tool for rapid scene classification. Gist has been modelled as average pooling of low-level biologically-inspired features (i.e. gabor-like features) over non-overlapping subregions arranged on a fixed grid. The term "spatial envelope" has been also used to refer to this very low dimensional representation of the [23]. Indeed, gist model bypasses the procedures that are usually applied in scene classification, such as segmentation and processing of individual objects.

The dominant spatial structure of a scene is represented in a set of perceptual dimensions (naturalness, openness, roughness, expansion, ruggedness). The gist model estimates these dimensions using spectral and coarsely localized information. To calculate the gist features, each image is divided into 16 bins, and then oriented Gabor filters (in 8 orientations) are applied over different scales (4 scales) in each bin. Finally, the average filter energy in each bin is calculated [24]. We followed this procedure and extracted 512-dimensioanl feature vector of GIST for each image.

LearnedSemantic-levelFeatures:For the purpose of semantic representation of the images, we extracted three object-based representation of the images: Classeme [29], Picodes [8], and CNN-based features [16]. In all these three features, each bit (element) of the feature vector represents the confidence of the presence of an object-category in the image, therefore they provide a semantic encoding of the images. The list of object categories is user-specified and not covering all object categories in the real world. Despite the limited number of categories in this type of modeling, these semantic encoding of images have shown remarkable results for the task of image search.

However, for learning these features, the object-categories are generic, mostly used for realistic images and are not specifically designed for the purpose of art. First two features are designed to capture the presence of a set of basic-level object categories as following: a list of entry-level categories (e.g. horse and cross) is used for downloading a large collection of images from the web. For each image a comprehensive set of low-level visual features are extracted and one classifier is learned for each category. For a given test image, these classifiers are applied on the image and the responses (confidences) make the final feature vector. We followed the implementation of [7] and for each image extracted a 2659 dimensional real-valued Classeme feature vector and a 2048 dimensional binary-value Picodes feature.

Convolutional Neural Networks (CNN) [17] showed a remarkable performance for the task of large-scale image categorization [16]. CNNs have four convolutional layers followed by three fully connected layers. Bar et al [4] showed that a combination of the output of these fully connected layers achieve a superior performance for the task of style classification of paintings. Following this observation, we used the last layer of a pre-trained CNN [16] (1000 dimensional real-valued vectors) as another feature vector.

3.4 Metric Learning

These aforementioned extracted visual features are meant to be used for real images, therefore, we should tune these features to perform reasonable on paintings as well. We consider using these features for the task of classifications in fine-art paintings, which is equivalent to put similar paintings close to each other. For the purpose of similarity measurement, we apply a list of metric learning approaches to find a reasonable approach. Metric learning is an active research area in the field of machine learning and we encourage interested readers to check surveys on this topic. In a formal notion, metric learning is defined as finding a real-valued mathematical function that assigns a score to each pair of its input. This score shows how similar are these items, where smaller number shows less difference and higher similarity. For this paper, we consider the following metric learning approaches:

Neighborhood Component Analysis (NCA): This approaches focuses on analyzing the nearest neighbors. This analysis is mainly based on putting neighbors of the same class (e.g. painting style in our study) close to each other.

Large Margin Nearest Neighbors (LMNN): LMNN [32] is an approach for learning a Mahalanobis distance, which is widely used because of its global optimum solution and superior performance in practice. The learning of this metric involves a set of constrains, all of which are defined locally. This means that LMNN enforces the k nearest neighbor of any training instance belonging to the same class (these instances are called "target neighbors"). This should be done while all the instances of other classes, referred as "impostors", should be far from this point. For finding the target neighbors, Euclidean distance has been applied to each pair of samples.

This metric learning approach is related to Support Vector Machines (SVM) in principle, which theoretically engages its usage along with SVM for the task of classification. Due to the popularity of LMNN, different variations of it have been introduced, including a non-linear version called gb-LMNN [32] which we used in our experiments as well. However, its performance for classification tasks was worse that linear LMNN. We assume this poor performance is rooted in the nature of visual features that we extract for paintings.

Boost Metric [27]***:*** The idea behind this approach follows this intuition: instead of learning a universal metric that works the best on all data, it might be better to learn and combine a set of weaker metrics that are not universal (giving the best performance across all data), but have a reasonable performance on a subset of the data. Shen et al [27] use this fact and instead of learning a metric directly, finds a set of metrics that can be combined and give the final metric. They treat each of these matrices as a Weak Learner, which is used in the literature of Boosting methods. The resulting algorithm applies the idea of AdaBoost to Mahalanobis distance, which has been shown to be quiet efficient in practice.

This method is particularly of our interest, because we can learn an individual metric for each style of paintings and finally merge these metrics to get a unique final metric. Theoretically the final metric can perform well to find similarities inside each style/genre of paintings as well.

Information Theory Metric Learning (ITML) [11]***:*** This metric learning algorithm is based on Information theory rather than numerical distances. In other words, the learning part of this metric is rooted in entropy measurement and probability models.

Metric Learning for Kernel Regression (MLKR): This approach performs similar to NCA, which minimizes the classification error. Weinberger and Tesauro [31] learn a metric by optimizing the leave-one-out error for the task of kernel regression. In kernel regression, there is an essential need for proper distances between points that will be used for weighting sample data. MLKR learn this distance by minimizing the leave-one-out error for regression on training data.

Although this metric learning method is designed for kernel regression, the resulted distance function can be used in variety of tasks.

4 Experiments

4.1 Experimental Setting

Visual Features: As we explained in section 3, we extract GIST features as low-level visual features and Classeme, Picodes and CNN-based features as the high-level semantic features. We followed the original implementation of Oliva and Torralba [23] to get a 512 dimensional feature vector. For Classeme and Picodes we used the implementation of Bergamo et al [29], resulting in 2659 dimensional Classeme features and 2048 dimensional Picodes features. We used the implementation of Vedaldi and Lenc [30] to extract 1000 dimensional feature vectors of the last layer of CNN. W

Object-based representations of the images produce feature vectors that are much higher in dimensionality than GIST descriptors. In the sake of a fair comparison of all types of features for the task of metric learning, we transformed all feature vectors to have the same size as GIST (512 dimensional). We did this by applying Principle Component Analysis (PCA) for each type and projecting the original features onto the first

512 eigenvectors (with biggest eigenvalues). In order to verify the quality of projection, we looked at the corresponding coefficients of eigenvalues for PCA projections. In-

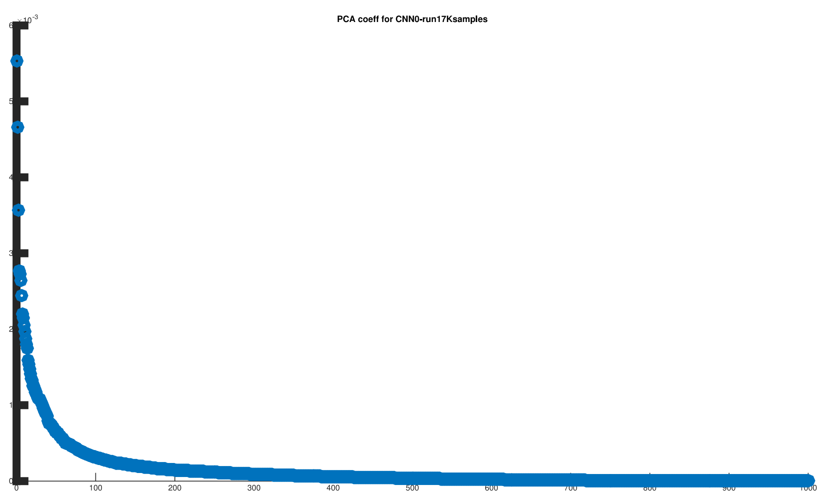

Figure 4: PCA coefficients for CNN features

dependent of feature type, the value of these coefficients drops significantly after the first 500 eigenvectors. For example, figure 4 plots these coefficients of PCA projection for CNN features. Summation of the first 500 coefficients is 95.88% of the total summation. This shows that our projections (with 512 eigenvectors) captures the true underlying space of the original features. Using these reduced features speeds up the metric learning process as well.

Metric Learning We used implementation of [32] to learn LMNN metric (both version of linear and non-linear) and MLKR. For the BoostMetric we slightly adjusted the implementation of [27]. For NCA we adopted its implementation by Fowlkes to work on large scale feature vectors smoothly. For the case of ITML metric learning, we followed the original implementation of authors with the default setting. For the rest of methods, parameters are chosen through a grid search that finds the minimum nearest neighbor classification. Regarding the training time, learning the ITML metric was the fastest and learning NCA and LMNN were the slowest ones. Due to computational constrains we set the parameters of LMNN metric to reduce the size of features to 100. NCA metric reduces the dimension of features to the number of categories for each tasks: 27 for style classification, 23 for artist classification and 10 for genre classification. We randomly picked 3000 samples, which we used for metric learning. These samples follow the same distribution as original data and are not used for classification experiments.

4.2 Classification Experiments

For the purpose of metric learning, we conducted experiments with labels for three different tasks of style, genre and artist prediction. In following sections, we investigate the performance of these metrics on different features for classification of aforementioned concepts.

We learned all the metrics in section 3 for all 27 styles of paintings in our dataset (e.g. Expressionism, Realism, etc.). However, we did not use all the genres for learning metrics. In fact, in our dataset we have 45 genres, some of which have less than 20 images. This makes the metric learning impractical and highly biased toward genres with larger number of paintings. Because of this issue, we focus on 10 genres with more than 1500 paintings. These genres are listed in table 1. In all experiments we conducted 3-fold cross validation and reported the average accuracy over all partitions. We found the best value for penalty term in SVM (which is equal to 10) by three-fold cross validation. In the next three sections, we explain settings and findings for each task independently.

Style Classification: Table 2 contains the result (accuracy percentage) of style classification (SVM) after applying different metrics on a set of features. Columns correspond to different features and rows are different metrics that are used for projecting features before learning style classifiers. In order to quantify the improvement by learning similarity metrics, we conducted a baseline experiment (first row in the table) as the following: For each type of features, we learn a set of one-vs-all classifiers on raw feature vectors. Generally, Boost metric learning and ITML approaches give the highest in accuracy for the task of style classification over different visual features. However, the greatest improvement over the baseline is gained by application of Boost metric on Classeme features. We visualized the confusion matrix for the task of style classification, when we learn Boost metric on Classeme features.

Figure 5 shows this matrix, where red represents higher values. Further analysis of some confusions that are captured in this matrix result in interesting findings. In the rest of this paragraph we explain some of these cases. First, we found that there is a big confusion between "Abstract expressionism" (first row) and "Action paintings" (second column). Art historians verify the fact that this confusion is meaningful and somehow expected. "Action painting" is a type or subgenre of "abstract expressionism" and are characterized by paintings created through a much more active process– drips, flung paint, stepping on the canvas.

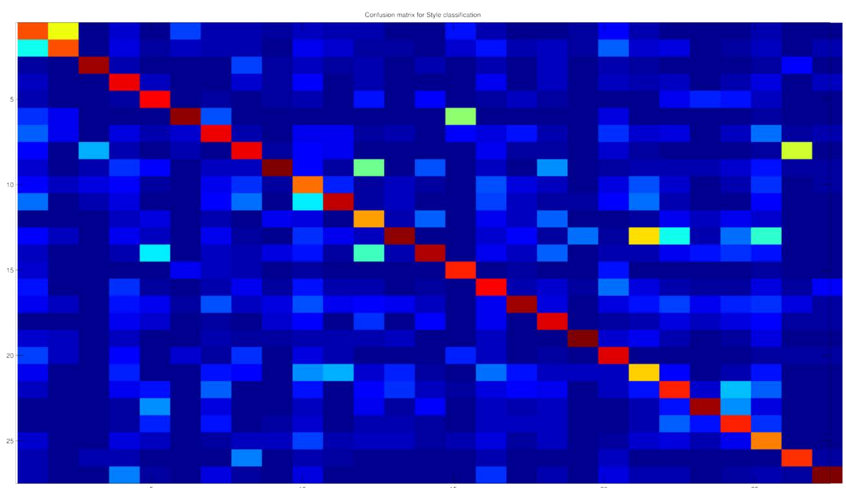

Figure 5: Confusion matrix for Style classification. Confusions are meaningful only when seen in color.

Another confusion happens between "Expressionism" (column 10) and "Fauvism" (row 11), which is actually expected based on art history literature. "Mannerism" (row 14) is a style of art during the (late) "Renaissance" (column 12), where they show unusual effect in scale and are less naturalistic than "Early Renaissance". This similarity between "Mannerism" (row 14) and "Renaissance" (column 12) is captured by our system as well where results in confusion during style classification. "Minimalism" (column 15) and "Color field paintings" (6th row) are mostly confused with each other. We can agree on this finding as we look at members of these styles and figure out the similarity in terms of simple form and distribution of colors. Lastly some of the confusions are completely acceptable based on the origins of these styles (art movements) that are noted in art history literature. For example, "Renaissance" (column 18) and "Early Renaissance" (row 9); "Post Impressionism" (column 21) and "Impressionism" (row 13); "Cubism" (8th row) and "Synthetic Cubism" (column 26). Synthetic cubism is the later act of cubism with more color continued usage of collage and pasted papers, but less linear perspective than cubism.

Genre Classification: We narrowed down the list of all genres in our dataset (45 in total) to get a reasonable number of samples for each genre (10 selected genres are listed

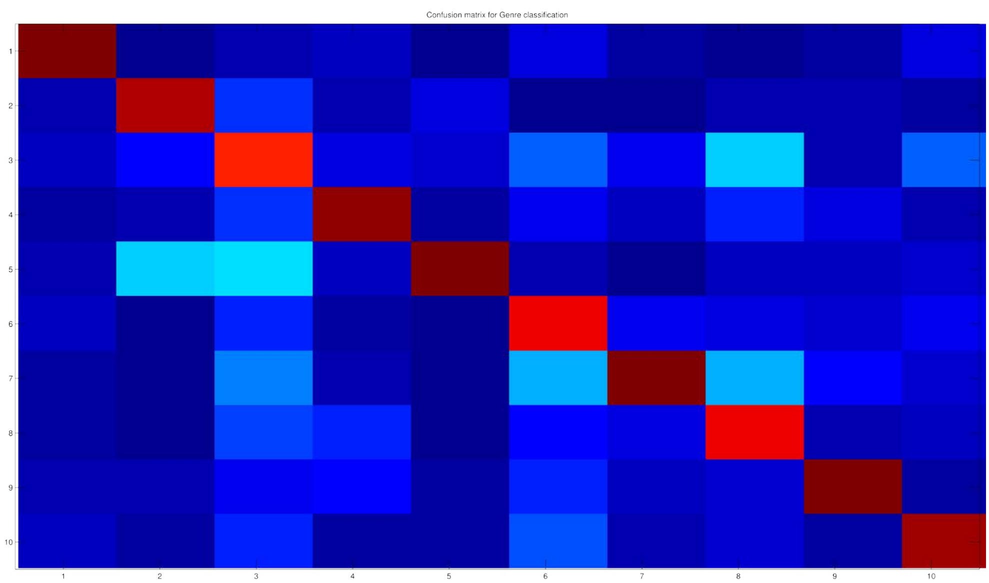

Figure 6: Confusion matrix for Genre classification. Confusions are meaningful only when seen in color.

in table 1). We trained ten one-vs-all SVM classifiers and compare their performance in Table 3. In this table columns represent different features and rows are different metric that we used to compute the distance. As table 3 shows we achieved the best performance for genre classification by learning Boost metric on top of Classeme features. Generally, the performance of these classifiers are better than classifiers that we trained for style classification. This is expected as the number of genres is less than the number of styles in our dataset.

Figure 6 shows the confusion matrix for classification of genre by learning Boost metric, when we used Classeme features. Investigating the confusions that we find in this matrix, reveals interesting results. For example, our system confuses "Landscape" (5th row) with "Cityspace" (2nd column) and "Genre paintings" (3rd column). However, this confusion is expected as art historians can find common elements in these genres. On one hand "Landscape" paintings usually show rivers, mountains and valleys and there is no significant figure in them; frequently very similar to "Genre paintings" as they capture daily life. The difference appears in the fact that despite the "Genre paintings", "Landscape" paintings are idealized. On the other hand, "Landscape" and "Cityspace" paintings are very similar as both have open space and use realistic color tonalities.

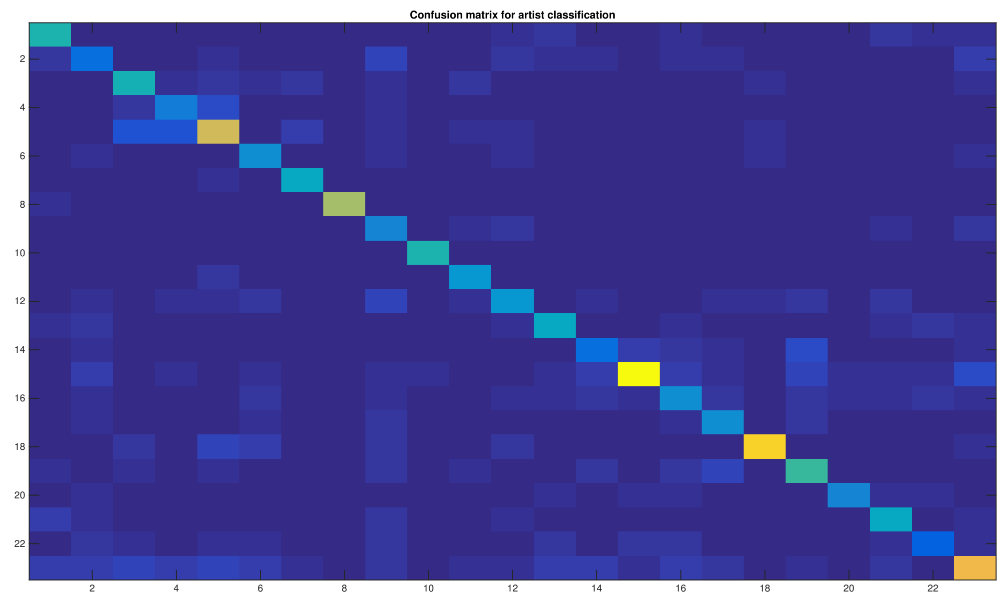

Figure 7: Confusion matrix for Artist classification. Confusions are meaningful only when seen in color.

Artist Classification: For the task of the artist classification, we trained one-vs-all SVM classifiers for each of 23 artists. For each test image, we determine its artist by finding the classifier that produces the maximum confidence. Table 4 shows the performance of different combinations of features and metrics for this task. In general learning Boost metric improves artist classification better than all other metrics, except the case of CNN features where learning ITML metric gained the best performance. We plotted the confusion matrix of this classification task in figure 7. In this plot, some confusions between artists are clearly reasonable. We investigated two cases:

First case, "Claude Monet" (5th row) and "Camille Pissaro"(3rd column). Both of these Impressionist artists who lived in the late nineteen and early twentieth centuries. Interestingly, based on art history literature Monet and Pissaro became friends when they both attended the" Acade´mie Suisse" in Paris. This friendship lasted for a long time and resulted in some noticeable interactions between them. Second case, paintings of "Childe Hassam" (4th row) are mostly confused with ones from "Monet" (5th column). This confusion is acceptable as Hassam is an American Impressionist, who declared himself as being influenced by French Impressionists. Hassam called himself an "Extreme Impressionist", who painted some flag-themed artworks similar to Monet.

By looking at reported performances in tables 2-4, we conclude that, all three classification tasks can benefit from learning the appropriate metric. This means that we can improve the accuracy of baseline classification by learning metrics independent of the type of visual feature or the concept that we are classifying painting based on. Experimental results show that, independent of the task, NCA and MLKR approaches are performing worse than other metrics. Additionally, Boost metric always gives the best or the second best results for all classification tasks.

Regarding analysis of importance of features, we can verify that Classeme and Picode features are better image representations for classification purposes. Based on these classification experiments, we claim that Classemes and Picodes features perform better than CNN features. This is rooted in the fact that amount of supervision for training Classeme and Picodes is more than CNN training. Also, unlike Classeme and Picodes, CNN feature is designed to categorize the object insides a given bounding box. However, in the case of paintings we cannot assume that all the bounding boxes around the objects are given.

Integration of Features and Metrics We investigated the performance of different metric learning approaches and visual features individually. In the next step, we find out the best performance for aforementioned classification tasks by combining different visual features. Toward this goal, we followed two strategies. First, for a given metric, we project visual features by applying the metric and concatenate these projected visual features together. Second, we fixed the type of visual feature that we use and project it with the application of different metrics and concatenate these projections all together. Having this larger feature vectors (either of two strategies), we train SVM classifiers for three tasks of Style, Genre and Artist classification. Table 6 shows the results of these experiments where we followed the earlier strategy and table 5 shows the results of the later case. In general we get better results by fixing the metric and concatenating the projected feature vectors (first strategy).

The work of Bar et al [4] is the most similar to ours and we compare our final results of these experiments with their reported performance. [4] only performed the task of style classification on half of the images in our dataset and achieved the accuracy of 43% by using two variations of PiCoDes features and two layers of CNN. However, we outperform their approach by achieving 45.97 % accuracy for the task of style classification when we used LMNN metric to project GIST, Classeme, PiCoDes and CNN features and concatenate them all together as it is reported in the third column of table 6.

Our contribution goes beyond outperforming state-of-the-art by learning a more compact feature representation. In this work, our best performance for style classification happens when we concatenate four 100-dimensional feature vectors. This results in a 400 dimensional feature vectors that we train SVM classifiers on top of them. However [4] extract a 3882 dimensional feature vector to their best reported performance. As a result we not only outperform the state-of-the-art, but presented a better image representation that reduces the amount of space by 90%. Our efficient feature vector is an extremely useful image representation that gains the best classification accuracy and we consider its application for the task of image retrieval as future work.

To qualitatively evaluate extracted visual features and learned metrics, we did a prototype image search task. As the feature fusion with application of LMNN metric gives the best performance for style classification, we used this setting as our similarity measurement model. Figure 8 shows some sample output of this image search task. For each pair, the image on the left is the query image, which we find the closest match (image on the right) to it based on LMNN and feature fusion. However, we force the system to pick the closest match that does not belong to the same style as the query image. This verifies that although we learn the metric based on style labels, the learned projection can find similarity across styles.

5 Conclusion and Future Works

In this paper we investigated the applicability of metric learning approaches and performance of different visual features for learning similarity in a collection of fine-art paintings. We implemented meaningful metrics for measuring similarity between paintings. These metrics are learned in a supervised manner to put paintings from one concept close to each other and far from others. In this work we used three concepts: Style, Genre and Artist. We used these learned metrics to transform raw visual features into another space that we can significantly improve the performance for three important tasks of Style, Genre and Artist classification. We conducted our comparative experiments on the largest publicly available dataset of fine-art paintings to evaluate the performance for the aforementioned tasks.

We conclude that:

1. Classeme features show the superior performance for all three tasks of Style, Genre or Artist classification. This superior performance is independent of the type of metric that has been learned.

2. In the case of working on individual type of visual features, Boost metric and Information Theoretic Metric Learning (ITML) approaches improve the accuracy of classification tasks across all features.

3. For the case of using different types of features all together (feature fusion), Large-Margin Nearest-Neighbor (LMNN) metric learning achieves the best performance for all classification experiments.

4. By learning LMNN metric on Classeme features, we find an optimized representation that not only outperforms state-of-the art for the task of style classification, but reduce the size of feature vector by 90%. We consider verification of applicability of this representation for the task of image retrieval and recommendation systems as future work. As other future works we would like to learn metrics based on other annotation (e.g. time period).

Bibliography

[1] A. E. Abdel-Hakim and A. A. Farag. C-sift: A sift descriptor with color invariant characteristics. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2006.

[2] R. Arnheim. Visual thinking. University of California Press, 1969.

[3] R. S. Arora and A. M. Elgammal. Towards automated classification of fine-art painting style: A comparative study. In ICPR, 2012.

[4] Y. Bar, N. Levy, and L. Wolf. Classification of artistic styles using binarized features derived from a deep neural network. 2014.

[5] A. Bentkowska-Kafel and J. Coddington. Computer Vision and Image Analysis of Art: Proceedings of the SPIE Electronic Imaging Symposium, San Jose Convention Center, 18-22 January 2010. PROCEEDINGS OF SPIE. 2010.
[6] I. E. Berezhnoy, E. O. Postma, and H. J. van den Herik. Automatic extraction of brushstroke orientation from paintings. Machine Vision and Applications, 20(1):1–9, 2009.
[7] A.Bergamo andL.Torresani. Classemes and other classifier-based features for efficient object categorization. IEEE Transactions on Pattern Analysis and Machine Intelligence, page 1, 2014.
[8] A. Bergamo, L. Torresani, and A. W. Fitzgibbon. Picodes: Learning a compact code for novel-category recognition. In Advances in Neural Information Processing Systems, pages 2088–2096, 2011.
[9] G. Carneiro, N. P. da Silva, A. D. Bue, and J. P. Costeira. Artistic image classification: An analysis on the printart database. In ECCV, 2012.
[10] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In International Conference on Computer Vision & Pattern Recognition, volume 2, pages 886–893, June 2005.
[11] J.V.Davis, B.Kulis, P.Jain, S.Sra, and I.S.Dhillon. Information-theoretic metric learning. In ICML, 2007.
[12] M. V. FahadShahbazKhan, JoostvandeWeijer. Whopaintedthis painting?, 2010.
[13] L. Fichner-Rathus. Foundations of Art and Design. Clark Baxter, 2008.
[14] J. Goldberger, S. Roweis, G. Hinton, and R. Salakhutdinov. Neighbourhood components analysis. In NIPS, 2004.
[15] C. R. Johnson, E. Hendriks, I. J. Berezhnoy, E. Brevdo, S. M. Hughes, I. Daubechies, J. Li, E. Postma, and J. Z. Wang. Image processing for artist identification. Signal Processing Magazine, IEEE, 25(4):37–48, 2008.
[16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
[17] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.
[18] J. Li and J. Z. Wang. Studying digital imagery of ancient paintings by mixtures of stochastic models. Image Processing, IEEE Transactions on, 13(3): 340–353, 2004.
[19] J. Li, L. Yao, E. Hendriks, and J. Z. Wang. Rhythmic brushstrokes distinguish van gogh from his contemporaries: Findings via automated brushstroke extraction. IEEE Trans. Pattern Anal. Mach. Intell., 2012.
[20] T. E. Lombardi. The classification of style in fine-art painting. ETD Collection for Pace University. Paper AAI3189084., 2005.
[21] D. G. Lowe. Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision, 2004.
[22] S. Lyu, D. Rockmore, and H. Farid. A digital technique for art authentication. Proceedings of the National Academy of Sciences of the United States of America, 101(49):17006–17010, 2004.

[23] A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. IJCV, 2001.
[24] G. Polatkan and S. Jafarpour and A. Brasoveanu and S. Hughes and L. Daubechies .Detection of forgery in paintings using supervised learning. In 16th IEEE International Conference on Image Processing (ICIP), 2009.
[25] R. Sablatnig, P. Kammerer, and E. Zolda. Hierarchical classification of paintings using face- and brush stroke models. 1998.
[26] B., K. Abe, and A. Elgammal. Knowledge discovery of artistic influences: A metric learning approach. In ICCC, 2014.
[27] C. Shen, J. Kim, L. Wang, and A. van den Hengel. Positive semi-definite metric learning using boosting-like algorithms. Journal of Machine Learning Research, 13:1007–1036, 2012.
[28] D. G. Stork. Computer vision and computer graphics analysis of paintings and drawings: An introduction to the literature. In Computer Analysis of Images and Patterns, pages 9–24. Springer, 2009.
[29] L. Torresani, M. Szummer, and A. Fitzgibbon. Efficient object category recognition using classemes. In ECCV, 2010.
[30] A. Vedaldi and K. Lenc. MatConvNet : convolutional neural networks for MATLAB. CoRR, abs/1412.4564, 2014.
[31] K. Weinberger and G. Tesauro. Metric learning for kernel regression. In Eleventh international conference on artificial intelligence and statistics, pages 608–615, 2007.
[32] K. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. JMLR, 2009.

Tables

Table 1: List of Styles, Genres and Artists in our collection of fine-art paintings. Numbers in the parenthesis are index of the row/column in confusion matrices 5, 6& 7 accordingly.

Task Name	List of Members
Style	Abstract Expressionism(1); Action Painting(2); Analytical Cubism(3); Art Nouveau-Modern Art(4); Baroque(5); Color Field Painting(6); Contemporary Realism(7); Cubism(8); Early Renaissance(9); Expressionism(10); Fauvism(11); High Renaissance(12); Impressionism(13); Mannerism-Late-Renaissance(14); Minimalism(15); Primitivism-Naive Art(16); New Realism(17); Northern Renaissance(18); Pointillism(19); Pop Art(20); Post Impressionism(21); Realism(22); Rococo(23); Romanticism(24); Symbolism(25); Synthetic Cubism(26); Ukiyo-e(27)
Genre	Abstract painting(1); Cityscape(2); Genre painting(3); Illustration(4); Landscape(5); Nude painting(6); Portrait(7); Religious painting(8); Sketch and Study(9); Still Life(10)
Artist	Albrecht Durer(1); Boris Kustodiev(2); Camille Pissarro(3); Childe Hassam(4); Claude Monet(5); Edgar Degas(6); Eugene Boudin(7); Gustave Dore(8); Ilya Repin(9); Ivan Aivazovsky(10); Ivan Shishkin(11); John Singer Sargent(12); Marc Chagall(13); Martiros Saryan(14); Nicholas Roerich(15); Pablo Picasso(16); Paul Cezanne(17); Pierre-Auguste Renoir(18); Pyotr Konchalovsky(19); Raphael Kirchner(20); Rembrandt(21); Salvador Dali(22); Vincent van Gogh(23)

Table 2: Accuracy for the task of style classification.

Metric / Feature	GIST	Classemes	Picodes	CNN	Dimension
Baseline	10.83	22.62	20.76	12.32	512
Boost	16.07	31.77	28.58	15.18	512
ITML	13.02	30.67	28.42	15.34	512
LMNN	12.54	27	24.14	16.83	100
MLKR	12.65	24.12	14.86	12.63	512
NCA	13.29	28.19	24.84	16.37	27

Table 3: Accuracy for the task of genre classification.

Metric / Feature	GIST	Classemes	Picodes	CNN	Dimension
Baseline	28.10	49.98	49.63	35.14	512
Boost	31.01	57.87	57.35	46.14	512
ITML	33.10	57.86	57.28	46.80	512
LMNN	39.06	54.96	54.42	49.98	100
MLKR	32.81	54.29	42.79	45.02	512
NCA	30.39	51.38	52.74	49.26	27

Table 4: Accuracy for the task of artist classification.

Metric / Feature	GIST	Classemes	Picodes	CNN	Dimension
Baseline	17.58	45.29	45.82	20.38	512
Boost	25.65	57.76	55.50	29.65	512
ITML	19.95	51.79	53.93	31.04	512
LMNN	20.41	53.99	53.92	30.92	100
MLKR	21.22	49.61	19.54	21.77	512
NCA	18.80	53.70	53.81	22.26	27

Table 5: Classification performance for metric fusion methodology.

Task / Feature	GIST	Classeme	Picodes	CNN
Style	20.21	37.33	33.27	21.99
Genre	35.94	58.29	56.09	47.05
Artist	30.37	59.37	55.65	33.62

Table 6: Classification performance for feature fusion methodology.

Task / Metric	Boost	ITML	LMNN	MKLR	NCA
Style	41.74	45.05	45.97	38.91	40.61
Genre	58.51	60.28	58.48	55.79	54.82
Artist	61.24	60.46	63.06	53.19	55.83

Table 7: Annotation of paintings in Figure 8. Each row corresponds to one pair of images, labeled with the name of painting, its style and its artist. First six rows correspond to the six pairs on the left in Figure 8 and next six rows correspond to the pairs

Art name	Artist	Style	Art name	Artist	Style
The marble staircase which leads up to S. Maria in Aracoeli in Rome	Christoffer Wilhelm Eckersberg	Neoclassicism	Corner Paleissingel Straat In Amsterdam	Cornelis Vreedenburgh	Impressionism
Brickworks at Eragny	Camille Pissarro	Pointillism	Countryside and Eragny Church and Farm	Camille Pissarro	Impressionism
At the races	Edgar Degas	Impressionism	Bayan	Viktor Vasnetsov	Romanticism
View towards the port of Hammamet	Paul Klee	Cubism	Sacrificial stone in Baalbek	Tivadar Kosztka Csontvary	Post-Impressionism
Ladies in a row	Walasse Ting	Pop Art	The four Apostles	Albrecht Durer	Northern Renaissance
Communion of dying	Alexey Venetsianov	Realism	Madonna Enthroned and ten saints	Rosso Fiorentino	Mannerism (Late Renaissance)

The lamentation	Alexey Venetsianov	Realism	Adoration of the shepherds	Bartolome Esteban Murillo	Baroque
A chestnut wood among the rocks	Camille Corot	Realism	Artist's children in garden	Max Slevogt	Impressionism
A road in the countryside, near lake leman	Camille Corot	Realism	Little Russian ox cart in winter	Ivan Aivazovsky	Romanticism
Hungarian gipsies	Endre Bartos	Expressionism	Composition with Romanian motifs	Corneliu Michailescu	Cubism
Lake Geneva from Chexbres	Ferdinand Hodler	Post-Impressionism	Sunset in the winter. A coast of the sea	Arkhip Kuindzhi	Impressionism
St. Jacques leads to martyrdom	Andrea Mantegna	High Renaissance	St. James the Great on his way to execution (painted on : 1448)	Andrea Mantegna	Early Renaissance

Babak Saleh is a PhD candidate in the department of computer science at Rutgers University, where he conducts research in the intersection of computer vision, machine learning, and human perception. Inspired by human visual perception, he has developed computational models for measuring typicality of an image and its application in learning more robust visual classifiers. He holds a MS in Computer Science and a second MS in Statistics from Rutgers University. He completed his undergraduate studies in Computer Science and Mathematics at Sharif University of Technology in Tehran, Iran. He is the recipient of outstanding student paper award from AAAI 2016, and NSF I-Corps award. His research has been recognized by major media and press outlets, including NBC News, PBS, New York Times, Washington Post, WIRED, Fast Company and IEEE MultiMedia.

Correspondence e-mail: babaks@cs.rutgers.edu

Dr. Ahmed Elgammal is an associate professor at the Department of Computer Science, Rutgers, the State University of New Jersey. He is a member of the Center for Computational Biomedicine Imaging and Modeling (CBIM) at Rutgers and affiliate member in Rutgers University Center for Cognitive Science (RUCCS.) and the director of the Art and Artificial Intelligence at Rutgers and the Human Motion Analysis Lab (HuMAn Lab.)

Correspondence e-mail: elgammal@cs.rutgers.edu

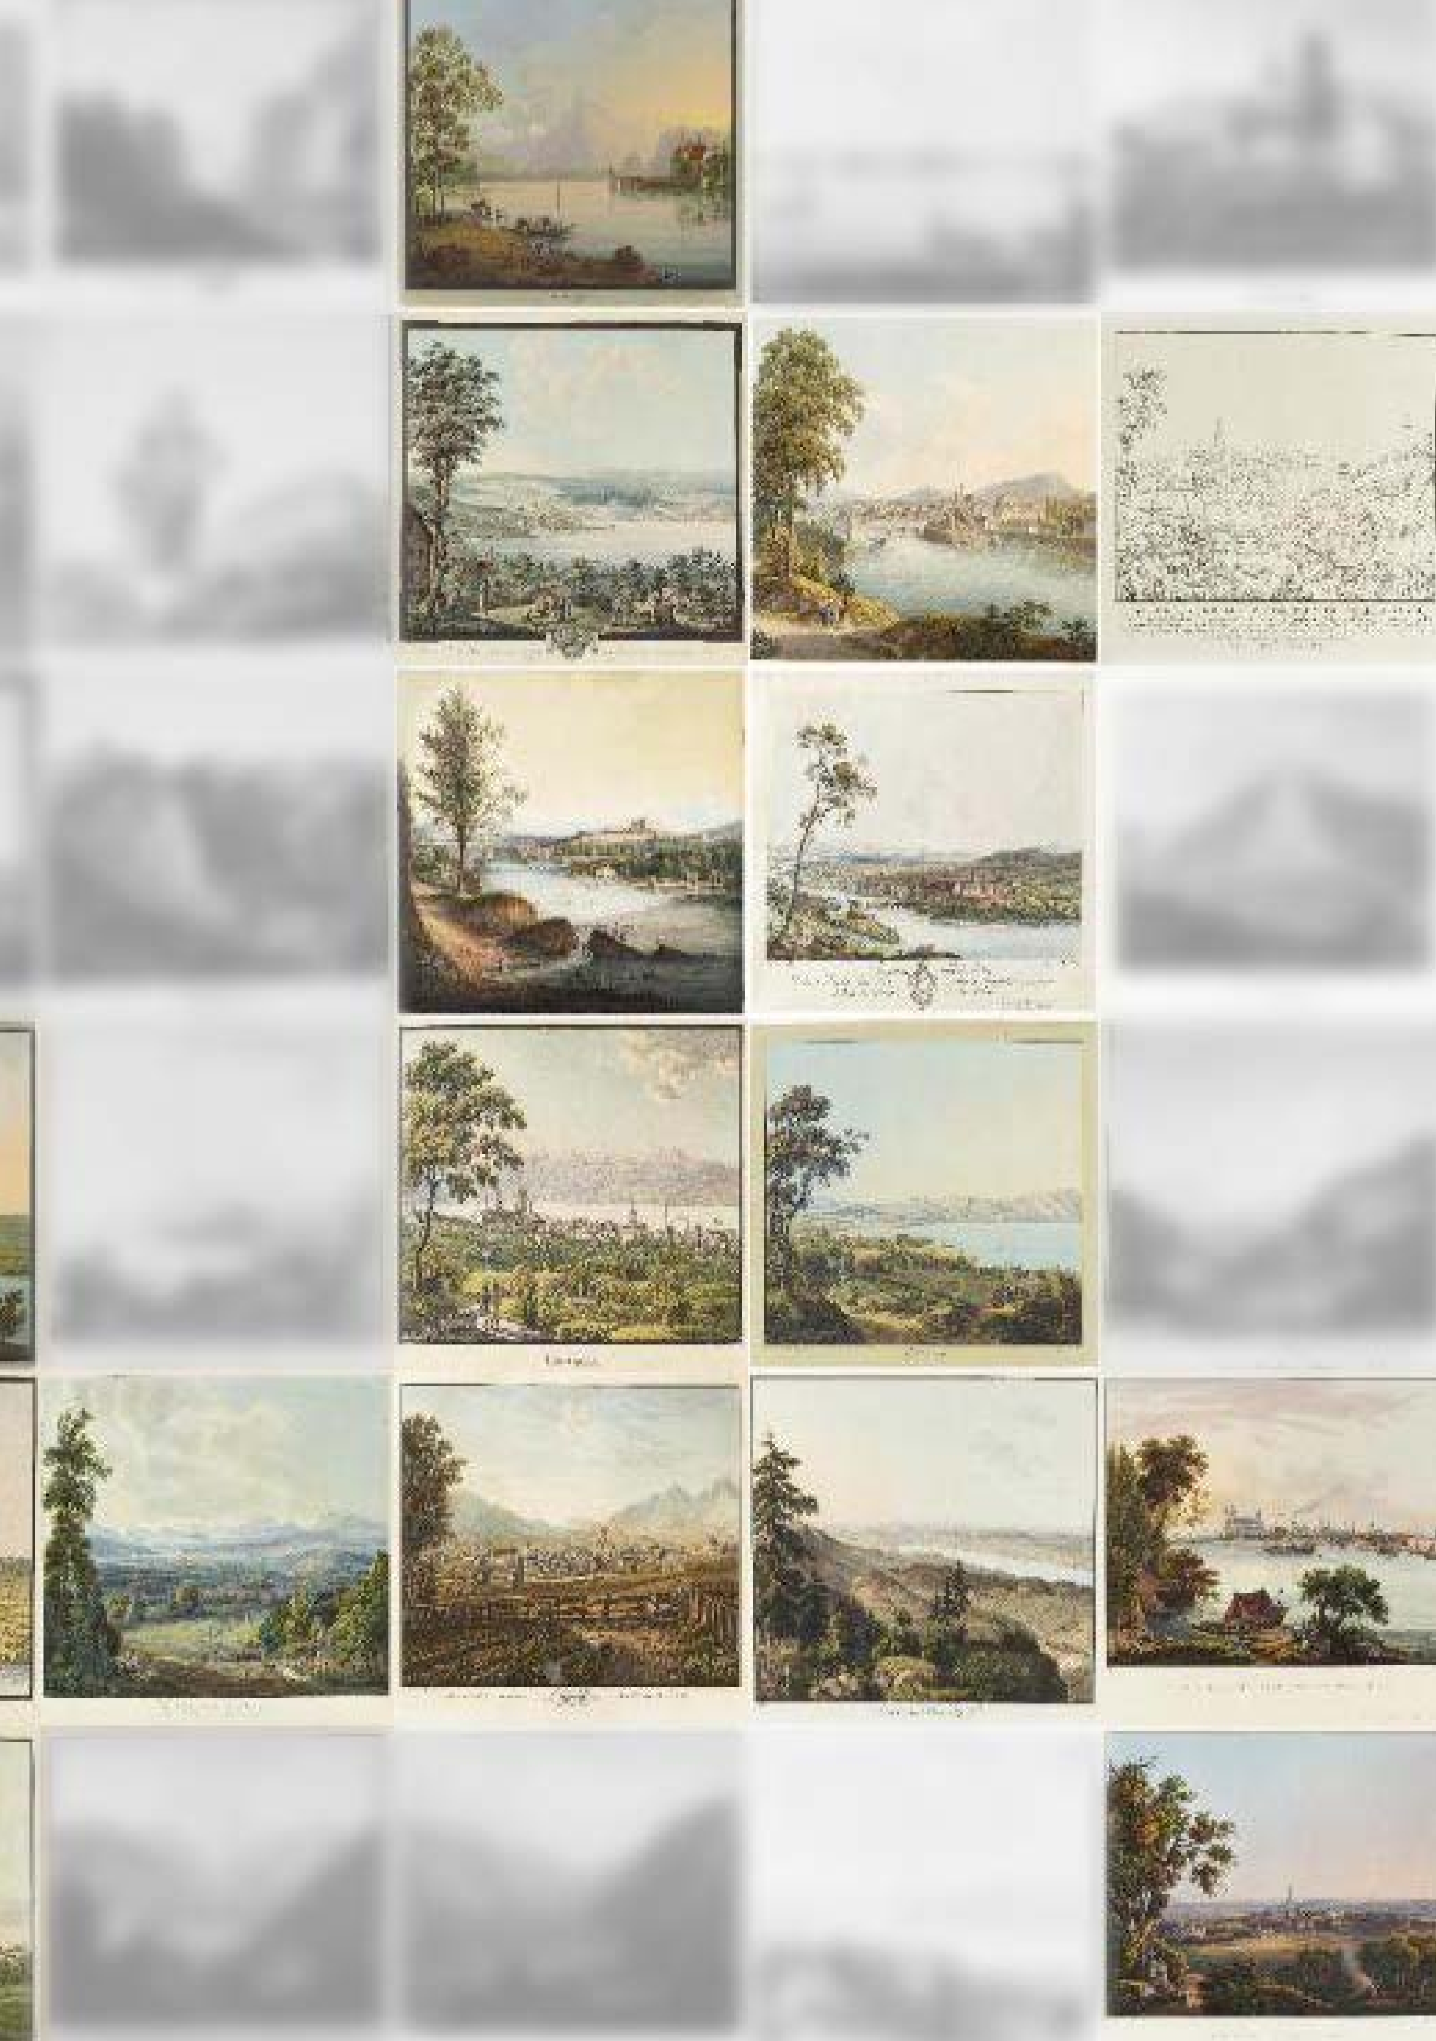

Peer-Reviewed

Gugelmann Galaxy. An Unexpected Journey through a collection of Schweizer Kleinmeister

Mathias Bernhard

Abstract: GLAM[1] institutions all over the world are digitizing their collections. As the number of items in such a collection amounts to tens or even hundreds of thousands, providing comprehensible access and presentation becomes increasingly difficult. At the same time, a steadily growing amount of this data is openly available. This gives rise to various projects approaching the hidden treasures in these collections with computational tools. The project presented here, Gugelmann Galaxy, lets the user explore an entire collection of digitized images and their textual metadata in an immersive three-dimensional cloud, whose configuration can be rearranged according to different criteria. The project questions traditional models of categorization and curating and implements alternative approaches prototypically.

Keywords: art collection, computer vision, machine learning, 3D visualization, curating

Introduction

The Gugelmann Galaxy project presented in this article touches two main issues of working with collections of art works. The first one is presentation and curation of the collection and the second is the understanding of the content of single items of the collection. Both these tasks are about organization, about the installation of rule sets and metrics to tell items apart and make subsets of items from the entire collection. In both areas exist many very interesting approaches. I would like to highlight a few of them here in order to situate the presented work.

As examples for the presentation of art collections serve the projects WikiArt (WikiArt 2010) and ArtStack (Konvitz 2012). These portals behave like static archives, showing the result of standard database queries as lists of thumbnails. Possible choices for a sorting on WikiArt are the hard coded categories style, genre and technique. Proposed similar items only include works by the same author or the same contributor. The platform artsy.net (Cleveland and Cwilich 2010) tries to go one step further with their Art Genome Project. A team of art historians and experts assigns tags to works out of a catalogue containing more than 1'000 different characteristics. They refer to these characteristics as genes.

Figure 1: Detail of Fig. 2, showing images with a tree in the left foreground.

The project Curiator (Erbuer and Boonstoppel 2012) joined forces with Curalytics, whose Steffon Davis defines curation on his blog (Davis 2013) as "Curation: the subjective selection, categorization, and arrangement of content". Curiator allows to create what they call rooms, individual collections curated by its users. The partner Curalytics applies a machine-learning algorithm called collaborative filtering, used by recommender systems known from media providers like Amazon or Netflix. It assumes that if two items are put in a room together by at least two different users, they have something in common. Analyzing all the rooms of all the users lets them on the one hand identify taste leaders and on the other hand control what shows up in the you-might-also-like-section.

Computer vision departments all over the world where approached in recent years with the question of image understanding in art works. The Visual Geometry Group at Oxford University dedicates itself to object detection in fine art paintings (Crowley and Zissermann 2014b). The results in terms of precision and recall are very promising. The downside is that only objects occurring very frequently in a lot of paintings, like trains, sheep, dogs, horses or chairs, offer a sufficient amount of samples for training a model. The range of detectable objects broadens when a combination of artworks and images from the Internet are used to train the model (Crowley and Zissermann 2014a). Researchers at the University of Heidelberg work on tasks such as object detection (Yarlagadda et al. 2012) and hand gesture estimation in medieval paintings (Schlecht, Carqué, and Ommer 2011). These tasks are highly complex, because the objects to be received (e.g. crowns) vary drastically in shape and color over the entire collection of images available. In addition, to be able to train a reliable classifier, a large collection of already labeled training samples must be available.

And finally somewhere in between the two – because they work both on features extracted from single items and the relationship network among all the items in the collection – are situated the researchers from Rutgers University. They made interesting studies on genre classification (Saleh and Elgammal 2015), creativity networks (Elgammal and Saleh 2015) or artistic influence (Saleh et al. 2014) in the WikiArt dataset. A combination of features and labels contained in available metadata and others extracted from the images by means of computer vision algorithms allows for training successful classification models.

Project

The Gugelmann Galaxy project grew out of a two days hackathon, a marathon for hackers. The First Swiss Open Cultural Data Hackathon was organized by opendata.ch, the Swiss chapter of the Open Knowledge Foundation. It took place on February

27th and 28th 2015 at the Swiss National Library (SNL) in Berne. Some 100 computer scientists, artists, researchers, and members of the heritage sector gathered to explore more than 30 open data sets, provided by over 20 different GLAM institutions with the aim to put this cultural data to wider use. One of the datasets provided by the SNL itself is the Gugelmann Collection.[2] The dataset was brought to the hackathon with the goal of raising its public awareness.

Our group, consisting of Nikola Marinčić, Jorge Orozco and myself[3] was joined by art historian Sonja Gasser to form an interdisciplinary team covering a broad field of interests and skills. The foundation of the project was developed during these two days by the team. Successive developments like the adoption to various platforms or the addition of alternative options are the result of my own research.

Problem

The collection was not easily accessible to the general public before. Even though the pictures in high resolution were available on Wikimedia Commons, one can only scroll through seemingly endless tables of thumbnails, ordered alphabetically by the authors' last names. In addition to the digital images, the collection provides an XML-file containing some metadata. The record for each item contains the records signature (corresponding to the filename and something like the unique identifier), title, author(s) and – if available – date, technique and a short description. Further fields hold information about the collector, the source and the legal permission but they are identical for all the items, as they all belong to the same collection. From the pure and unprocessed metadata alone, the items could linearly be sorted by date with the goal to unveil developments over time, similar to what Florian Kräutli (Kräutli 2015) showed with the MoMA collection (MoMA 2015). But the majority of the items in the Gugelmann Collection are from a relatively narrow range of time between 1750 and 1850 and so to investigate a development over time is difficult. Alternatively, one could try to group the items by technique. Because there did not seem to be a uniform naming convention, this cannot be done directly but needs some further processing of the raw data as we will see later.

Question

How can a user get an overview of the content of this collection, when one has to flick through these tables clicking "next" twelve times? Our goal is to provide a different form of access to this body of work by rethinking conventional ways of sorting and questioning concepts of curating. Items need no longer be stringed together linearly along a time or alphabetical axis. There is no necessity for items to be classified in pre-labeled drawers, obeying fixed, hierarchical ontologies. Instead of hard-edged categories, configurations can arrange freely, densifying to clusters at some points while spanning open voids at others. Binary decisions whether or not an item is for example colorful, are replaced by assigning to it a certain degree of colorfulness.

Method

The project Gugelmann Galaxy provides four different criteria for arrangement, based on four different similarity measures between the items. Two are based on textual data contained in the metadata (TECHNIQUE and DESCRIPTION) and two are based on attributes[4] extracted from the digital images using computer vision algorithms (COLOR and COMPOSITION). Each of these four criteria, its attempt and method of feature engineering shall now be described.

TECHNIQUE attempts a topological grouping of the information contained in the "Technik" field. Some annotations are very consistent, as there is 605 times the entry "Aquarell", 478 times "Umrissradierung, koloriert" or 456 times "Aquatinta, koloriert". But there is a long tail of a total of 82 different descriptions that are not supposed to make up a category on their own. For example one item's technique is described with "schwarze und weisse Kreide auf blauem Papier" and another one's "Kohle und weisse Kreide auf grau eingefäbtem Papier". A human reader can easily concede a high degree of similarity between the two, since they both contain "weiss" (white), "Kreide" (chalk/crayon) and are drawn on colored paper. For a computer, this is less obvious. Therefore, we compiled a list of 22 words[5] and checked, whether

Figure 2: Clustering of images by compositional concepts, 2d grid plot.

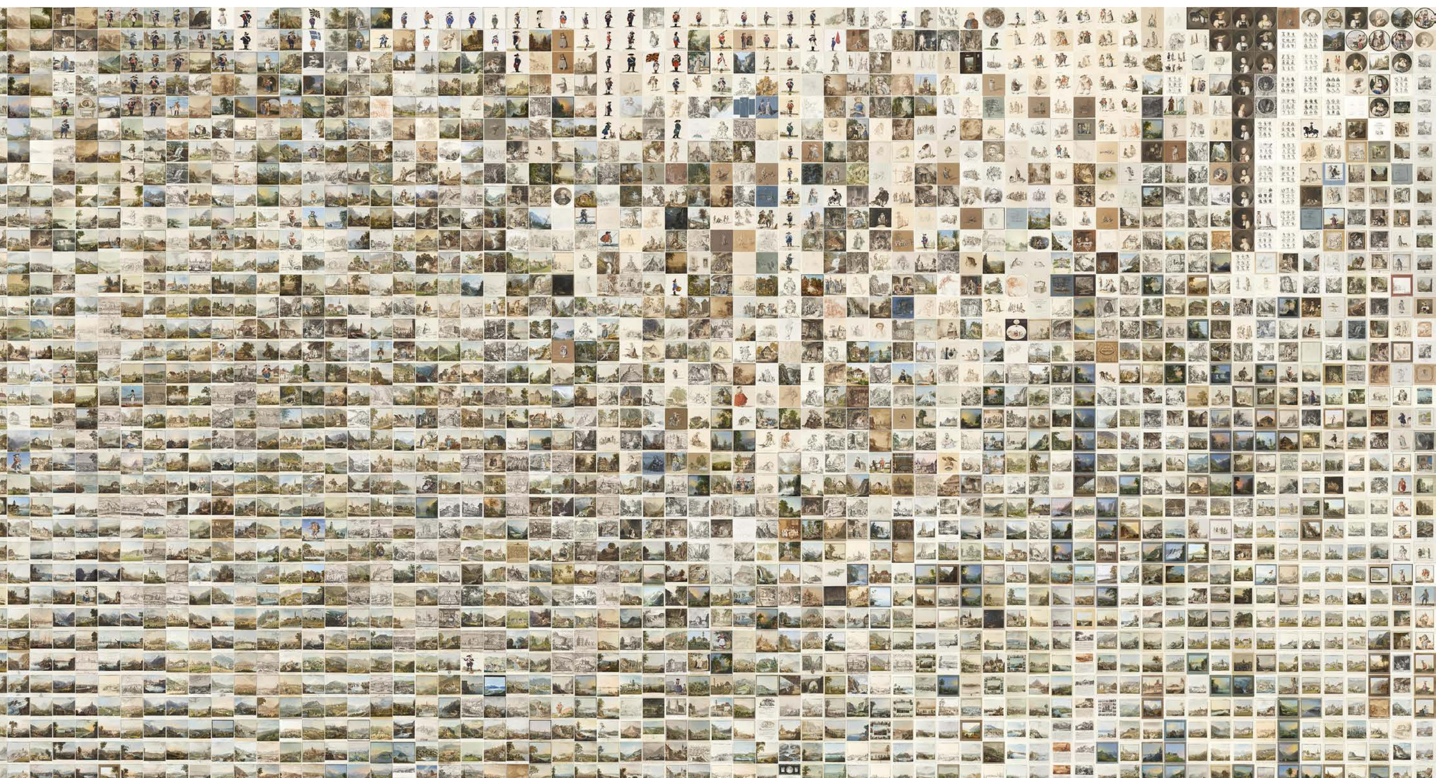

or not that word was contained in the description. This resulted in a 22 dimensional binary vector for each item, having a 1 for the presence and 0 for the absence of every word in the list. A 22-dimensional vector, each item a point in 22-dimensional space, thus represents the items. The selection of words is a first guess. One possible improvement is to work with word parts instead of full words only. This would allow an item described as “Kreidelithographie” to get a check for both “Kreide” and “Lithographie”. Another improvement could be to automatically assign more global tags like drawing to items containing “Tusche”, “Rötel”, “Kreide” or “Bleistift”, even if the word “Zeichnung” is not contained in the text. The elaboration of this list requires a lot of expert knowledge from art historians and can hardly be automated. A well composed set of words leads to a more fine grained distinction between items and thereby sheds some lights on what techniques were often combined together, instead of assigning two separate categories. Works produced using “Aquatinta, unkoloriert” are in proximity of the ones using “Aquatinta, koloriert” who again are close to the ones using “Umrissradierung, koloriert”.

The DESCRIPTION sorting aims to extract concepts from the three fields title, description and place. For this purpose, the text of all the three fields is concatenated, punctuation and numbers removed and the individual words extracted[6]. After removal of stop words (von, und, mit, der, die, das etc.), the list of the most frequent words is Bern (373), Kirche (323), Pfarrhaus (307), Blick (151), links / rechts (both 133), Schloss (128) and Ansicht (127). Words that occur in almost every description are not good indicators to measure the difference between two items. A method frequently used in natural language processing (NLP) called TF-IDF[7] is applied to account for that. Items are described as roughly 2300-dimensional vectors of reciprocal distance measures. The sorting resulting from this similarity measure reveals higher level concepts represented in the images, independent of author, date or technique.

The COLOR sorting is an attempt to compare images not only by their average color (one single value for hue or saturation) but rather the distribution of different colors. The feature vector extracted for that purpose is 300-dimensional, namely the red, green and blue channel of a thumbnail image down-sampled to 10 by 10 pixels. This sorting can distinguish coarse categories like landscapes (bluish in the upper, greenish in the lower part) or traditional dresses studies (column of colors in the center on bright background). It also unveils clusters of very specific palettes used by one painter to represent different landscapes.

And finally, COMPOSITION describes every image as a series of eight orientations with their relative intensity in 64 different regions (8x8 tiles) of the image, resulting in a

512-dimensional vector. This is called a HOG descriptor, for histogram of oriented gradients (Dalal and Triggs 2005). This feature vector describes an image by the directions dominant in different regions, independent from colors and allows therefore clustering the collection by composition of the images.

For every item, various sets of relevant features have been extracted that describe it as an n-dimensional vector. As was shown, anything can become a dimension, easily accessible numerical data like date or size, over low level features like color distribution or word counts up to high level features like edges or textures. Each item constitutes one data point in this n-dimensional space, with similar items being closer to each other. We are used to seeing two-dimensional scatterplots, where one feature (read: dimension) of

Figure 3: (l) Image with its eight closest neighbors in color distribution embedding: same palette, different places. (r) Same image with its eight closest neighbors in description embedding: Lucerne thematic cluster,

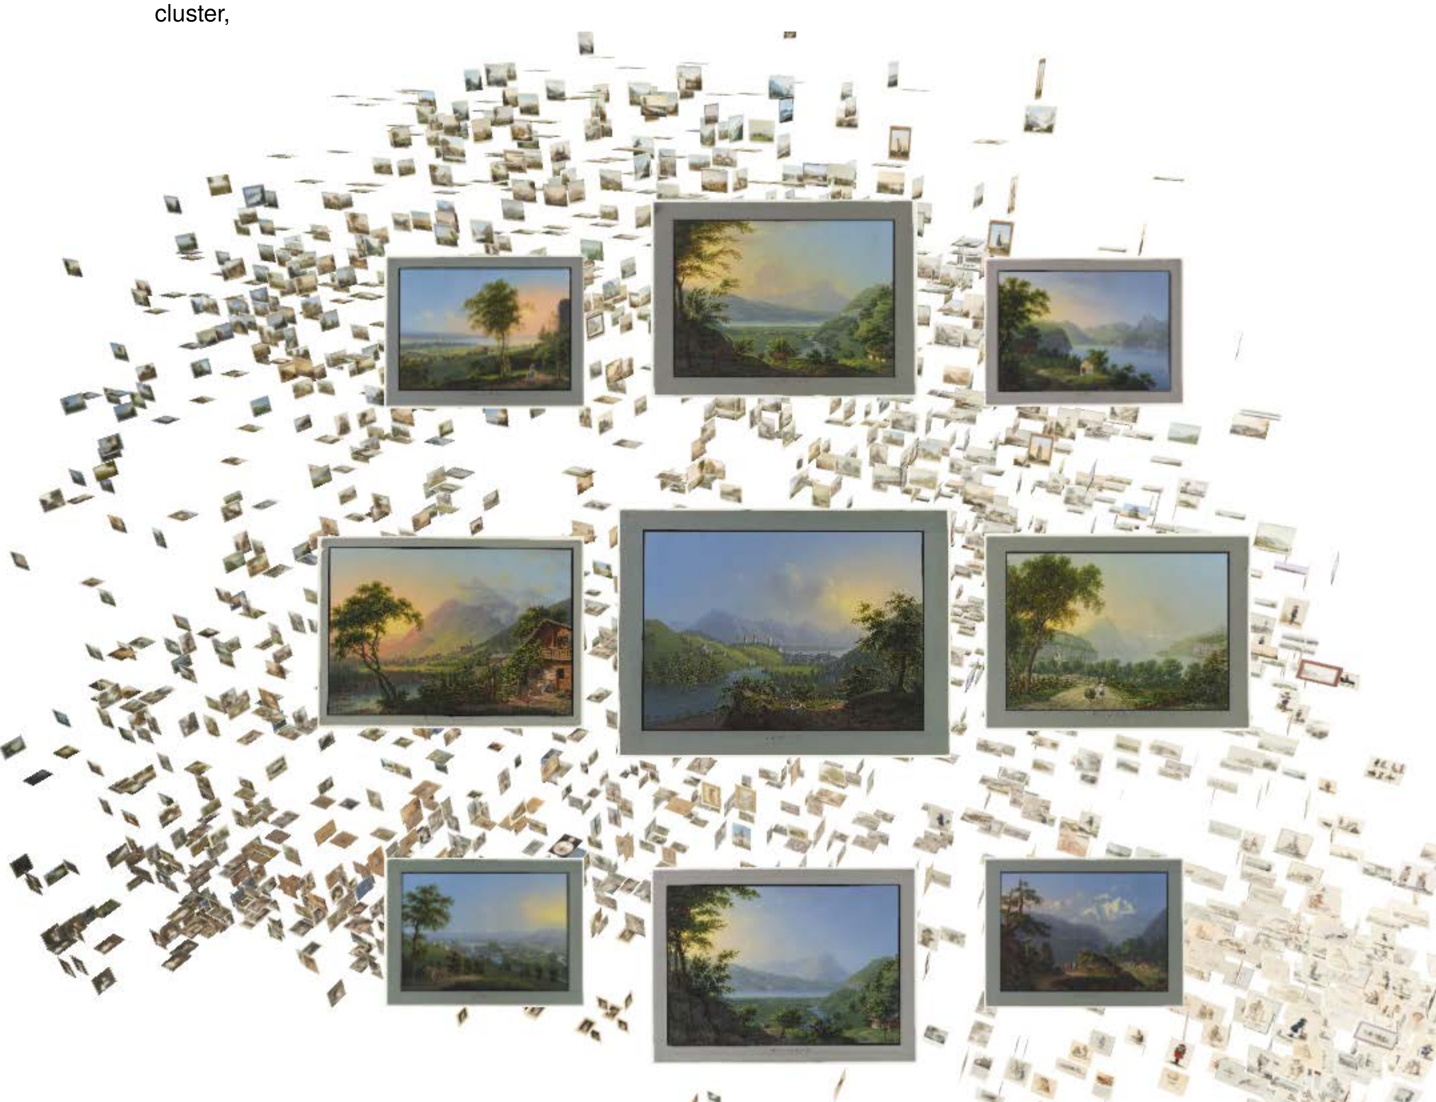

every item is plotted against one other feature. But how can 2300-dimensional space (for the DESCRIPTION sorting) be represented, visualized and made accessible? How can a user explore this space? Sophisticated algorithms from the domain of manifold learning (Pedregosa et al. 2011) are used to project the original high-dimensional space onto a lower-dimensional space, in this case three dimensions. For DESCRIPTION and COMPOSITION, the algorithm used is t-SNE (van der Maarten and Hinton 2008). For the other two, Isomap (Tenenbaum, de Silva, and Langford 2000) proved to give good results. Proximity relations are maintained so that neighbors in the original high-dimensional space remain neighbors in the lower dimensional embedding. It is important to note that the axes of the resulting space don't have any meaning. They do not map individual numerical at-

tributes like hue against saturation or size against time. The position of an item in that space therefore doesn't say anything about that very item. But instead items are expressed in terms of the relational network of neighborhoods they are embedded in. Niches are not treated as anomalies but preserved as richness. Class affiliations can be deduced but are not presumed. Also, there is no intention to extract means from the entire collection like the average painting. In order to illustrate this emergence of clusters beyond categories, I would like to explain in a bit more detail one of the findings of this method, for the COMPOSITION sorting.

Figure 2 is not a visualization from the Gugelmann Galaxy, but shows the two-dimensional embedding with the points to an orthogonal grid (to avoid mutual occlusion) using a rather

Figure 4: (l) Image with eight others of the same technique. (r) Same image with its eight closest neighbors in composition embedding: upright figure (person or waterfall) in the center.

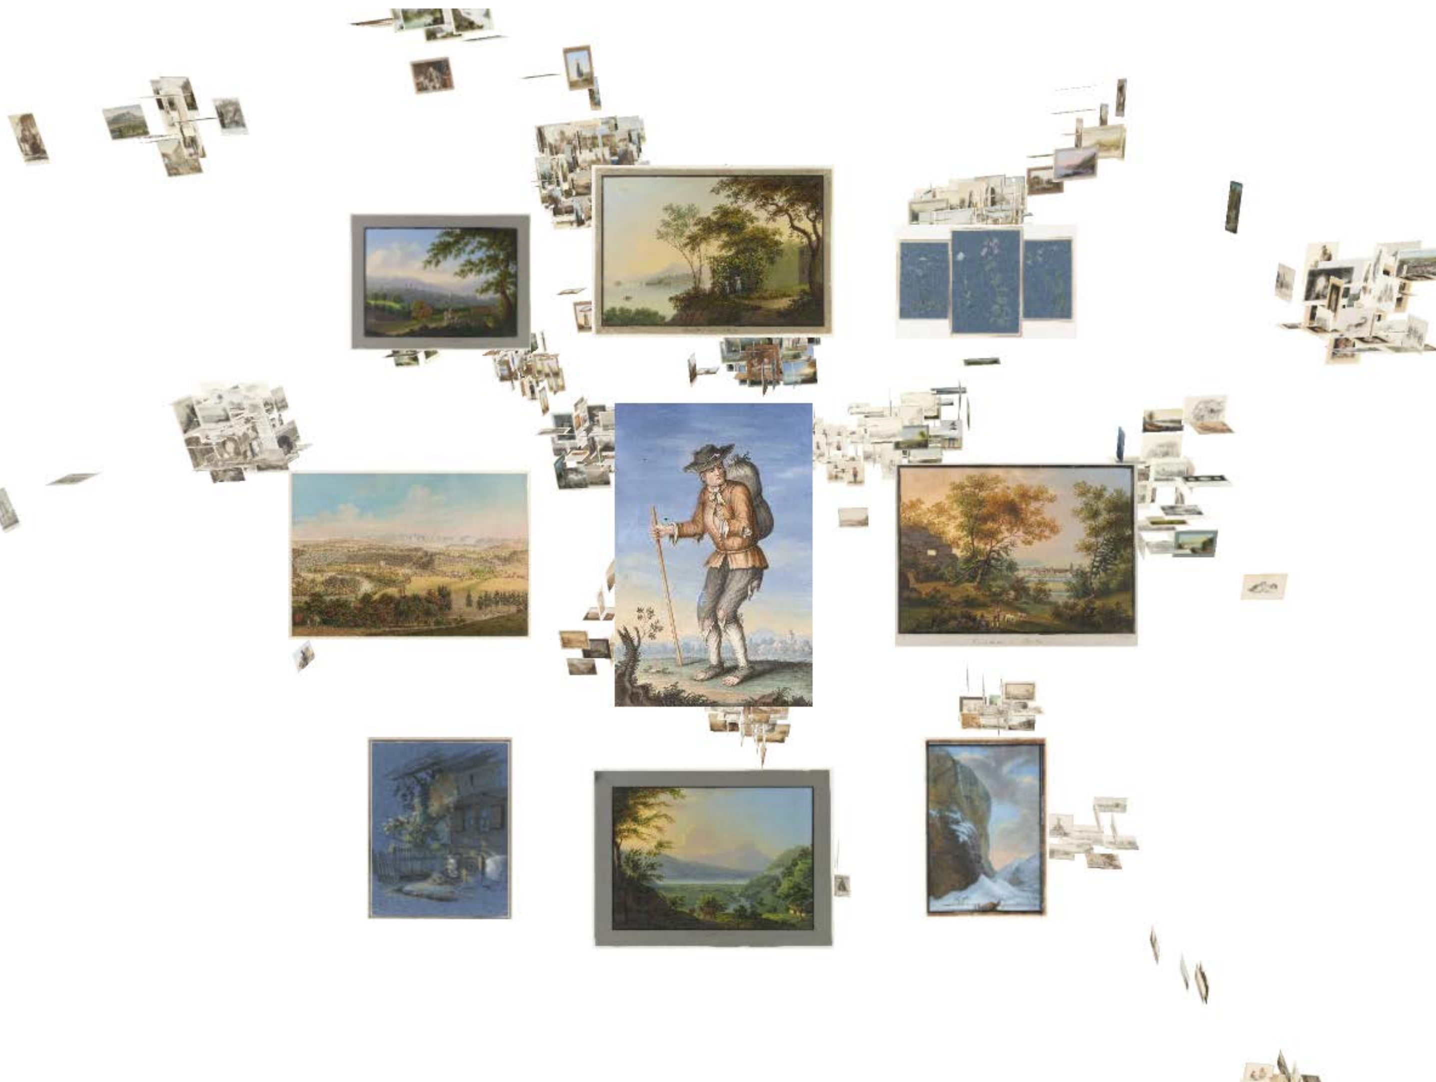

simple method. There exist much more elaborate algorithms for that purpose, e.g. RasterFairy (Klingemann 2016) or IsoMatch (Fried et al. 2015) that I was unfortunately not aware of at the time of working on the project.

Some compositional concepts become obvious at first sight, such as the circular vignettes in the top right corner, upright figures on a neutral background a bit farer left, wide landscapes in the bottom left or images maintaining a frame despite the cropping in the bottom right. Some of these clusters (landscapes, soldiers on white background or the women portraits in folk costumes) already emerged from the sorting according to color distribution. A notable improvement can be illustrated by a group of images at the bottom, a bit left of the center.

All the colored thumbnails[8] show a tree on the left side in the foreground and an open landscape, often with a lake, in the background. The similarity of these images' composition could by no means be discovered in the metadata. The images are by different authors, date from different years and represent different places. Also the color clustering would not arrange them in close proximity because they are very different in sky color and overall tonality. Even a very light pencil sketch (2nd row, right) can be found in this cluster, something that from a computer vision system point of view – that is, dealing uniquely with a matrix of numbers for red, green and blue values – really is something different than a full colored painting.

Figure 5: Stills from YouTube video (https://youtu.be/3O6OfSyn7_4): (l) The author wearing a VR headset. (r) Render from within the goggles with right and left eye view.

Results

As a final product, these three-dimensional projections are rendered as a cosmos of free floating images. The first version was called Schweizer Kleinmeister – An Unexpected Journey (Bernhard et al. 2015), was developed on the occasion of the hackathon and is a desktop application created in Processing[9]. The user can freely navigate by rotating, zooming and panning through the cloud. Figures 3 and 4 show exemplary stills from that version. Once an image is selected, it can be displayed surrounded by its eight closest neighbors in the currently selected embedding.

With the aim of making the experience even more immersive and physical, the second version was made in a game engine called Unity[10]. The Gugelmann Galaxy can now be explored by means of a virtual reality headset called Oculus Rift.

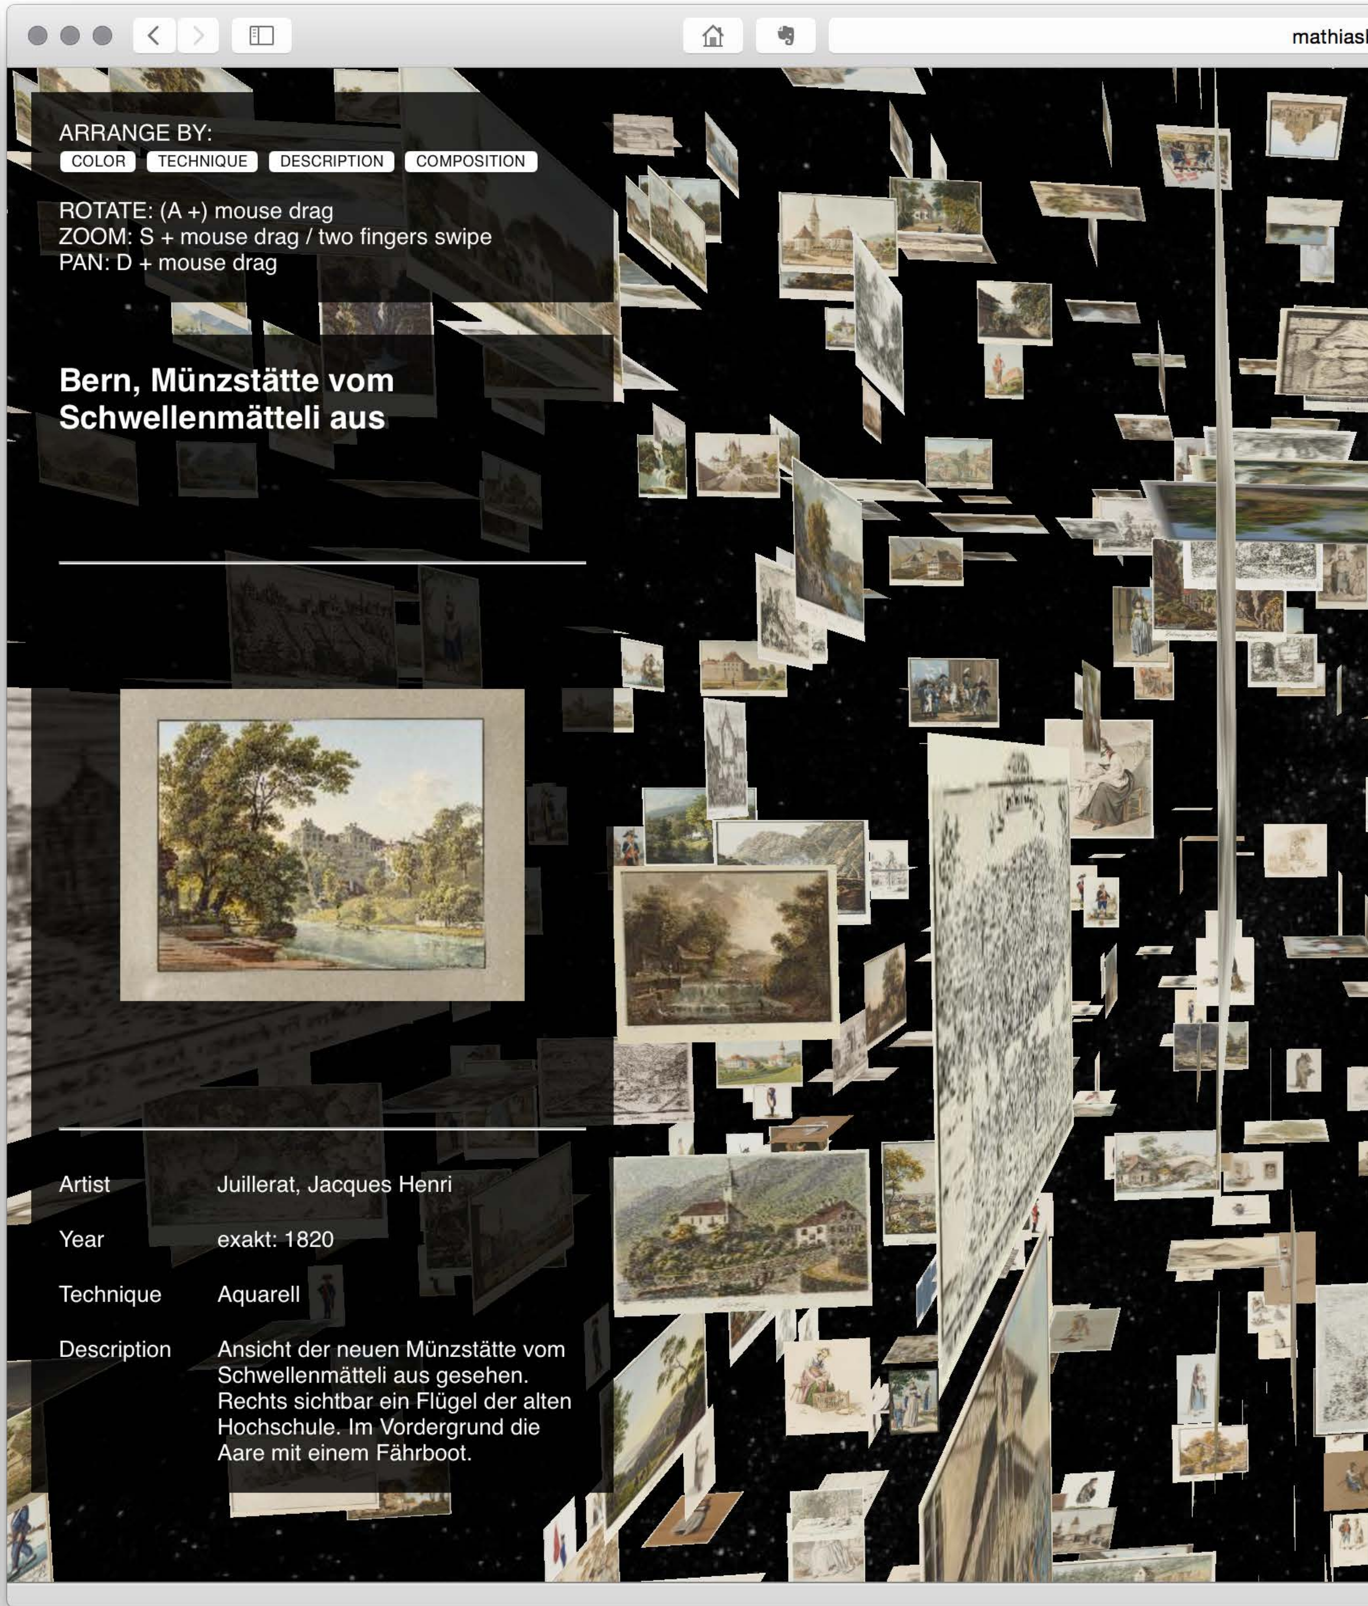

Figure 6: Screenshot of browser application “Gugelmann Galaxy“


bernhard.ch

To meet the demand for increased accessibility, I created a third version of the Gugelmann Galaxy running in the online. It can be accessed from anywhere by anyone without the need to install software, just by pointing the browser to the address www.mathiasbernhard.ch/gugelmann. The user can navigate through the galaxy by rotating, zooming and panning and can rearrange the configuration according to the four different criteria described above. When the mouse hovers over an item in space, it is highlighted and a frame on the left lists a higher resolution version of the image and the available metadata (see Figure 6: screenshot of browser application "Gugelmann Galaxy").

The four different views on the collection (TECHNIQUE, DESCRIPTION, COLOR and COMPOSITION) are meant to be only prototypical implementations, demonstrating a methodology. Many more questions can of course be formulated, limited solely by the curiosity of the users. Whatever the interest, the collection can be projected into an arrangement specific for that very query.

Conclusion

Elaine Gurian begins her paper The Essential Museum (Gurian 2006) with the question "What if our profession created a museum in which visitors could comfortably search for answers to their own questions regardless of the importance placed on such questions by others?" She goes on in describing it as "another kind of museum, one that arises not from organized presentations by those in control, but one that puts control into the hands of the user. [...]Unfettered browsing of objects will be the main organizing motif in this museum and to facilitate that, the majority of the museum's objects will be on view." Such a proposal questions the role of the curator and the installed hierarchies. This idea is further elaborated in our magazine article ANY-FOLD: On Curation, Literacy and Space (Bernhard, Marinčić, and Orozco 2015). A project like the Gugelmann Galaxy can in many ways provide loose ends to connect to in order to pursue this goal. It is not bound to a physical location like a museum building – a potential mental barrier for many audiences. It does not impose one specific and unchangeable organization. Instead, it facilitates access and empowers curios minds to dig for personal nuggets. It does not assume an indefeasible set of categories but relies on gradual statistical correlations. Statistics show up in artists' work every now and then. There is for example the collective Guerilla Girls, who furnish evidence for the gender gap with percent figures. (Freeland 2001) Or the Russian-American conceptualist artist duo Vitaly Komar and Alexander Melamid who made professional polls in various countries asking people for their favorite color, format or motif in an artwork. Evaluating the results, they created the most wanted and most unwanted paintings of these countries, co-authored by the majority vote.

But while inevitably a certain weird interestingness has to be attested to the lonely George Washington standing in front of a lake with meadowing deers (Komar and Melamid 1994), probably nobody really finds his or her particular interest reflected in it because everybody's interest is supposed to be reflected. The assumption leading to this meaninglessness is that the features are all independent. They even publish the numerical results on their website but only the sums. Looking at correlation of the various features one could make more specific conclusions. The absurdity of making everything flat and unspecific is what the artists play with and should be taken with a grain of salt. With projects like the Gugelmann Galaxy, statistical approaches to an entire collection of artworks, the interest is neither in the examination of a single item, nor in boiling it down to some global insight. The marvel such a framework offers us is to observe individuality always within and defined by the position within all the others.

Little Big Data

The focus of this second issue of the DAH Journel is Visualizing Big Image Data. I think it is necessary to relativize the hyped, ubiquitous and watery term of big data. The Gugelmann collection, which serves as a base for the work presented here, contains roughly 2300 items. In orders of magnitudes that is 10^3 items. Other collections of art works available are in the range of 10^5 items. There are 101'086 items in the WikiArt (formerly WikiPaintings) dataset (Karayev 2013), which for Saleh and Elgammal (2015) "to the best of our knowledge, is the largest online public collection of digitized artworks". Further, there are the roughly 200'000 items in the YourPaintings (BBC 2014) collection or the 120'000 records on items at the Museum of Modern Art (MoMA 2015) recently published on GitHub. The Gugelmann collection dataset is hence a rather small one among the art collection datasets. However, the datasets used by search engines or social media companies to train their classifiers are again 2-4 orders of magnitude bigger. The dataset provided for the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) contains 14'197'122 (10^7) images labeled with thousands of categories (ImageNet 2010). The body of image raw material Facebook can work with is estimated to amount to 200'000 images uploaded every minute (Horaczek 2013). So to summarize it drastically, twice as many images are uploaded to Facebook every minute than WikiArt collected from 800 years of art history. As Peter Norvig, head of Google research points out in his lecture The Unreasonable Effectiveness of Data (Norvig 2011), for a lot of tasks the search engine giant has to deal with, more is more. The scientific method, feature engineering and more conventional machine learning algorithms that work well for smaller datasets get easily outperformed by algorithms such as multi-layer neural networks through the mere number of training samples available.

Outlook

User behavior like detail views or fly routes can be recorded to further improve the network of links between the items, similar to how search engines like Google learn from user clicks what a suitable answer to a query might be. After having looked at or even liked, starred, favored (what ever the evaluation metrics in place) enough paintings, such a system could even learn my taste and make a ranking or proposals beyond any category. Why am I only given the choice to love or hate all of photography? Why do I have to decide upon one specific century to make a choice – let alone how should I know what –ism[11] suits me best before I know anything about styles? To keep the motivation high, the user needs the feeling to be in control. It is rewarding to be able to make new discoveries. No absolute treasure is hidden in the archives and collections until we ask our very specific questions. As more and more researches from different fields join forces and create more and more elaborate systems, digital art history heads for a very interesting future.

Notes

1 Galleries, libraries, archives and museums

2 Since 1982, the SNL has been home to the Gugelmann Collection, consisting of over 2300 drawings, prints and paintings by the Schweizer Kleinmeister - Swiss 18th century masters - assembled by Annemarie Gugelmann and her brother Rudolf. It is one of the most valuable donations the NL has ever received. This unique collection is continually being expanded with significant new acquisitions, and constitutes the essential core of Swiss iconography.

3 All three are architects, computer scientists and PhD students at the chair for Computer Aided Architectural Design at ETH Zurich.

4 Following the proposal of Lev Manovich in the first issue of this journal (Manovich 2015), I will also refer to these attributes as features. Features can be integers, floating point numbers or text and they are organized in the data table's columns. Individual images are called items or samples of the dataset and are organized in rows. In a later section, I will not write about features but about dimensions instead. This should facilitate the understanding of the concept multi-dimensional spaces and vectors.

5 The words are: aquarell, aquatinta, gouachiert, koloriert, unkoloriert, bleistift, roetel, farbkupferstich, federzeichnung, laviert, gouachemalerei, holzschnitt, kohle, kreide, kupferstich, lithographie, pinselzeichnung, radierung, sepia, stahlstich, tuschezeichnung and umrissradierung

6 Similar as for TECHNIK described above, 1 if the word is present, 0 if it is not

7 TF-IDF short for term frequency–inverse document frequency, is a numerical statistic that is intended to reflect how important a word is to a document in a collection or corpus. It is often used as a weighting factor in information retrieval and text mining. The TF-IDF value increases proportionally to the number of times a word appears in the document, but is offset by the frequency of the word in the corpus, which helps to adjust for the fact that some words appear more frequently in general. (Wikipedia)

8 The grayed and blurry images are similar as well, maybe without the tree. Their appearance could be the result of the inferior quality algorithm used to push the items to a grid layout.

9 www.processing.org Processing is a flexible software sketchbook and a language for learning how to code within the context of the visual arts. (Processing website)

10 www.unity3d.com Unity is a cross-platform game engine developed by Unity Technologies and used to develop video games for PC, consoles, mobile devices and websites. (Wikipedia)

11 There are 83 –isms among a total of 164 styles distinguished by the WikiArt platform.

Bibliography

BBC. 2014. "YourPaintings." Accessed 6. Aug. http://www.bbc.co.uk/arts/yourpaintings/paintings/search.

Bernhard, Mathias, Nikola Marinčić, and Jorge Orozco. 2015. "ANY-FOLD: On Curation, Literacy & Space." *trans*, 84-87.

Bernhard, Mathias, Nikola Marinčić, Jorge Orozco, and Sonja Gasser. 2015. "Schweizer Kleinmeister: An Unexpected Journey." opendata.ch. http://make.opendata.ch/wiki/project:schweizer_kleinmeister:an_unexpected_journey.

Cleveland, Carter, and Sebastian Cwilich. 2010. "The Art Genome Project." https://www.artsy.net/about/the-art-genome-project.

Crowley, Elliot J., and Andrew Zissermann. 2014a. "In Search of Art." In *Workshop on Computer Vision for Art Analysis, ECCV*.

Crowley, Elliot J., and Andrew Zissermann. 2014b. "The State of the Art: Object Retrieval in Paintings using Discriminative Regions." British Machine Vision Conference.

Dalal, N, and B Triggs. 2005. "Histograms of Oriented Gradients for Human Detection." IEEE Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, CA, USA.

Elgammal, Ahmed, and Babak Saleh. 2015. "Quantifying Creativity in Art Networks." The Art and Artificial Intelligence Laboratory, Department of Computer Science, Rutgers University.

Erbuer, Moenen, and Tobias Boonstoppel. 2012. Curiator - The World's Greatest Collaborative Art Collection New York, NY.

Freeland, Cynthia A. 2001. *But Is It Art?* Oxford, New York: Oxford University Press.

Fried, Ohad, Stephen DiVerdi, Maciej Halber, Elena Sizikova, and Adam Finkelstein. 2015. "IsoMatch: Creating Informative Grid Layouts." *Computer Graphics Forum (Proc. Eurographics)* 34 (2).

Gurian, Elaine Heumann. 2006. "The Essential Museum." Museums Aotearoa Conference, Hawke's Bay, New Zealand.

Horaczek, Stan. 2013. "How Many Photos Are Uploaded to The Internet Every Minute?" *Popular Photography*, 6. Aug. http://www.popphoto.com/news/2013/05/how-many-photos-are-uploaded-to-internet-every-minute.

ImageNet. 2010. "Summary and Statistics." Last Modified April 30, 2010 Accessed 6. Aug. http://www.image-net.org/about-stats.

Karayev, Sergey. 2013. "Vislab - Computer Vision Group." Berkeley, University of California Accessed 6. Aug. http://vislab.berkeleyvision.org/tutorial.html.

RasterFairy 1.0.2. Klingemann, Mario (Quasimondo), GitHub.

Komar, Vitaly, and Alexander Melamid. 1994. UNITED STATES: MOST WANTED PAINTING. Dia Center for the Arts.

Konvitz, Ezra. 2012. ArtStack.

Kräutli, Florian. 2015. "MoMA on GitHub." *YYYY-MM-DD Time/Data/Visualization*. http://research.kraeutli.com/index.php/2015/09/moma-on-github/.

Manovich, Lev. 2015. "Data Science and Digital Art History." *International Journal for Digital Art History* (1).

MoMA. 2015. The Museum of Modern Art (MoMA) collection data. doi:10.5281/zenodo.21147.

Norvig, Peter. 2011. *The Unreasonable Effectiveness of Data*. YouTube: UBCCPSC.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. "Scikit-learn: Machine Learning in Python." *Journal of Machine Learning Research* 12:2825-2830.

Saleh, Babak, Kanako Abe, Ravneet Singh Arora, and Ahmed Elgammal. 2014. "Toward Automated Discovery of Artistic Influence." In *Multimedia Tools and Applications*. Springer US.

Saleh, Babak, and Ahmed Elgammal. 2015. "Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature." Department of Computer Science, Rutgers University.

Schlecht, Joseph, Bernd Carqué, and Björn Ommer. 2011. "Detecting Gestures in Medieval Images." IEEE International Conferene on Image Processing (ICIP 2011), Brussels, Belgium.

Tenenbaum, Joshua B., Vin de Silva, and John C. Langford. 2000. "A Global Geometric Framework for Nonlinear Dimensionality Reduction." *Science* 290 (5500):2319-2323. doi: 10.1126/science.290.5500.2319.

van der Maarten, L.J.P., and G. E. Hinton. 2008. "Visualizing High-Dimensional Data Using t-SNE." *Journal of Machine Learning Research* 9:2579-2605.

WikiArt. 2010. WikiArt - Encyclopedia of fine arts. http://www.wikiart.org.

Yarlagadda, Pradeep, Antonio Monroy, Bernd Carqué, and Björn Ommer. 2012. "Towards a Computer-Based Understanding of Medieval Images." In *Scientific Computing and Cultural Heritage*.

Mathias Bernhard is a PhD student at the chair for Computer Aided Architectural Design (CAAD), ETH Zurich. His research focuses on the synthesis of novel artifacts by recombining and learning from vast collections of precedents. He studied architecture at the École Polytechnique Fédérale de Lausanne (EPFL) and at the Eidgenössische Technische Hochschule Zürich (ETH). After his diploma, he worked as an assistant at the rapid architectural prototyping laboratory (RAPLAB) on the application of digital fabrication in research and teaching. After finishing a master of advanced studies (MAS) and two years in a interdisciplinary team applying state of the art building technology research on the planning of a real world building on the ETH campus (ArchTecLab), he started his doctoral studies at the chair for CAAD. He has also been teaching classes for bachelor, master and post-graduate students ever since starting his work at the ETH Zurich.

Correspondence e-mail: bernhard@arch.ethz.ch

Gugelmann Galaxy

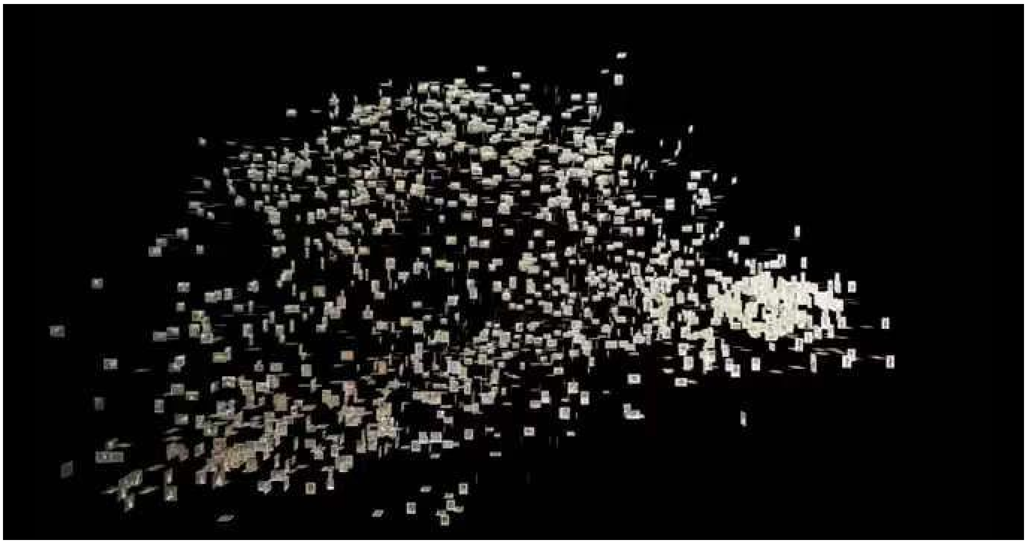

Artistic Data and Network Analysis

Figure 1. Aby Warburg's Panel 45 with the color version of the images mapped over the black-and-white photographic reproductions.

Peer-Reviewed

Images as Data: Cultural Analytics and Aby Warburg's *Mnemosyne*

Stefka Hristova

Abstract: In this paper, by extending the methodology of media archaeology to the praxis of Cultural Analytics/Media Visualization I ask how have we compared multitude of diverse images and what can we learn about the narratives that these comparisons allow? I turn to the work of Aby Warburg who attempted to organize close to two thousand images in his Mnemosyne Atlas. In comparing contemporary methods of image data visualization through cultural analytics method of remapping and the turn of the century methodology developed by Warburg under the working title of the "iconology of intervals," I examine the shifts and continuities that have shaped informational aesthetics as well as data-driven narratives. Furthermore, in drawing parallels between contemporary Cultural Analytics/Media Visualization techniques, and Aby Warburg's Atlas, I argue that contextual and image color data knowledge should continue to be important for digital art history. More specifically, I take the case study of Warburg's Panel 45 in order to explore what we can learn through different visualization techniques about the role of color in the representation of violence and the promise of prosperous civil society.

Keywords: images as data, Aby Warburg, Cultural Analytics, color, visualization, violence, reconciliation

In their current state, the methods of Cultural Analytics and Digital Humanities provide two radically different ways of interpreting images. Cultural Analytics calls upon the understanding of images as objects with features, and more specifically as image-data in and of themselves. The potential of this computational method to digital art history was presented in detail in Manovich's essay "Data Science and Digital Art History."[1] Digital Humanities methodologies on the other hand rarely analyze images per se. Instead, they tend to focus on the metadata: historical and cultural information about the artifact. In this project, I take on a hybrid Digital Art Historical methodology that combines Cultural Analytics as articulated by Lev Manovich with Digital Humanities paradigms. I apply this method to the case study of Aby Warburg's *Mnemosyne Atlas* from 1924, and to Panel 45 more specifically. The twenty-four images in Panel 45, many of which were created by the Italian fresco painter Domenico Ghirlandaio, comprise the objects for this study. I selected Panel 45 because it grapples with the relationship between color and violence. Using Cultural Analytics, I render Warburg's images as color-data. I argue that this color-data, situated in the theoretical framework provided by Christopher Johnson's 2013 Digital Humanities project about Warburg titled *Mnemosyne: Meandering through Aby Warburg's Atlas*[2] and in reference with traditional art historical texts about Domenico Ghirlandaio, speaks to the larger tropes of violence and civil reconciliation. I demonstrate that color-data becomes art historical knowledge through an in-depth conversation with digital humanities practices as well as with already established disciplinary knowledge.

Warburg's *Mnemosyne*

Aby Warburg worked on the *Mnemosyne Atlas* project from 1924 until his death in 1929. The *Atlas* consisted of over two thousand black and white photographs of works of art arranged on "sixty-three wooden boards, measuring approximately 150x200cm, covered with black cloth."[3] According to Warburg, these arrangements of diverse types of images attempted to produce "first of all an inventory of pre-coined classical forms that impacted upon the stylistic development of the representation of life in motion in the age of the Renaissance."[4] It aimed to offer a "comparative analysis" of visual and cultural tropes spanning from Classic Antiquity to Renaissance, and engaged the early 20th century as well. This analysis was built upon an understanding of not only of the formal elements of the images, but also of their cultural meaning.[5] While his panels contained no textual descriptions, Warburg kept notes that listed provisional titles as well as the overarching theme of each panel.

The black-and-white photographs of artworks were pinned either directly to the black fabric, or framed over a white mate. They varied in size, proportion of the depicted image, as well as distance from each other. Warburg sometimes photographed the work of art in its entirety. In other instances, he focused on a detail and included only a close-up. Through the use of the close up, he cropped and framed elements that can be then positioned as points of emphasis within the panels. The panels themselves have not survived and what we have left instead are black and white photographs of these panels instead.

Panel 45

Christopher Johnson conducted an extensive study of Warburg's work. His research resulted in both a book manuscript titled *Memory, Metaphor, and Aby Warburg's Atlas of Images* and a Digital Humanities Project titled *Mnemosyne: Meandering through Aby Warburg's Atlas.* The book illuminates Warburg's use of the metaphor as a structuring element. Here Johnson argues that Panel 45 is a "study in contrasts.[6]" As cited in Johnson book, Warburg's notes on this panel read:

> Superlatives of the language of gestures. Wantonness of self consciousness. Individual heroes emerging out of the typological grisaille. Loss of the 'how' of metaphor. At the center of the panel's twenty-four images are two frescoes by Ghirlandaio from the Torenabouni Chapel, Massacre of the Innocents and Apparition of the Angel to Zechariah.[7]

The Digital Humanities project *Mnemosyne* offers further insight into the specific images that constitute the panels.[8] Here the twenty-four black-and-white photographs of artworks that constitute Panel 45 are named and numbered. They are also linked to the color version of each artwork. Johnson's interpretation of the panel reads:

> A study in extremes, panel 45, builds on the sequence of panels 41, 41a, 42, 43, & 44. But it does so to signal a perilous 'loss' of metaphoric distance. Here frescos by Ghirlandaio emblemize the 'afterlife' of classical 'expressive values' in the Renaissance. Yet Ghirlandaio's use of the *grisaille* technique is not able to fully moderate or mediate the intensity of the passions. Thus even as the *Massacre of the Innocents* metonymically yields to *The Blood of the Redeemer* (while also anticipating the Eucharistic theme of panel 79), and even as the serving girl in *Birth of John the Baptist* heralds the all-important theme of the *nymph* whose 'life in motion' [*bewegtes Leben*] animates Warburg's own thinking, the panel heightens the lethal threats (e.g., plague, tyranny, war) against the possibility of achieving psychological 'balance' [*Ausgleich*].[7]

According to Johnson, this panel explores the tension between violence and passion on one hand and the aesthetic formal properties of *grisaille* – a technique that relies on varied in brightness gray palettes and desaturated colors - on the other.[9] Here Warburg explored the role of grisaille technique as well as color in conveying violence.

Building upon Johnson's work, I argue that Panel 45 is indeed a comparative study of contrasts – contrasts that engage both visual and cultural tropes. Warburg explores violence and promise of civil reconciliation rather than passion; grisaille and brightly saturated color, Classic Antiquity and Renaissance. Warburg articulates these dyads by positioning images of rebirth on the left side of the panel – such as the *Birth of John the Baptist* by Domenico Ghirlandaio, and images of violence on the right – such as *Massacre of the Innocents* by Matteo di Giovanni. As the two frescoes represented in the center of Panel 45 show, the possibility of the social is rooted in a long-standing tradition of violence.

Panel 45 in Color

The extensive documentation provided by the Digital Humanities *Mnemosyne* Project (2013) allows for the reconstitution of this panel in color (Figure 1). Cultural Analytics further confirms the importance of the question of color in Warburg's work. I argue that color is an important agent of meaning in Panel 45 not only because of its primacy as a feature for data extraction in Cultural Analytics, but also because it was the primary diving principle in the organization of Warburg's thought. Warburg questions the transience of violence and the promise of social rebirth in the Classic as well as Renaissance periods through an exploration of both color and *grisaille.* Although panel 45 consists of black and white photographs, I argue that he was working with and conceptualizing through works of art in color. The black-and-white effect is due to the photographic reproduction technology deployed in the constitution of the panels. The *Atlas* images for Warburg were pictures of intense color.

Digital Humanities projects, such as Johnsons's *Mnemosyne,* engage with the context in which Panel 45 with its images exists. They offer relevant information on the level of *metadata* about images such as: author, period, title, etc. Remapping visualization techniques allow for a closer look at the *data* that constitutes the images: namely brightness and saturation. Image visualization thus allows us to further analyze the relationship between color and gray scale in relation to violence and possibility of civil reconciliation that Warburg positioned in first place as his motivation behind the composition of Panel 45.

Cultural Analytics as a Methodology

Having translated Panel 45 into color, I turn to the methodology of Cultural Analytics in order to convert its images into color-data. Cultural Analytics is a methodology spearheaded by Lev Manovich and the Software Studies Initiative, which "allows the users to work with different kids of data and media all shown together."[10] It adapted Information Visualization techniques used in STEM to the fields of Art History, Film and Media Studies, and Popular Culture, to name a few. Information visualization, as defined by Lev Manovich, is "a widely used as a tool for understanding data – i.e. discovering patterns, connections, and structure" and as such delivers "new knowledge about the world through systematic methods – such as experimentation, mathematical modeling, simulation." Furthermore, this new method is invested in design, as "it involves the visual presentation of data in a way that facilitates the perception of patterns."[11]

Deploying visual databases and graphic plotting software (*ImageJ*, which originally was developed by the National Institute for Health), the Cultural Analytics approach allows for the comparison of large sets of visual data through techniques previously deployed by STEM: namely "visualization, visual analytics, data mining, and information visualization."[11] A prerequisite for the methodology of Cultural Analytics is an understanding of "culture as data" that can be "mined and visualized" – a move that distills visual culture into a set of data patterns.[12] In their seminal essay "How to Compare One Million Images?" Manovich, Douglass, and Zepel write that "having at our disposal very large cultural data sets which can be analyzed automatically with computers and explored using interactive interfaces and visualization" opens up new ways of understanding culture.[13] Tracing the grayscale, brightness, hue, and saturation of images on macro as well as micro scale allows for the articulation of trends across contemporary visual landscapes as well as across historical periods. In short, Cultural Analytics provides a computational algorithmic answer to the question: "How to Compare One Million [Diverse] Images?" It is important for new paradigm of exploring humanistic big data however to continue to emphasize the cultural context of this new visual landscape in order to offer a critical interpretation of the significance of this algorithmic knowledge.

In comparing images, whose origin lie in user-generated media, television, film, as well as photography, and art broadly defined, Lev Manovich proposes three major strategies: (1) Collection Montage – or the creation of image-grids where images are arranged in neat rows and columns with uniform size and spacing through the Montage function of the *ImageJ* software; (2) Temporal and Spatial Sampling in which the archive of images is sampled either based on the date of the image (temporally) or based on the selection of portions of images (spatially); (3) Remapping – the plotting of images in order to identify patters across temporal media artifacts or series of artifacts. Thus, Cultural Analytics, with its methodology for media visualization, "is based on three operations: zooming out to see the whole collection (image montage), temporal and spatial sampling, and remapping (re-arranging the samples of media in new configurations)."[14] In this project, I am particularly interested in the possibilities for shaping digital art historical methodologies centered on color that Remapping offers, as it is allows for a wide range of visual landscapes to be drafted.

Remapping Panel 45

In the context of Cultural Analytics, color becomes one of the key feature components through which analysis is conducted and data attained. Coupled with art historical knowledge, it provides a crucial look at the role of color in the articulation of violence as well as the possibility of civility (Figure 2). In remapping Panel 45, I worked with the

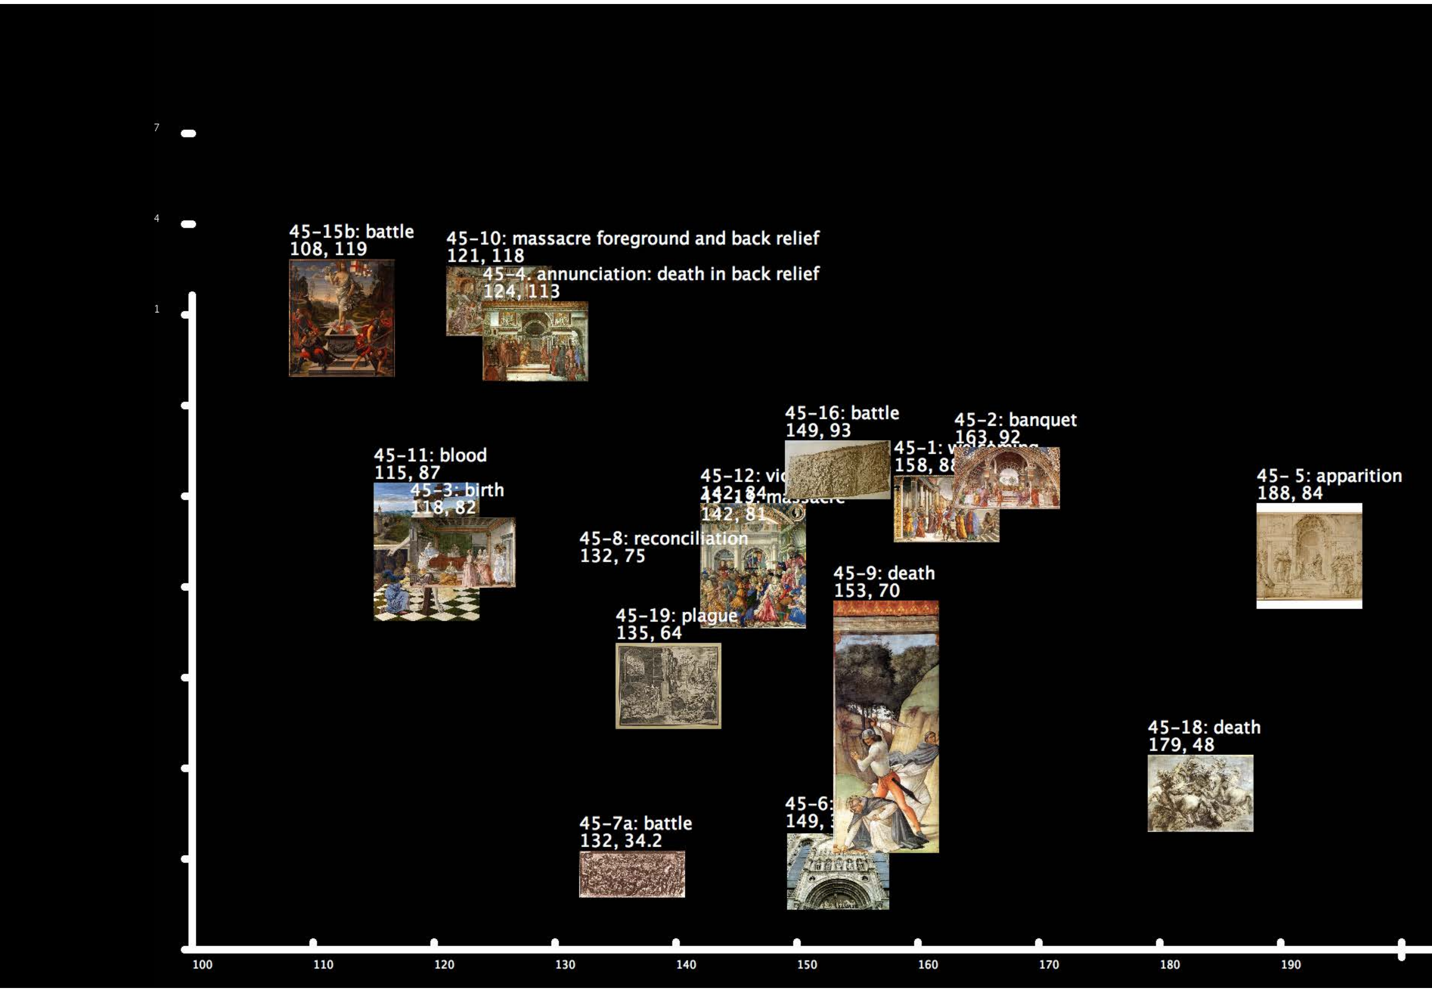

Figure 2. Remapping the images of Panel 45, using the ImageMeasure and ImagePlot modules of the ImageJ Software.

color versions of the images. Using the ImageJ's module ImagePlot, developed by the Software Studies Institute at UCSD and CUNY Graduate School, I first ran analytics on the following features: Grayscale, Brightness, Saturation, and Hue. Exceptions were made in the case of image 17, which I was unable to locate in color, and image 14, which included too much black or white negative space around the actual round-shaped image and thus escaped accurate numerical representation. This data was then mapped onto a black canvas with the same size as the photographs of Warburg's panel – namely 650x847 pixels (the canvas size is a combination of the plotting area of 550x747 pixels and 50px border space). Next, whereas Warburg had the ability to choose images with different size and scale, ImagePlot's insists on a uniform image size. I calculated the size of the images in the photograph of Warburg's panel and then reduced that size in half for legibility purposes. I arrived at the average width of 65 pixels for the size of all images in the remapping. The X-axis was used for the distribution of the brightness via *brightness_median* (average of grey scale values for the pixels in an image) while the Y-axis was used to display the saturation via *saturation_median* (average of saturation, i.e. purity of color for each pixel in an image). Even though *brightness_median* and *saturation_median* range between the values of 0 and 255 according to the standard RGB schema, where as 255 designates highest brightness or white and purest saturation, I chose a smaller range that reflected the subset of color available in Panel 45. I limited the X-axis brightness range to 100-200 and the Y-axis saturation range to 20-150. This restriction indicates that all images included fall in the mid range in terms of brightness and in the lower to mid-range of saturation, given the overall digital RGB color spectrum. In prefacing my argument about the need of art historical context for understanding image-data, I included to additional variables. First, I added a parameter indicating whether the title of the artwork signals associations with violence or civility and reconciliation. Second, I included the image numbers in the plotting of the data in order to keep reference to the title and author of the works. The image number and metadata correlation can be found in the Appendix to this article.

The remapping plane became a landscape of color in which the images coming closest to a *grisaille* technique lie in its lower right corner, while the images with highest brightness and saturation lie in its upper left corner. Thus *grisaille* images occupied the space of the lower Y-axis, while images with bright colors clustered around the upper Y-axis. As described earlier in the article, the *grisaille* technique attempted to subdue color by decreasing saturation to a point of gray monochrome composition and by working within varied attributes of brightness within this single shade in order to express meaning. In other words grisaille artwork can exhibit black, white, as well as shades of gray as colors and thus tends to stay altogether in the low spectrum of the saturation range –hence on the lower Y-axis. Grisaille images at the same time vary in brightness and thus span along the horizontal X-axis. The images coming closest to a *grisaille* technique on the remapped canvas lie in the lower level of the Y-axis. In this case they are 45-7a – color image of a bronze fresco, 45-6 – color image of sculpture, 45-18 – color image of an engraving, and 45-19 – color image of copperplate engraving. These images lie on the outer rim, while 45-4 and 45-10 take the central stage.

Surprisingly, the two images with highest saturation and brightness levels are positioned adjacent to each other both at Warburg's Panel 45 as well as in the Remapping of the panel. Images 45-10 (*Massacre of the Innocents*, Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel Santa Maria Novella) and 45-4 – same as 4a and 4b – (*Annunciation to Zacharias*, Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel, Santa Maria Novella) lie at the heart of the Warburg's Panel 45. They are positioned in the vertical and horizontal center of the panel. These two images are of the same size. And both are larger, thus more prominent than the rest of the images included in Warburg's panel (See Figure 1). In the remapped canvas, they are positioned together at the top left corner. This positioning indicates that they carry the highest brightness and saturation values of all images included in Warburg's original canvas. One image represents violence, the other – civil reconciliation.

Art Historical Account of Ghirlandaio's frescoes

Massacre of the Innocents and *Annunciation to Zacharias* are among the six frescoes coming out of Tornabuoni Chapel, Santa Maria Novella in Florence, Italy. Both were created by Domenico Ghirlandaio and his helpers in 1485. As Art Historian Steffi Roettgen writes,

> Iconographically, the *Massacre of the Innocents* forms an antithetical counterpart of the *Adoration of the Kings,* the one constituting one of the seven joys of the Virgin, the other one of her seven sorrows....In the Bible narrative of the massacre of the children in Bethlehem... it took place a year after the magi had returned home and Joseph and Mary had fled to Egypt. Herod, having heard that a king of the Jews had been born in Bethlehem, had all of the year-old male infants of the city killed – among them, unwittingly, his own son, who had been entrusted to a nurse.[15]

This dramatic fresco portrays the massacre both in the foreground through the main scene as well as on the background as part of the architectural relief included in the scene. Taking a closer look at this artwork demonstrates that Ghirlandaio deploys *grisaille technique* in both planes. In the background plane *grisaille* is associated with architecture. In the foreground plane *grisaille* comes to represent the living dead: both people and animals appear through desaturation pale and bloodless. In the foreground plane, in contrast to the *grisaille* of flesh, dress appears to be in bold colors such as gold, blue, purple, and crimson flow across the image.

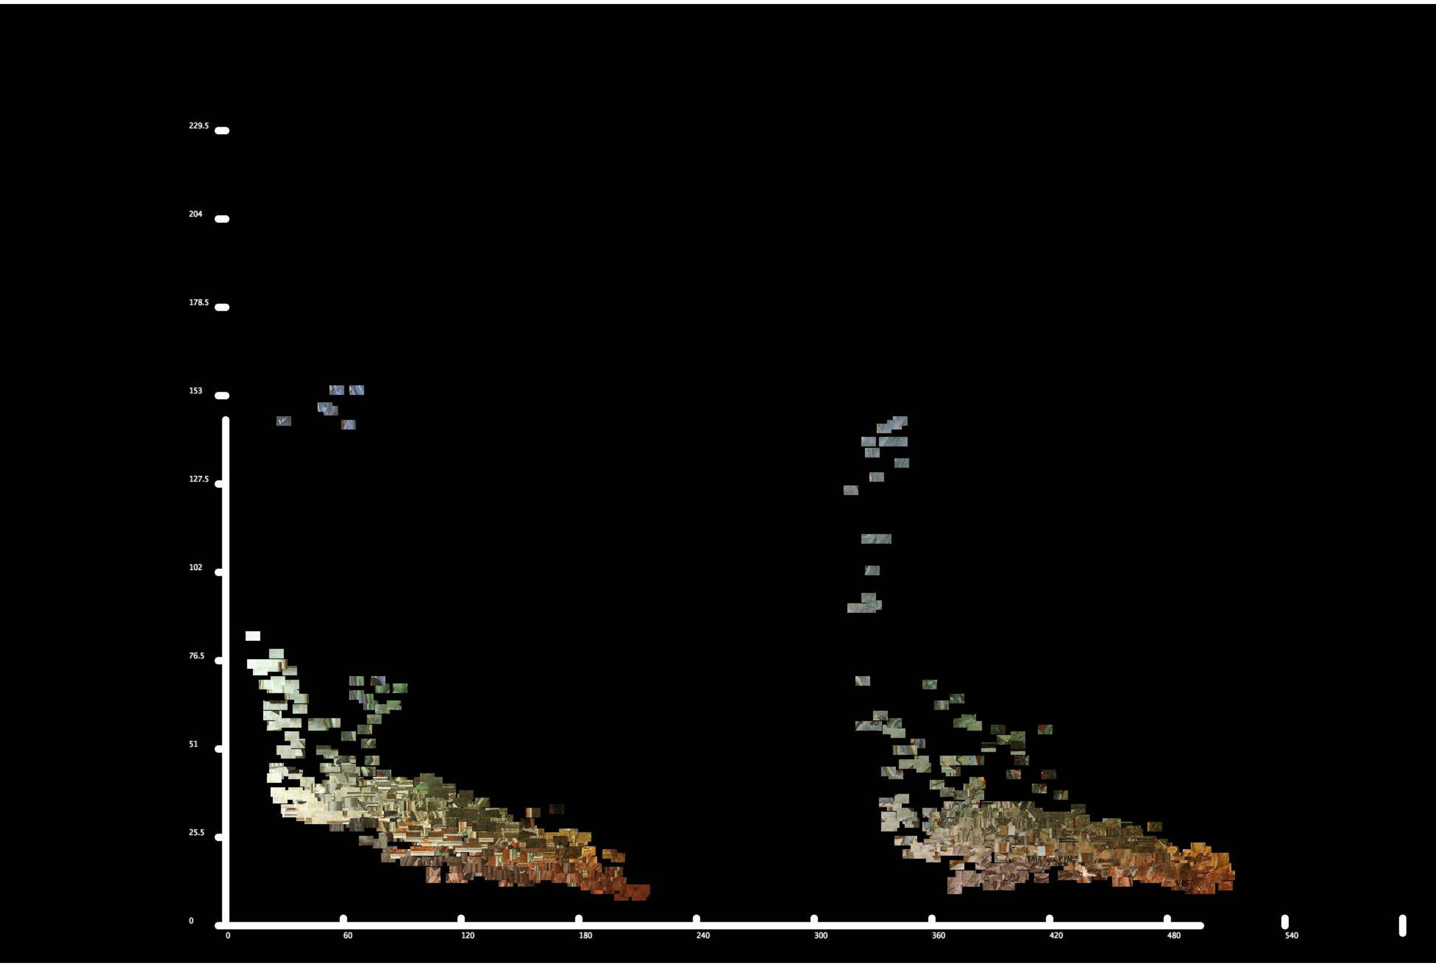

Figure 3. Visualization the hue and saturation in images 45-4 (*Annunciation*) on the left and 45-10 (*Massacre*) on the right

Where as the *Massacre* is a fresco representing totalizing violence, *Annunciation to Zacharias* signals the beginning of salvation and the possibility of civil reconciliation. This fresco, featuring the portraits of "several generations of the Tornabuoni family" as representatives of Florence, celebrates the prosperity of 1490's Florence as well as the assurance that with God's will salvation will come.[16] According to the story of Zacharias, here, the *promise of salvation* lies in prophesy of his unborn son. Warburg included three versions of this image – one of which presents a tightly cropped version of the fresco to exclude the architectural framework of the chapel, and another, which under the name Angel Appearing to Zacharias presents a drawing which was presented to the patron by the painter for approval, before beginning to paint.[17]

Yet the prosperity of the city and the promise of salvation is founded upon death and violence. This triumph is evident in the inscription on the background architectural detail of the city wall, which reads: "In the year 1490, when the most beautiful cities, owing its wealth, its conquests, its undertakings and buildings, enjoyed prosperity and peace".[18] Florence perceived itself as the new Rome and in order to assert its prominence and legitimacy, sought to situate itself within the Classical tradition. Warburg hints at this tie as well by positioning Ghirlandaio's fresco of *Brutus, Scaevola, Camillus,* which signals precisely the establishment of Rome, immediately to the right of the *Annunciation* in Panel 45. The Classical tradition, portrayed in the background architectural details through *grisaille*, includes violent scenes of triumph in keeping up with the legends of the establishment of Rome. The civility and prosperity of Florence as the new Rome is thus secured first through the legitimacy of the violet triumph in Classical Tradition, thus contemporaneously, as well as simultaneously through the destruction of the enemy as seen in the *Massacre of the Innocents* fresco.

Ghirlandaio's frescoes *Massacre of the Innocents* and *Annunciation to Zacharias* lie at the heart of Warburg's panel as well, with fragments of the *Annunciation* repeated and scaled three times. The centrality of these two works in Warburg's panel as well as their positioning in relation color, thus *grisaille* in the Remapping, articulate violence and civility/prosperity, not violence and passion as Johnson argues as the major tropes of this panel. Where as Johnson argues that "Ghirlandaio's use of the *grisaille* technique is not able to fully moderate or mediate the intensity of the passions," I argue based on the data provided by the Remapping, that Ghirlandaio, deployed similar techniques in the portrayal of both sorrow and joy, death and rebirth. While these observations are meaningful, taking a closer look at the art historical background provided by Warburg of both of Ghirlandaio's frescoes: *Massacre of the Innocents* and *Annunciation to Zacharias,* reveals a powerful critique of both the institution and visual representation of civil society founded upon violence, waged by Aby Warburg.

Situated in the context of art history, the image-data produced through the remapping of Panel 45 demonstrates that images representing violence as well as civility/prosperity used color in similar fashion. The proximity of *Massacre of the Innocents* and *Annunciation to Zacharias* confirms Warburg's inquiry into the value of color in representing violence. Reading these two images as image-data prompts a budding observation: in the age of Renaissance, death and civil reconciliation do appear to be portrayed with similar color footprint. Were we to expect grisaille being more present in images of death or destruction, reserving color for the heroes that emerge out of total violence? The remapping of Panel 45 illustrates one of Warburg's central concerns: "Individual heroes emerging out of the typological grisaille." And further, imbricates this concern in the aesthetics of color and the tropes of violence and civil reconciliation.

Color-Data in Art History

Traditional art historical accounts allows us to understand the techniques and material basis of the colors, while Cultural Analytics provides an indispensible overview of the quantity or amassment of hues used. In painting the chapel, Ghirlandaio used *all of the colors available:* "di pingere facere et exornare cum omnibus coloribus ut vulgariter dicitur posti in frescho."[19] As Julia DeLancey argues, the artist used "a palette of fully saturated, pure hues, employing few pastel or dark colour and generally only white in order to create relief, lending the cycle a vibrant and legible appearance."[20] It is interesting that Ghirlandaio's work in Tornabuoni Chapel is notable in terms of its pure hues and bright saturations. Yet two of Ghirlando's brightly colored frescos were chosen by Warburg in a discussion about *grisaille* – hence grayness and desaturation. Cultural Analytics sheds additional light into this paradox.

In translating Ghirlandaio's *Massacre of the Innocents* and *Annunciation to Zacharias* into color-data in and of themselves, patters of similarities between the two images are revealed. I created a second remapping that focused on the color visualization of these two images, based on the features of hue and saturation. I first resized them to the same width of 800 pixels. I then divided each image into 1024 pieces (each piece containing 25 pixels of color data, both images amounting to 2048 pieces. These 2048 segments became the objects of the second remapping that I conducted. The color-data in each segment was measured and mapped accordingly: Y-axis represented hue (*hue_median*) on the scale of 0 to 255, while the X-axis showed the saturation (*saturation_median*).

Both Ghirlandaio's *Massacre of the Innocents* (representing total violence) and *Annunciation to Zacharias* (depicting the possibility of civil reconciliation) exhibit similar hue/saturation distributions (Figure 3). Their color blueprints signal dominance of red, orange, and green highly saturated hues, as well as green-blue highly desaturated elements. This color-data gains significance when interpreted through the framework fo art history. The color footprints demonstrate not only articulations of the tropes of violence and the promise of civil rebirth in terms of foreground and background planes, but also in terms of temporal dissonance: Renaissance and Classic Antiquity in

Renaissance thought respectively. Our movement across the horizontal X-axis thus becomes movement through time. Desaturated *grisaille,* almost white chips come to signal the perception of the Classic Antiquity in the late 1400s and the deeper reds and oranges represent the full saturation hues of the Renaissance. This visualization illuminates comparatively the way color has come to structure violence and the possibility of reconstruction across two different temporal periods (Figure 3).

The two frescoes indeed feature *grisaille* architectural elements that operate as backdrop to the scenes that mirror both in content as well as technique sculptures or engravings of violence. The backdrop of both frescoes projects Classic Antiquity into the age of Renaissance. In representing "classical motifs and inscriptions, battles, cavalcades and warriors," these frescoes allow for an investigation of the contemporaneity of *grisaille.*[21] This dualism is theorized by Christopher Johnson. He articulates tension between form and content is manifested in the image itself where "the slaughter colorfully depicted in the foreground is repeated monochromatically in the scenes portrayed on the triumphal arch in the background."[22] Through repetition of visual motifs both within the image as well as within the panel - a close up fragment of this same image was pinned on the panel as well, Warburg illustrates how "the Renaissance language of gestures is prefigured in classic antiquity."[23] Cultural Analytics further extends this argument to mirroring by providing a comparative framework of the images as color-data.

Color, Violence, and the Promise of Rebirth of Civil Society

Moving both literary as well as symbolically back and forth between *grisaille* representations of violence in the Classical tradition and pure saturated colors in the Renaissance, Warburg's Panel 45 poses important question about the role of color in articulating memory and history as well as distance from the past in which contemplation about the present and belief in the future can take place. As Giorgio Agamben has argued, for Warburg "just as the creation and enjoyment of art require the fusion of two psychic attitudes that exclude each other ('a passionate surrender to the self leading to a complete identification with the present – and a cool and detached serenity which belongs to the categorizing contemplation of things])."[24] The examination of *grisaille* in relation to color through projects informed by Digital Humanities and Cultural Analytics techniques prompts us to further ask questions: what is the role of color in allowing for or precluding contemplation; how are we to understand contemporary acts of totalizing violence; how do we represent hope for the possibility of social and well as spiritual humanistic reconstruction?

The oscillation between identification and contemplation that Warburg proposes structures the comparative methodology that accounts for both form and content: hence both *data* and *metadata* associated are imbricated in a framework of duality. It is in this movement within images as well as in-between images, in other words in the intervals within and in-between the visual, that critical thinking can be provoked. In the context of Cultural Analytics, this space of contemplation too is visible in the movement in-between bits of visual data. But as much as we try to break down images into unique molecules, into the smallest units of information possible, even those miniscule units are complex, multilayered, and dynamic, thus not stable and self-evident. This is particularly relevant to the study of visual media, as even a single pixel is invested in color and every color carries within itself complex historical and theoretical implications.[25] Relational visual knowledge is thus manifested not only in the *edges* (lines in-between plotted data) that connect the *nodes* (the dots that plot data) in practices of information/media/visualization, but within the images that constitute these nodes in and of themselves.

Contextual Visual Knowledge

In articulating Digital Art History through methodologies that analyzes data and metadata, I attempted to integrate cultural and historical knowledge with data science. While adapting Big Data quantitative paradigms to the study of art, media, and popular culture allows for important new ways of seeing the still and moving image, I argue that this mode of visualization is always and already imbricated in a complex network: one that is not purely algorithmic, but also socio-economic, political, and last but not least historical. I argue that Digital Art History should continue its investment in contextual visual knowledge by combining quantitative image-data paradigms with traditional art history in order to foster critical interpretations of visual culture.

Appendix

Panel 45. Image 1. Mary enters the Temple, Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel, Santa Maria Novella.

Panel 45. Image 2. Herod's Banquet, Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel, Santa Maria Novella.

Panel 45. Image 3. Birth of John the Baptist, Domenico Ghirlandaio, fresco, 1485-90;

Florence, Tornabuoni Chapel, Santa Maria Novella.

Panel 45. Image 4a&b. An Angel appears to Zacharias, [Annunciation to Zacharias] Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel, Santa Maria Novella.

Panel 45. Image 5. Angel appears to Zacharias, Domenico Ghirlandaio, preparatory sketch, ca.1485; Vienna, Albertina Museum.

Panel 45. Image 6. Offering in the Temple with a Hercules Relief (?) on the Altar, Jacopo and Tommaso Rodari, sculpture, 1491-1509; Como, Como Cathedral.

Panel 45. Image 7.a. Battle (with Hercules), Bertoldo di Giovanni, bronze sculpture, ca. 1478; Florence, Museo Nazionale del Bargello.

Panel 45. Image 7.b. Battle (with Hercules), Bertoldo di Giovanni, bronze sculpture, ca. 1478; Florence, Museo Nazionale del Bargello.

Panel 45. Image 8. Reconciliation of the Romans and Sabines, Bartolomeo di Giovanni, spalliera painting, late 15th cent.; Rome, Galleria Colonna.

Panel 45. Image 9. Death of Peter Martyr (Saint Peter of Verona), Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel, Santa Maria Novella

Panel 45. Image 10. Massacre of the Innocents, Domenico Ghirlandaio, fresco, 1485-90; Florence, Tornabuoni Chapel Santa Maria Novella.

Panel 45. Image 11. The Blood of the Redeemer, Giovanni Bellini, oil on wood, ca. 1460-65; London, National Gallery.

Panel 45. Image 12. Brutus, Scaevola, Camillus, Domenico Ghirlandaio, fresco, 1481-85; Florence, Sala dei Gigli in the Palazzo Vecchio.

Panel 45. Image 13. Massacre of the Innocents, Matteo di Giovanni, tempera on panel, 1482; Naples, Museo Nazionale di Capodimonte.

Panel 45. Image 14.1. Mehmed II, Bertoldo di Giovanni, bronze medal, 1482. Helsinki, National Museum of Finland.

Panel 45. Image 14.2. Mehmed II, Bertoldo di Giovanni, bronze medal, 1482. Helsinki, National Museum of Finland.

Panel 45. Image 15.a. Resurrection of Christ, Domenico Ghirlandaio (design), with Davide and Benedetto Ghirlandaio Gemäldegalerie, tempera on wood, ca.1494; Berlin, Gemäldegalerie.

Panel 45. Image 15.b. Resurrection of Christ, Domenico Ghirlandaio (design), with Davide and Benedetto Ghirlandaio Gemäldegalerie, tempera on wood, ca.1494; Berlin, Gemäldegalerie.

Panel 45. Image 16.a. Drawings of scenes from a relief on Trajan's Column, Workshop of Domenico Ghirlandaio [both scenes depict the flight of the Dacians], ca. 1490; from a sketchbook, *Codice Escurialensis*, 28-II-12, fol. 63.

Panel 45. Image 17. A prisoner is led before the emperor, Maso di Bartolomeo, relief sculpture, 1452; Florence, Palazzo Medici-Riccardi.

Panel 45. Image 18. The Battle at Anghiari, Lorenzo Zacchia, (after Leonardo), engraving, 1558.

Panel 45. Image 19. The Phrygian Plague (Il Morbetto), Marcantonio Raimondi (after Raphael), copperplate engraving, ca. 1520.

Notes

1 Lev Manovich "Data Science and Digital Art History" *International Journal for Digital Art History.* Vol 1. No. 1. 2015. p. 12-35.
2 Christopher Johnson *Mnemosyne. 2013.* http://warburg.library.cornell.edu/
3 Christopher Johnson. *Memory.* P.9
4 Aby Warburg. "The Absorption of the Expressive Values of the Past." Trans. Matthew Rampley. *Art In Translation* 1.2 (July 2009). P. 227-8.
5 Ibid.
6 Johnson, p. 98
7 Ibid.
8 Christopher Johnson *Mnemosyne. 2013.* http://warburg.library.cornell.edu/http://warburg.library.cornell.edu
9 Grisaille, is a technique of monochrome painting in grey or greyish color that Renaissance painters used for effects of sculpture. See Irene Earls. *Renaissance Art: A Topical Dictionary.* Greenwood Press, New York (1987) p. 130
10 Lev Manovich,. "Cultural Analytics: Analysis and Visualization of Large Cultural Data Sets" [2008]. www.manovich.net/cultural_analytics.pdf (accessed September 1, 2014).
11 Manovich "Cultural Analytics" 2008.
12 Ibid.
13 Manovich et al. "How to Compare One Million Images?" [2011] http://manovich.net/index.php/projects/how-to-compare (accessed September 1, 2014).
14 Lev Manovich. "Media Visualization: Visual Techniques for Exploring Large Media Collections"*Media Studies Futures*, ed. Kelly Gates. (Blackwell, 2012).
15 Steffi Roettgen. *Italian Frescoes: The Flowering of the Renaissance, 1470-1510.* (New York, London, Paris: Abbeville Press Publishers, 1997), p. 174.
16 Ibid.
17 Jean K Cadogan. *Domenico Ghirlandaio: Artist and Artisan* (New York and London: Yale University Press, 2000), p. 46.
18 Roettgen 1997, 174.
19 Roettgen 1997, 166.
20 Julia DeLancey. "Before Michelangelo: Colour usage in Domenico Ghirlandaio and Filippino Lippi" *Apollo* 145 (1997), p. 18.
21 Micheletti 1990, 48
22 Emma Micheletti,. *Domenico Ghirlandaio* Florence. Italy: SCALA, 1990). P. 48.
23 Ibid.
24 Giorgio Agamben. "Aby Warburg and the Nameless Science." *Potentialities: Collected Essays in Philosophy.* Trans. Daniel Heller-Roazen. (Stanford: Stanford University Press, 1999), P. 94.
25 The complex history of color is evident in the works of Manlio Brusatin, Philip Ball, and Micheal Pastureau to name a few color historians.

Bibliography

Agamben, Giorgio. "Aby Warburg and the Nameless Science." *Potentialities: Collected Essays in Philosophy*. Trans. Daniel Heller-Roazen. Stanford: Stanford University Press, 1999: 89-103.

Cadogan, Jean K. *Domenico Ghirlandaio: Artist and Artisan* Yale. University Press, New York and London (2000).

DeLancey, Julia. "Before Michelangelo: Colour usage in Domenico Ghirlandaio and Filippino Lippi" *Apollo* 145 (1997): 18.

Huhtamo, Erkki and Parikka, Jussi. "Introduction: An Archaeology of Media Archaeology" *Media Archaeology: Approaches, Applications, and Implications*. University of California Press, Berkeley and Los Angeles (2011).

Johnson, Christopher. *Memory, Metaphor, and Aby Warburg's Atlas of Images*. Cornell University Press, Cornell (2012).

Manovich, Lev. "Cultural Analytics: Analysis and Visualization of Large Cultural Data Sets" [2008]. http://www.manovich.net/cultural_analytics.pdf (accessed September 1, 2014).

Manovich, Lev. "Foreword" in Miguel Lima's *Visual Complexity* Princeton Architectural Press, Princeton (2011).

Manovich, Lev. "Media Visualization: Visual Techniques for Exploring Large Media Collections" *Media Studies Futures*, ed. Kelly Gates. Blackwell, 2012.

Manovich, Lev. "Museums Without Walls, Art History Without Names: Visualization Methods for Humanities and Media Studies" *Oxford Handbook of Sound and Image in Digital Media*, ed. Carol Vernallis. Oxford University Press, Oxford (2012).

Manovich, Lev and Douglass, Jeremy and Zepel, Tara "How to Compare One Million Images?" [2011] http://manovich.net/index.php/projects/how-to-compare (accessed September 1, 2014).

Manovich Lev. "Data Science and Digital Art History" *International Journal for Digital Art History*. Vol 1. No. 1. (2015): 12-35.

Michaud Philippe-Alain. *Aby Warburg and the Image in Motion*. Zone Books, New York [2004].

Micheletti, Emma. *Domenico Ghirlandaio* Florence. SCALA, Italy (1990).

Roettgen, Steffi. *Italian Frescoes: The Flowering of the Renaissance, 1470-1510*. Abbeville Press Publishers, New York, London, Paris (1997).

Warburg, Aby M. "The Absorption of the Expressive Values of the Past." Trans. Matthew Rampley. *Art In Translation* 1.2 (July 2009): 273-283.

Stefka Hristova, Ph.D., is an Assistant Professor of Digital Media at Michigan Technological University. She holds a PhD in Visual Studies from the University of California Irvine. Her research examines the digital visual cultures of war and displacement.Stefka teaches in the Communication, Culture, & Media Undergraduate Program and the Rhetoric, Theory, and Culture Graduate Program.

Correspondence e-mail: shristov@mtu.edu

Figure 1. Jan Collaert I after Jan van der Straet, “The Invention of Copperplate Engraving”, in *Nova reperta*, published by Philips Galle, c. 1600. Engraving, 27 x 20 cm. The Metropolitan Museum of Art, The Elisha Whittelsey Collection, The Elisha Whittelsey Fund, 1949.

Peer-Reviewed

Social Network Centralization Dynamics in Print Production in the Low Countries, 1550-1750

Matthew D. Lincoln

Abstract: The development of a professionalized, highly centralized printmaking industry in northern Europe during the mid-sixteenth century has been argued to be the inevitable result of prints' efficacy at reproducing images, and thus encouraging mass production. However, it is unclear whether such a centralized structure was truly inevitable, and if it persisted through the seventeenth century. This paper uses network analysis to infer these historical print production networks from two large databases of existing prints in order to characterize whether and how centralization of printmaking networks changed over the course of this period, and how these changes may have influenced individual printmakers.

Keywords: printmaking, network analysis

Introduction

The development of a professionalized printmaking industry in northern Europe during the mid-sixteenth century has been singled out as a turning point in the history of early modern reproductive prints. In their landmark overview of European printmaking, *The Renaissance Print*, David Landau and Peter Parshall trace two parallel trends in artistic print production in sixteenth-century Europe: individual, highly innovative printmakers who acted as their own printers and distributors; and, on the other hand, an emerging class of professionalized publishers who coordinated the print production pipeline.[1] By the mid-sixteenth century, "industrialized" houses such as that of Hieronymus Cock in Antwerp and Antonio Lafréri in Rome began to dominate the artistic print landscape. Landau and Parshall posited that the practical considerations of printmaking inevitably led to this kind of centralized production. To make one's business out of selling relatively cheap images, one had to operate at scale. Publishers needed a wide-ranging network of buyers, yes, but also of talented platecutters, suppliers of paper, plates, and presses, as well as artistic collaborators for creating designs for single prints, series, or large book illustration projects. (Activities illustrated by Jan van der Straet in Figure 1) Modern firms separated the roles of *inventor* and *sculptor* as a way to increase the efficiency and scale of production. Large, highly-centralized firms could take full advantage of these affordances of scale. A print publisher who was able to position himself in the core of the larger web of the print market by accumulating enough artistic, material, and social capital would have rosy prospects, indeed.

But does this theory hold for the seventeenth-century century? Current overviews of seventeenth-century print production in the Low Countries offer conflicting assessments.[2] This paper takes a quantitative approach to the question of centralization by applying formal network analysis methods to study databases of existing artistic prints from this period, looking at both network-wide centralization as well as the changing positions of individual printmakers and publishers. The results suggest how the simple incentives of a professionalized production market could result in unexpectedly complex repercussions.

Why centralization mattered

A key structural property of any network is its network centralization, a measure of how evenly or unevenly ties are distributed between its members. In a centralized network, a few key individuals occupy powerful and flexible broker positions. Actors in these positions are able to

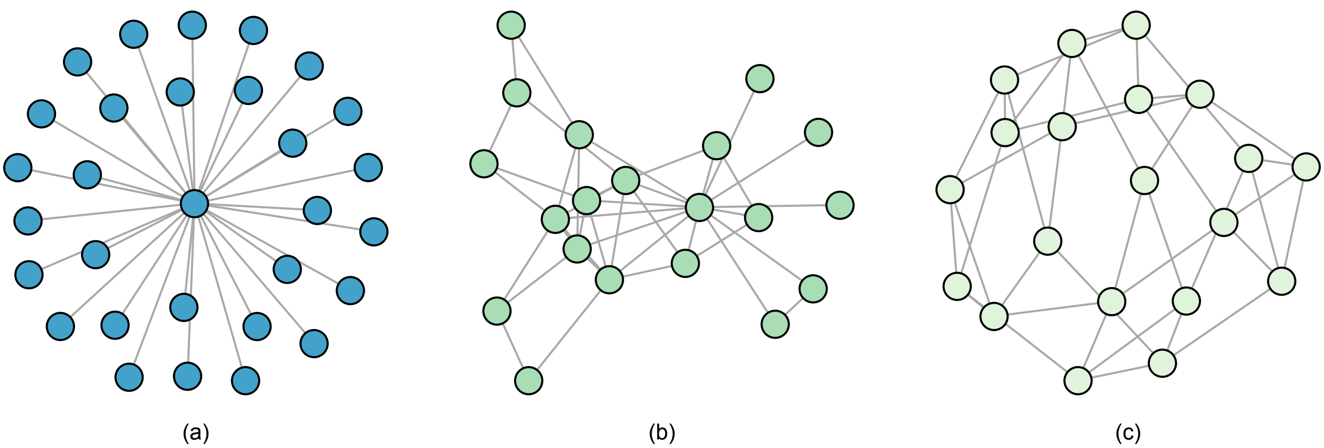

Figure 2. Network centralization examples. (a) is a highly-centralized “star” graph, where one actor receives all the connections. (b) is a relatively centralized graph, where a few nodes receive most of the connections. (c) is a relatively decentralized graph, where most nodes have the same number of connections.

initiate contact with a wide range of individuals they already know, and are granted easy access *through* those immediate contacts to the rest of the network. (Figure 2(a) and 2(b)) Centralized networks are not equal networks, however. The average individual in that same network is unlikely to know multiple well-connected actors, meaning that their access to the rest of the network is mediated and easily cut off by those few highly-central individuals. Conversely, in a more decentralized network (Figure 2(c)), where connections are spread more evenly, a given individual has a better chance of knowing more than one well-connected actor, reducing his or her distance to the rest of the network and making it easier to forge new and diverse connections.

This characteristic of networks is directly relevant to the making of artistic prints. On the one hand, the medium of printing demanded a set of artistic and technical skills, not to mention a set of social connections and financial capital, that presented a barrier to new entrants into the printmaking world. A printmaker could only get so far based on their individual talent with the block knife, burin, etching needle, or mezzotint rocker. Printmakers also needed social capital in the form of connections to publishers who could pay them for their plates, and to artists with artworks or designs whose reproductions were marketable. Publishers gained a competitive advantage by having a large pool of contacts, both with artists as well as with distributors and buyers in domestic markets and at the international book fairs[3] Designers likewise relied on printmakers and publishers to promote and disseminate their artworks All these requirements may have made it more likely that a few centrally-connected individuals would continually increase their number of contacts in a rich-get-richer pattern of increasing centralization.

On the other hand, several factors at both the individual as well as the societal level may have instead encouraged a less centralized printmaking network. A nascent printmaking community with a relatively small population of printmakers and publishers might well have had just a few relatively skilled, experienced, and professionally-established individuals occupying very central positions. But such a structure could be transient. As the number of printmakers increased, those experienced, highly-central players would necessarily take on apprentices and commission less-experienced collaborators. Knowledge is not a fixed quantity; as printmaking opportunities expanded, more and more of those less experienced artists could have learned the technical skills necessary to succeed in the medium. Likewise, a greater number of knowledgeable publishers would be able to create their own local connections without relying on the established knowledge and social connections of a printmaking "elite". As a result, what began as a highly-centralized network could, over time, evolve into a much more distributed one.

A burgeoning Dutch economy in the first half of the seventeenth century may have contributed to this decentralization by supporting print markets in smaller Dutch towns, thus enabling a more decentralized network of print producers with a wider geographic spread. Michael Montias has shown how economic prosperity in the early decades of the seventeenth century may have promoted a burst in painting activity between 1630 and 1650 in Delft, a market that was relatively decentralized, with little institutional patronage compared to other contemporary artistic centers.[4] Might the same effect have changed the production pattern of prints as well?

On paper, these contradictory incentives are both plausible; indeed, they both may have been operating in parallel between 1550 and 1750. Can we determine, though, which incentive (if either) won out? Using the empirical evidence offered by the British Museum and Rijksmuseum print databases, I will show how the balance of these centralizing and decentralizing incentives may have played out over the seventeenth century.

Data on artistic prints

The rich collections of Dutch and Flemish prints in the British Museum (hereafter BM) and the Rijksmuseum (hereafter RKM) present an excellent opportunity to bring quantitative methods to bear on data concerning prints. It must be noted that while these institutions have especially rich holdings, particularly in Dutch prints, no one print collection can be said to be perfectly "representative" of the full range of actual connections between printmakers in the sixteenth and seventeenth century. For example, while these databases have rich information on *constprenten*, or fine art prints, engraved illustrations for books will likely be undercounted. Missing almost entirely from these data are information about the lowest end of print production in illustrated broadsides, playing cards, calendars, and cheap devotional prints.[5] Therefore, the claims of this study will be restricted to the production of fine art prints. Each of these museums also has its own unique collecting history. This distortion can be mitigated, however, by comparing two distinct sources of data about the same phenomenon.[6] By running the same analysis on both datasets, it will be possible to easily compare the results offered by both sources. Where similar results are returned by both the BM and RKM datasets, we can at least reject the claim that the results are *solely* artifacts of collecting preferences specific to each museum.

Second to the source-specific biases of each dataset is the larger question of historical distortion inevitably shared by *both* institutions. This study will have to contend with the same unknown unknowns that plague any historical investigation. Because of paper impressions' fragility, prints likely have a very uneven survival rate, and it would be unwise to take modern day collections as a proxy for the absolute sizes of editions and print runs in the sixteenth and seventeenth centuries.[7] That said, the reproducibility of prints does grant one advantage. Peter Parshall has endorsed the idea that, if the exact sizes of *print runs* are not accurately represented in modern day print collections, we can nevertheless understand a great deal about overall *patterns* of production because the survival rate of a given print that has been reproduced several times is much greater than that of other non-reproducible artifacts. In 1998, Parshall suggested that the voluminous evidence of today's museum print rooms, could be invaluable for research if only it could be aggregated and analyzed fluently.[8] We now have that capability.

Both the BM and the RKM have published the cataloged information for their collections as structured digital data.[9] These data include object-level descriptions such as creators and their various roles, title, date, medium, dimensions, and subject matter. For printed artworks, each database details (when known) the artists who produced the original design for the print, the printmaker who cut the woodblock or plate, and, when applicable, the publisher who printed and distributed the artwork. The BM database describes 14,821 print impressions dated between 1550 and 1750 by Dutch or Flemish artists, while the RKM database contains 19,980 of the same. These numbers represent records that have been assigned dates, and which also have at least two identified creators.[10] Each database also contains biographical information on the creators associated with these prints, including life dates and classifications by nationality.

Inferring print production networks

From these individual artwork records, it is possible to construct a digital model of a network that represented the inferred social connections between these artists. In this network, artists (the nodes of the graph) are connected when they are associated with the production of a print in one of three general roles: *designer* (either as an active participant, or simply "made after"), *printmaker*, and/or *publisher*.[11] A single print may thus provide a basis for connecting designer, engraver, and publisher nodes at a particular point in time. Dozens or hundreds of prints support the construction of a larger network. The network will be projected like so:

1. Create small sections, or time-slices, of the production network as it may have existed at different points in time.

2. From these time-slices, calculate the centralization of the network as a whole.

3. For each individual within each of these time-slices, calculate various metrics for that given individual at that point in time.

Rather than construct a large network looking at every print in the BM and RKM databases at the same time, I will use a rolling window approach to construct many "slices" of the network as it existed at different points in time. (Figure 3) For example, a slice of the network between 1640 and 1650 would include all artists who were alive at some point during that ten year interval.[12] These artists will

Figure 3. Visualizing the method for creating time slices of the historical print production network. Each node in the graph is an artist, engraver, or publisher, and links are formed when two artists both worked on the same print.

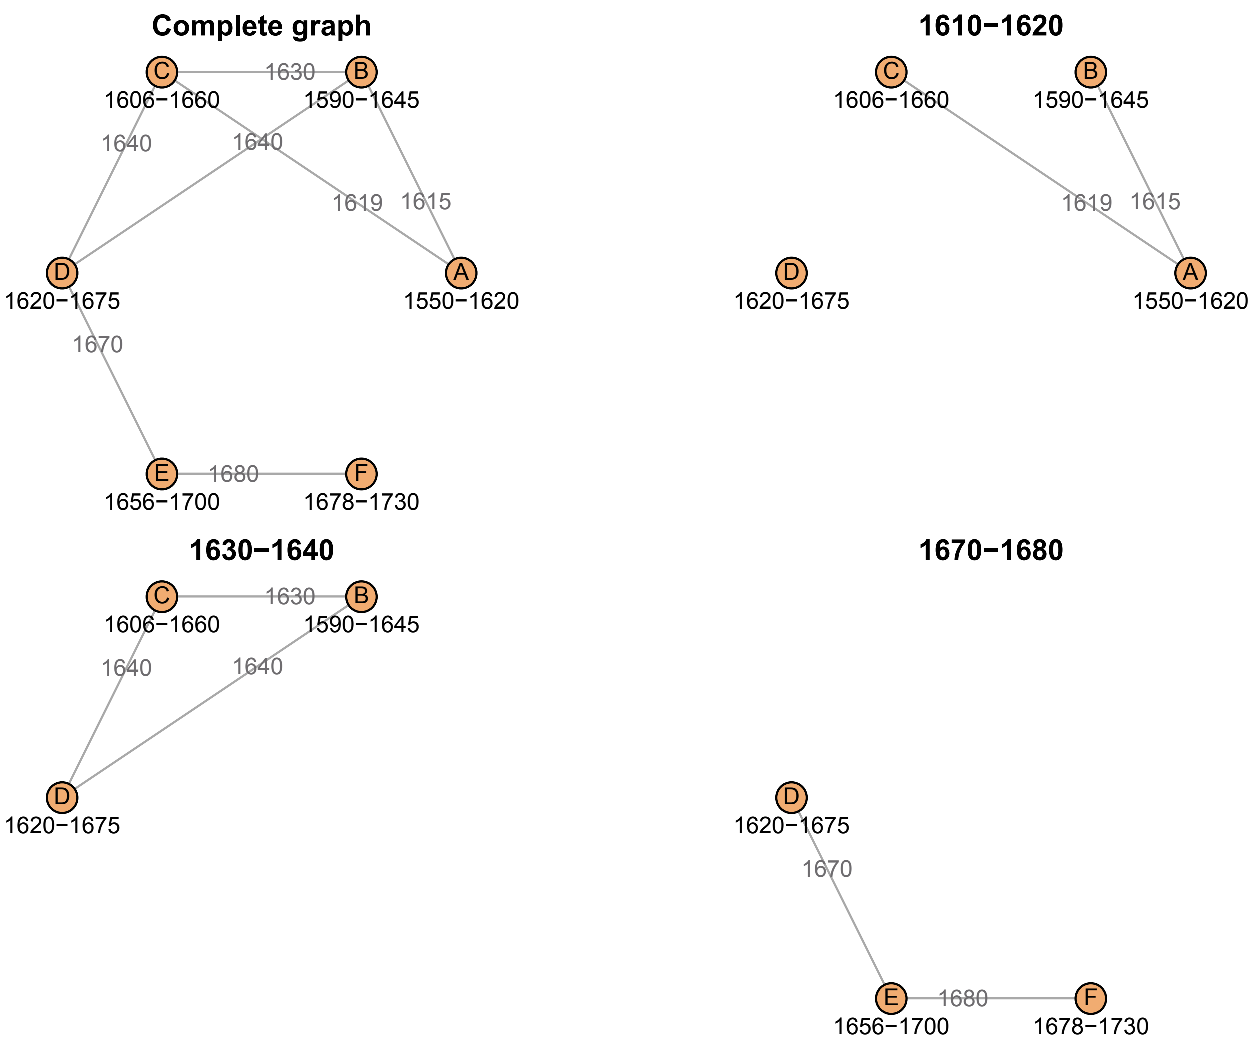

only be connected by edges derived from what prints were being produced during that same 1640–1650 interval.[13] (Figure 4) The resulting network will be unweighted and undirected.[14]

Having created national sub-networks from each time slice, we can compute their graph centralization scores.[15] The higher score, the more centralized a graph is; the lower the score, the more distributed it is. It is also possible to characterize centralization at the scale of the individual. An individual node's degree centrality characterizes how many different connections it has to other members of the network.[16]

Results: A mixed message

Figure 5 displays the changing network centralization score for the Dutch and Flemish communities of print producers between 1550–1750, contrasting that metric with the numbers of nodes and edges in each network over the same period. Results from both the BM and RKM datasets are overlaid. While there are several local differences between the BM and RKM results, they are generally consistent with each other, providing some measure of confidence that the result we are seeing is partially

Figure 4. Visualizations of the core components of two Dutch network time slices.

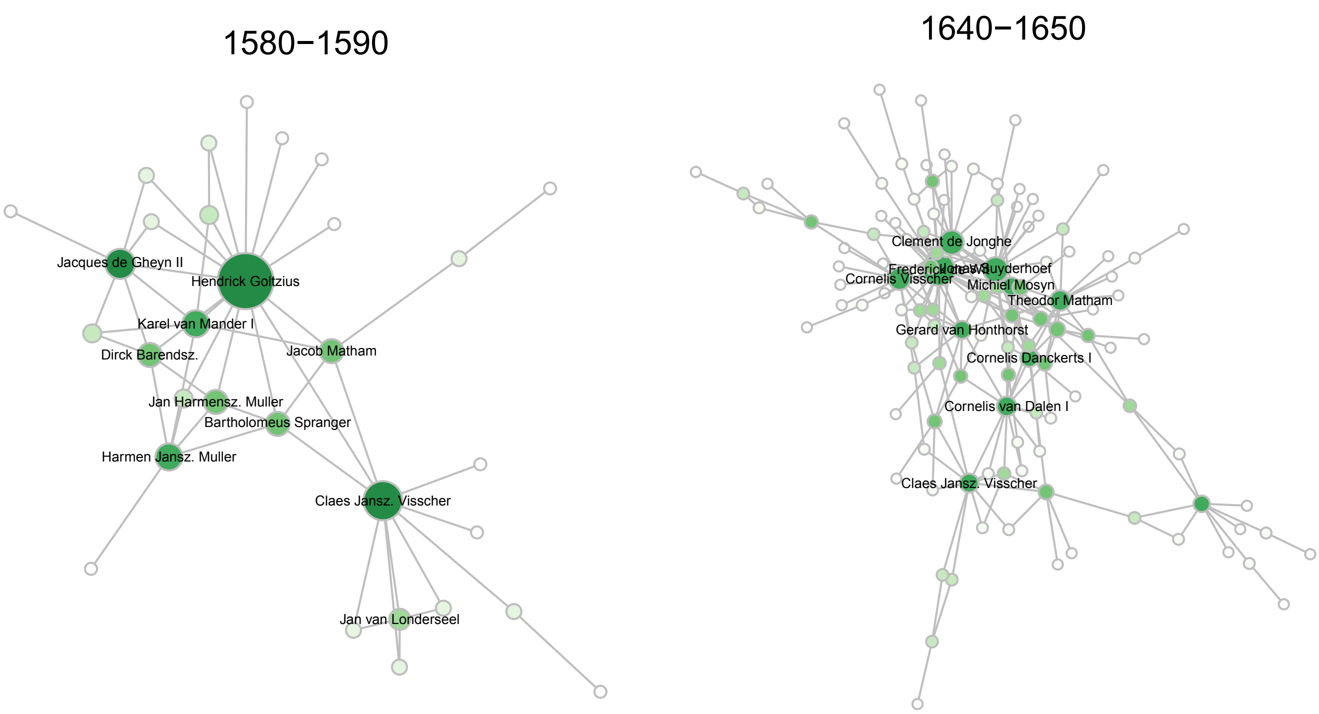

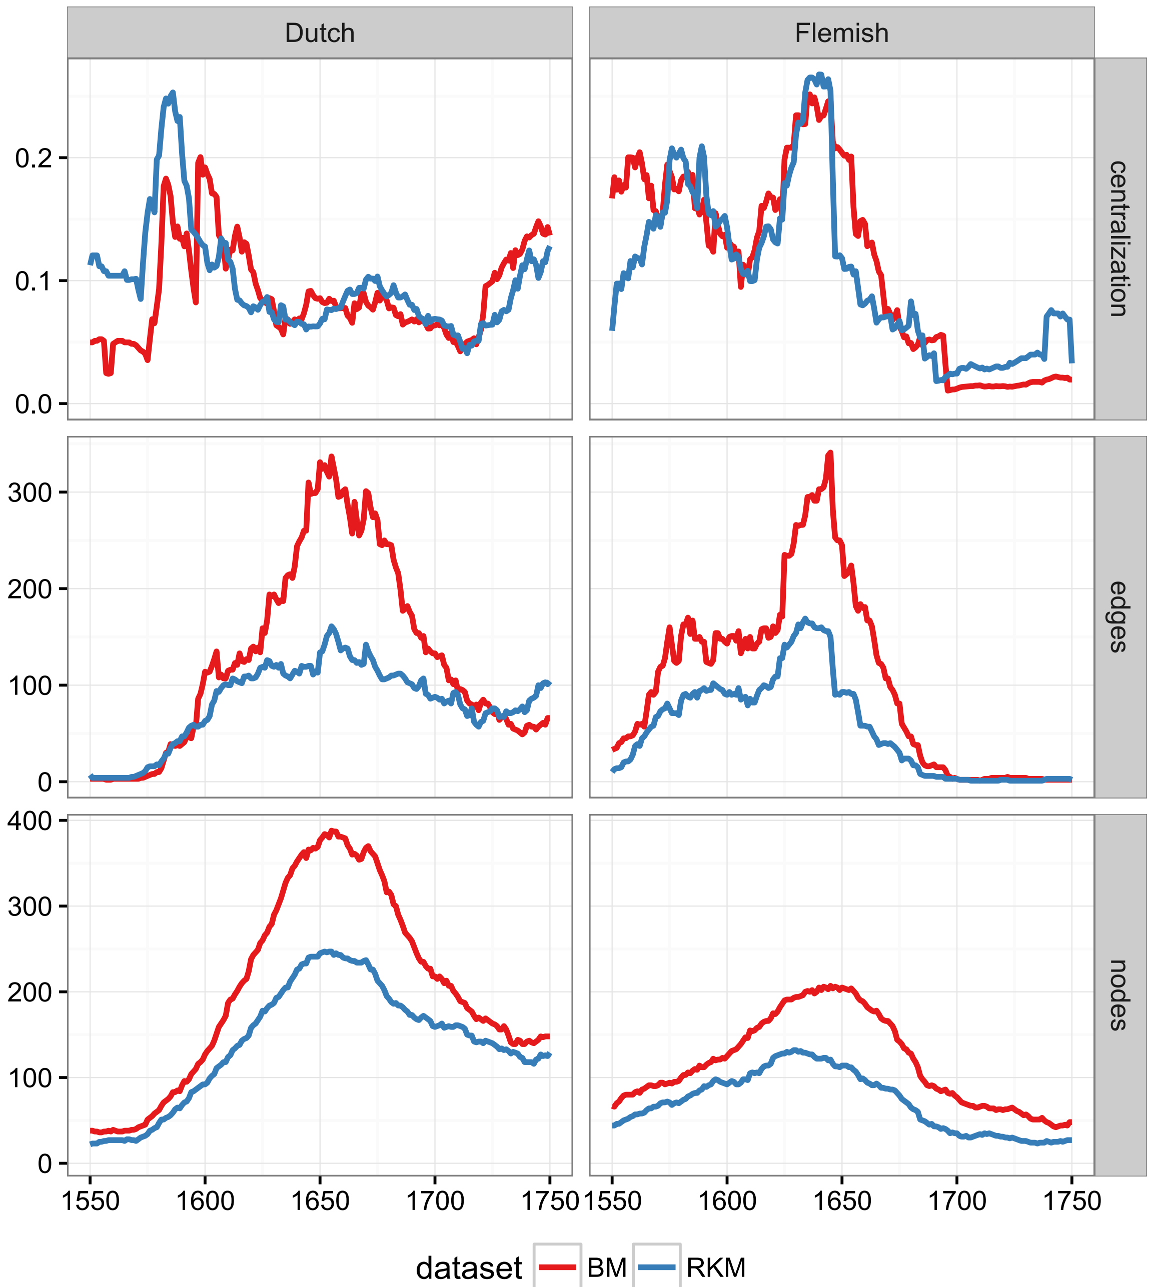

Figure 5. Comparison of the changing centralization of the Dutch and Flemish print production network between 1550–1750 with the changing number of actors and the changing number of edges connecting them over the same period.

representative of actual historical trends, and not solely an artifact of the individual collecting practices of either the BM and RKM.

A gradual rise and fall in the number of active printmakers is observed over the sixteenth and seventeenth centuries. However, the networks of both Dutch and Flemish printmakers fluctuated sharply between relative centralization and decentralization during the same period. In the northern Netherlands, centralization was generally low through most of the sixteenth-century, spiking rapidly around 1575, with another spike at 1600, before falling quickly down to earlier levels by about 1650. With small variations, the Dutch network centralization remained at relatively consistent levels from then until about 1720, when it once again quickly increased.

The southern Netherlands, by comparison, exhibited a much higher level of centralization in the mid-sixteenth century than did the Dutch. In 1550, the Flemish network also had more participants than did the Dutch. However, the population of Flemish printmakers increased only gradually through 1650, the number of active Dutch printmakers surpassing them in 1600. By 1650, the number of Flemish printmakers began to decline. Despite the smaller number of active printmakers, the print production of the southern Netherlands easily rivaled that of the north, experiencing an especially sharp increase between 1625 and 1645. Flemish centralization gradually declined until shortly after 1600, when that network experienced its own sharp spike in centralization at the same time as it saw a dramatic increase in the number of prints being made by Flemish artists. This spike was short-lived, however, and by 1650 the Flemish print production underwent relatively quick decentralization as soon as this brief printmaking boom wore off, dropping to very low levels by 1675 as it diminished greatly in size and activity.

Both the Dutch and Flemish results suggest that both the centralizing and decentralizing incentives posited above *did* have their effect in the sixteenth and seventeenth centuries, but that the incentives towards *decentralization* won out—at least in the long run. This would seem to confirm Landau and Parshall's claim that the highly-centralized model that Hieronymus Cock constructed in Antwerp was an immense success in its own time. But these initial results also undermine the notion that printmaking would inevitably necessitate highly centralized production in the following centuries. If printmaking, as a medium, encouraged high centralization at the local level, with printmaker-publishers amassing a wide range of contacts, why did this not result in persistently high centralization at the regional level? What is more, these results also suggest the surprising speed with which these structural shifts could occur. The gradual changes seen in the network population would indeed appear to

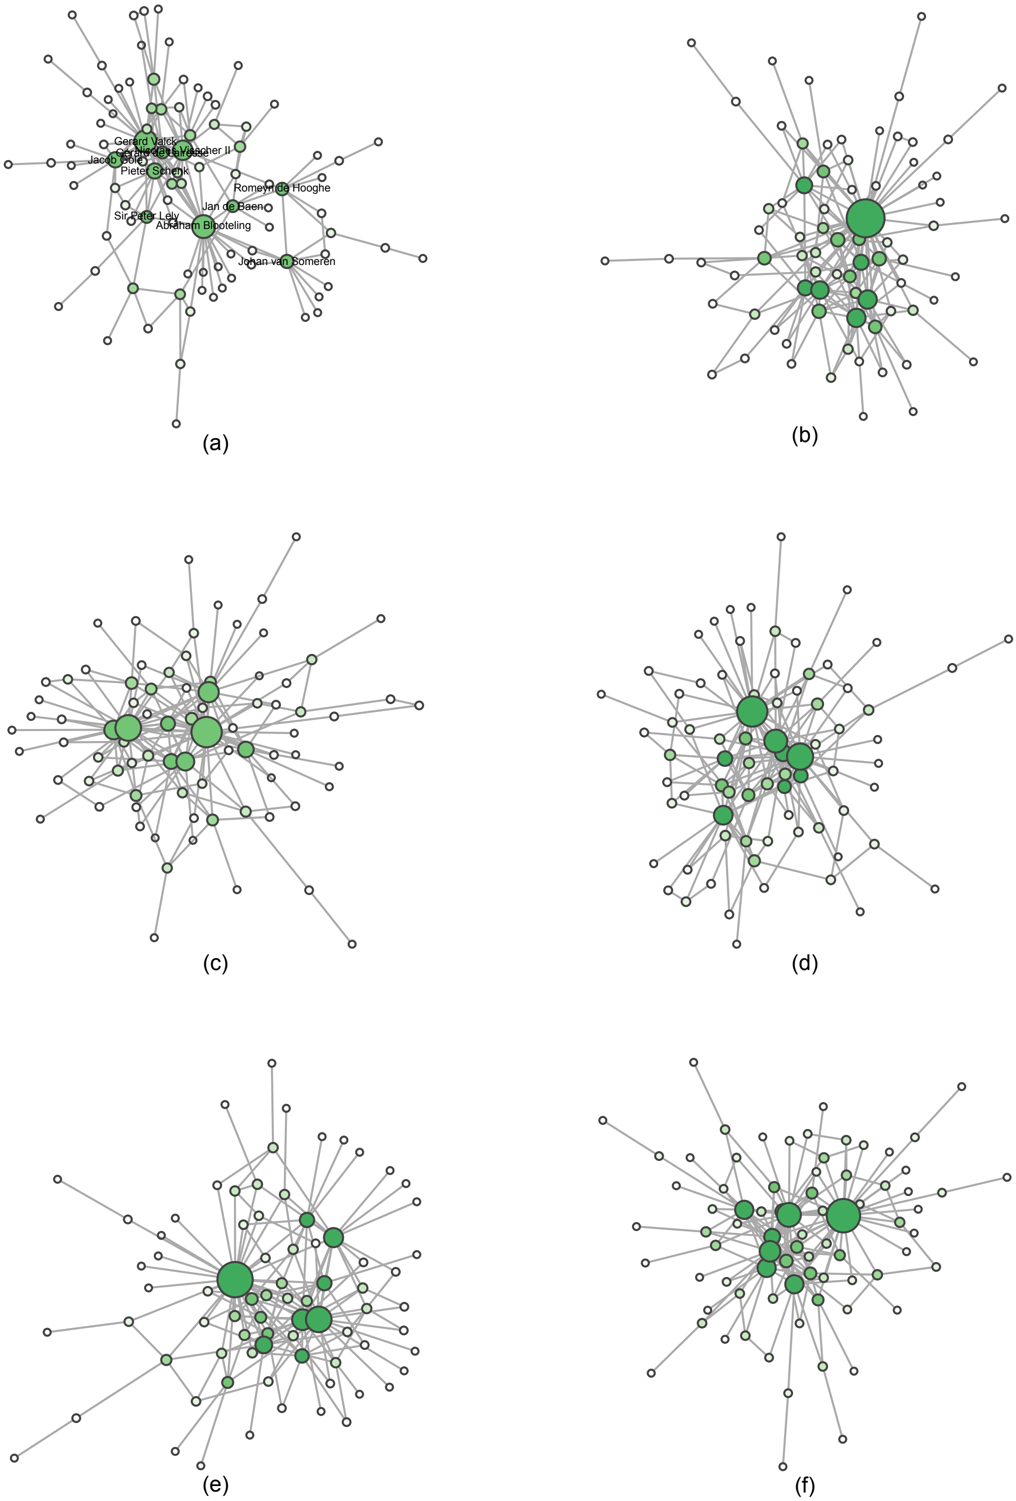

Figure 6. A visualization of random graph generation. (a) is the network generated from the BM data between 1640–1650. (b-f) are five randomly generated networks with the same number of nodes and edges as (a).

mask more dynamic upheavals in actual network concentration. The change in centralization, on the other hand, was far more abrupt.

Unaccustomed to thinking of our subject matter in terms of unpredictable systems, art historians may be surprised to find such rapid structural changes in networks involving hundreds of artists, printmakers, and publishers. This type of "phase change" behavior, where marginal changes in one set of characteristics catalyze a much more dramatic change in another measurement, are common characteristics of complex systems like social networks.[17] This raises the question: which of these changes might signify the influence of some outside event or other fundamental change in the ways in which these designers, printmakers, and publishers connected to each other? (e.g. were certain publishers or printmakers able to attract far more students or collaborators than we might expect given the size of that network?) And which changes are just the kinds of levels we might find in *any network* of the same size that follows a similar, rich-get-richer pattern of connection?[18]

We can do this by running the same centrality measurement on a random network of the same size as each network time slice (that is, with the same number of nodes and the same number of edges) derived from the BM or RKM data. Links between any two nodes in this graph are generated based on a power law probability distribution. It produces a graph such that a handful of nodes have a very large number of connections, and the majority of nodes make very few connections.[19] This distribution of edges creates a network similar to the kind that we have observed in the print production networks: a rich-get-richer scenario, where a few actors make and receive the majority of connections. Because random graph generation is stochastic, the same simulation run twice with the same inputs will produce slightly different outputs. (Figure 6 (b-f)) Run many times, a randomized simulation will tend to produce values that fall within a certain range, with many iterations producing values close to some average, and a few iterations producing outliers.

Figure 7 compares the centralization values returned by these random graph models to those found for the Dutch and Flemish communities in both the BM and RKM datasets. The shaded bands indicate the range occupied by 90% of the most central values produced by random graph sampling (thus excluding the most extreme outliers). The black trend line represents the actual centralization value measured from the data at each year. These bands indicate how centralized random networks of the same size tend to be.

Both the BM and RKM datasets are largely congruent. Moreover, we find that many of the sudden changes in Dutch and Flemish centralization through the sixteenth and seventeenth centuries are surprisingly consistent

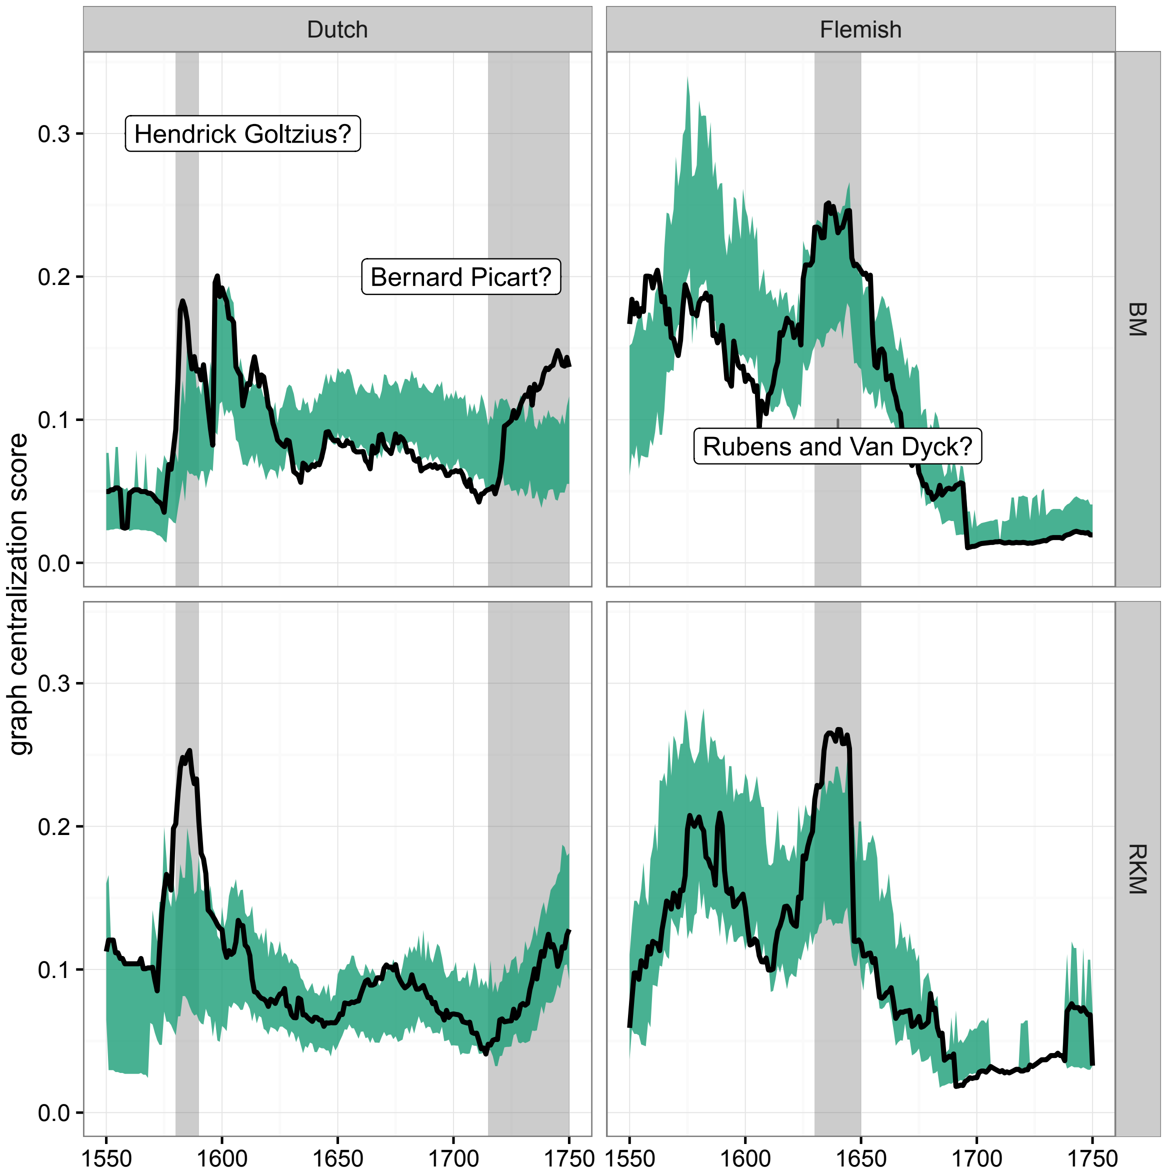

Figure 7. Comparing observed centralization results with centralization returned by random graph sampling.

with the results returned by random graph sampling. Even as observed levels of centralization shifted dramatically over this period, the basic attractiveness of well-connected individuals (the exponent of the power law probability distribution) remained essentially constant between 1550–1700. In other words, for all the apparent fluctuations in the printmaking networks in the northern and southern Netherlands, the simple incentive for printmakers to seek out well-connected collaborators appears to have been nearly constant through the end of the seventeenth century. This constant, coupled with the gradual changes in the sizes of these networks, could result in surprisingly sudden shifts in network structure.

Not every year of empirical data lines up with the random graph results, however. Not all these changes, in other words, were due to the shifting size of these communities alone. The first deviation, which appears in both the BM and RKM results, occurs in the Dutch network around 1580, with a spike in centralization that does not occur in the random graphs with the same number of nodes and edges. Another major aberration occurs in the BM results (though not the RKM model), where, around 1720, the model underpredicts the jump in centralization experienced by the Dutch network. In the Flemish network, a large spike is observed around 1640 that exceeds the centralization shown by random graph sampling.

Case studies in centrality

Having calculated the overall level of centralization in the Dutch and Flemish print production networks above, it is useful to disaggregate these larger networks and look at the changing positions of the individual artists within them. Figure 8 shows the most central members of the Dutch print production network at different points in time.[20] Many of the most central artists in each of these years are well-known names, such as Hendrick Goltzius, Claes Jansz Visscher, Hendrick Hondius, and Frederick de Wit. However, other highly central artists are relatively unknown, such as Abraham Blooteling and Jonas Suyderhoef.

Hendrick Goltzius made an enduring contribution to engraving styles and techniques, as well as to the overall artistic standing of prints, in the Netherlands and beyond in the seventeenth century—this much has never been in dispute.[21] But more than that, these results suggest that his firm also had a dramatic impact on the future structure of the Dutch printmaking network—an impact due in large part to the timing of Goltzius' career. Goltzius' well-deserved reputation as a masterful printmaker attracted apprentice engravers from many different cities in the northern and southern Netherlands. This attractive power was truly exceptional; it is the rise of Goltzius and his studio that likely explain that remarkable surge in centralization seen around 1580 that well exceeded levels seen in randomized networks of the same size. When Goltzius opened his firm in 1582, there were no Dutch competitors as centrally-placed as he was, with connections not only to his own group of Dutch printmakers, but also to foreign publishers and commissions. Any aspiring printmaker or publisher looking for a partner in the northern Netherlands would have had few choices as attractive as Goltzius' Haarlem studio at that time.

The highly centralized "star power" that had attracted so many aspiring engravers to apprentice with Goltzius in the last decades of the 1600s likely primed the network for its speedy decentralization in the following decades. The following generation of printmakers had enough built up

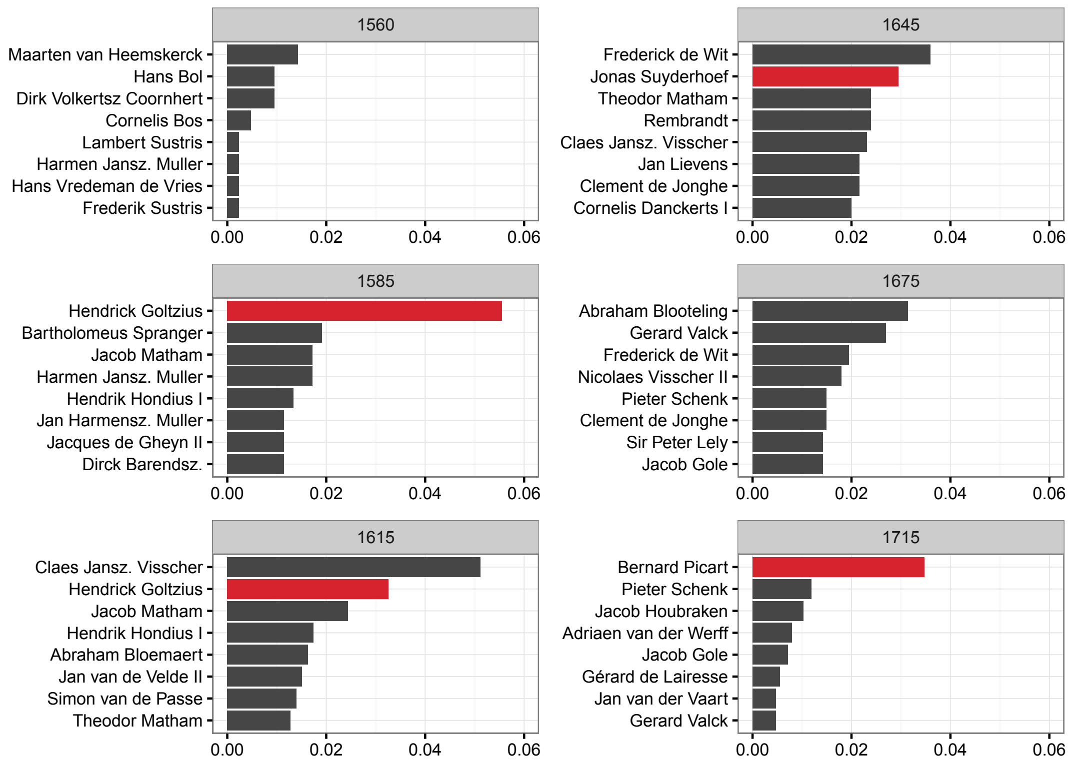

Figure 8. The top most central members of the Dutch printmaking network in the years 1560, 1585, 1615, 1645, 1675, and 1715, ranked by their degree centrality, with some notable examples highlighted. (BM Dataset)

artistic and technical knowledge to allow the print network to spread out, increasing the number of relatively competitive printmakers and thus decentralizing the overall network. For example, in 1640, Jacob Matham's son, Theodor (1605/6–1656) had risen to become the most well-connected member of the Dutch network, working in Amsterdam and specializing in engraving large figural scenes.[22] However, even as Matham, or members of the Visscher family, would assume powerful positions brokering ties between a large array of other artists, printmakers, and publishers, they would have more well-placed competition than Goltzius ever did. For the remainder of the century, no single individual would occupy as privileged a position in the Dutch network of printmakers as Goltzius did in the 1580s.

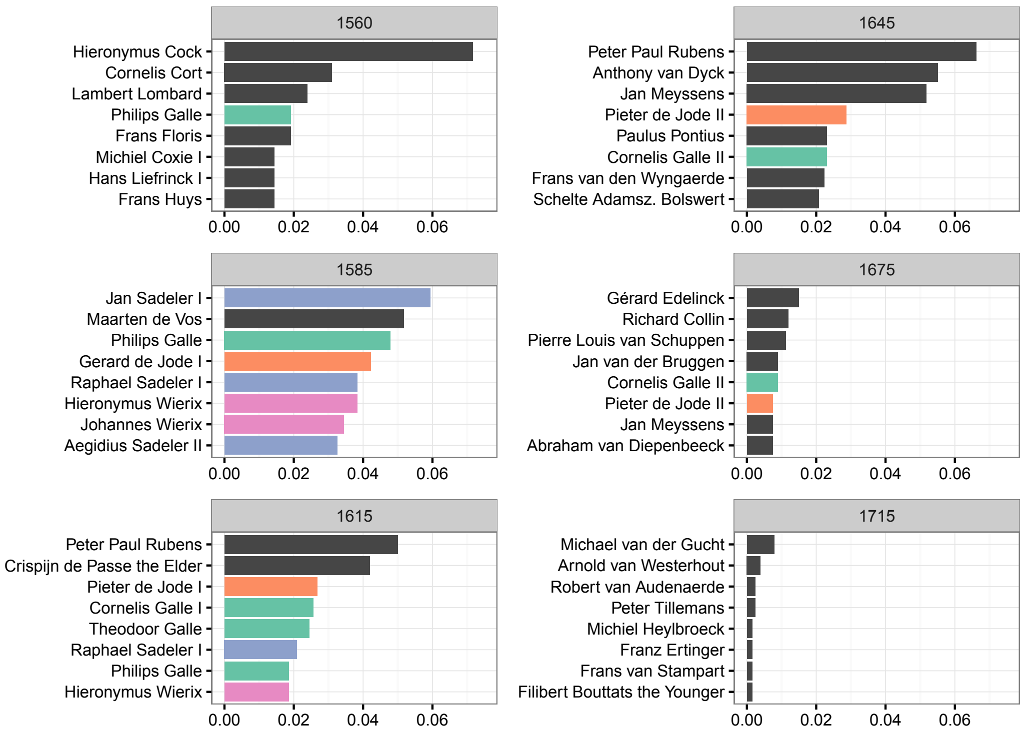

Figure 9. The top ten most central members of the Flemish printmaking network in the years 1560, 1585, 1615, 1645, 1675, and 1715, raked by their degree centrality. Members of family dynasties are highlighted. (BM dataset).

The other exceptional spike in Dutch centralization occurs in the early eighteenth century, where the French-born Bernard Picart dominated the print trade in Amsterdam soon after moving there in 1711 at the age of 38. Picart was a prolific printmaker himself, and would establish his own publishing house for producing both fine art prints as well as book illustrations. Picart trained many Dutch engravers in a formally-regulated academy, the *Amsterdamse tekenschool*, emulating his own training the French academic manner.[23] This centralized school—unique to the Dutch print production world up to that point—may account for the fact that the Dutch printmaking netowrk began to centralize far more than would be expected from its size alone in this period.[24]

It is unsurprising to find names like Goltzius or Picart at the top of the Dutch centralization charts. Far less known is the printmaker Jonas Suyderhoef, who ranks as one of the central hubs of Dutch printmaking in 1640. Jonas Suyderhoef was born in Haarlem to Andreas Pietersz Suyderhoef, secretary to the Dutch Ambassador to Constantinople.[25] He was an active member of the Guild of St. Luke, and would become dean in 1678. Suyderhoef used a combination of engraving and etching to produce prints after a remarkably wide range of artists, from formal portraits by Frans Hals, Rembrandt van Rijn, Anthony van Dyck, and Pieter Dubordieu, to history paintings by Rubens, to peasant scenes by Adriaen van Ostade, and Italianate landscapes by Jan Both. What is more, Suyderhoef worked with an impressively broad range of publishers, not only in his hometown of Haarlem, but also in Amsterdam and Leiden.

At first glance, it is easy to understand why Suyderhoef has been overlooked in literature on Dutch printmaking. Because his only known works are reproductions, he has never received the kind of close attention given to engravers from the same period who made original works as well as reproductions, such as Hendrick Hondius. Moreover, Suyderhoef's prints are often workmanlike in character. In other words, Suyderhoef had little to contribute to the aesthetic achievements of Dutch printmaking. That said, there was clearly a broad market for prints of this quality, and it is for that reason that his career deserves some attention. Suyderhoef appears to have fulfilled a demand for freelance engravers who were able to produce reasonably good engravings after almost any artist or genre, and who were willing to work with a wide array of publishers. Unlike Goltzius or Hondius, Suyderhoef does not appear to have ever tried to establish his own publishing business. Perhaps this was because there was already too much established competition among publishers in Haarlem by the 1630s when Suyderhoef would have began producing prints for the market. Conversely, the rising number of active publishers, both in Haarlem and in other Dutch cities, may have made it possible for Suyderhoef to make a comfortable living as a freelance engraver alone. In the decentralized network that existed in the mid-seventeenth century northern Netherlands, it may have made professional sense not to work exclusively with a single publisher, given the wide array of reasonably well-connected publishers that were available as clients.

Like the Dutch, the Flemish print production network underwent sudden structural fluctuations during the seventeenth century. As can be seen in Figure 5, these fluctuations did not occur in tandem. Disaggregating these network-wide statistics in order to examine individuals within this network sheds light on this difference. (Figure 9) Much as Goltzius did in the Dutch network of 1585, Hieronymus Cock unsurprisingly dominated the Flemish network of 1550. The other top-ranking participants in that year were either Cock's artistic sources (Lambert Lombard, Frans Floris, Pieter Bruegel the Elder, Michiel Coxie) or engravers with whom he contracted (Cornelis Cort, Philips Galle).[26] By 1585, the Flemish network had transitioned from a unipolar one with Cock at the center, to a multi-centric network with a large array of engraver/publishers who had worked for either Cock or Plantin early in their careers, but then transitioned into publishing works by themselves and after others. These included Jan and Rapahel Sadeler, Philips Galle (a former student of Cock's), Hieronymus and Johannes Wierix, and Gerard de Jode. Like the transition underwent by the Dutch network around 1610, the Flemish network trended towards a more decentralized structure as the overall skill level of its inhabitants increased. While one or two well-connected hub individuals might have maintained their dominance for a few decades, the small number of new printmakers that they enabled soon were influential enough to become hubs in their own right—albeit smaller in scale, and with more competition.

Yet something occurred in the south that did not happen in the North. The print production network suddenly *re*-centralized. Starting around 1625, the gradual decline in Flemish centralization reversed sharply, if briefly, to become more centralized than ever before for about two decades—a surge that was far greater than expected for a similar network of that size. (Figure 7) Sitting at the heart of this re-centered Flemish network were Rubens and Van Dyck, both of whom had a long, sustained engagement with printmakers, and whose studios generated huge demand for reproductive prints.[27] During this period of greatest network concentration, the number of prints being made surged with no accompanying increase in the number of actual participants in the network. While Flemish printmakers and publishers were producing more prints than their northern counterparts at this time, note they did so with little more than half the number of participants in the network. At the same time that the

United Provinces were experiencing an economic boom that attracted many artists, the southern Netherlands' economy was still suffering from the effects of Spanish invasion at the end of the sixteenth century.[28] The closure of the Scheldt river and the exodus of prosperous Protestant merchants had led to economic stagnation in Antwerp. Though shining stars like Rubens and Van Dyck, not to mention engravers like Aegidus Sadeler, could still attract international commissions, Antwerp itself was not an attractor for young artistic talent, by in large. In a small network with few new entrants, a handful of influential individuals could still maintain the kinds of highly centralized positions that were no longer available in the more distributed northern provinces. This focused and demanding environment would have favored highly professionalized printmakers who were able to produce prints consistently and prolifically. This may explain the predominance of printmaking *families* among the list of central Flemish print producers like the Weirixes, the De Jodes the Galles, or the Sadelers. They all established businesses that were able to produce massive, multi-generational stocks of plates. With few newcomers in the seventeenth century tempted to try their luck at starting an engraving or publishing practice in Antwerp, already established firms were easily sustained over multiple generations.

More so than any of their counterparts in the northern Netherlands, these Flemish printmaking families were able to establish long-lived dynasties that dominated Antwerp print publishing for more than a century. Without the persistent influx of new talent and competition into Antwerp, these established Flemish families could easily maintain their businesses by passing down artistic knowledge, social connections, and (perhaps most importantly), large inventories of plates that provided a stable foundation for young heirs to the business.

Conclusion: Simple rules for a complex system

In one sense, this analysis has confirmed Parshall and Landau's core argument: that the highly centralized form of print production originated by Hieronymus Cock, Antonio Lafréri, and other sixteenth-century publishers would govern printmaking well into the future. This system of professionalized publication would tend to favor well-connected individuals able to marshall commissions, clients, and labor from a wide spectrum of the market. However, network analysis has also revealed that the ramifications of this simple rule could be unexpectedly complex. With a handful of exceptions where combinations of economic circumstances and extraordinary individuals like Hendrick Goltzius, Bernard Picart, or Peter Paul Rubens caused unique spikes in centralization, many of the apparently major changes in Dutch and Flemish printmaking centralization were simply manifestations of the gradual rise and fall in the number of active printmakers over this period. A network-based approach to studying print production has also highlighted otherwise neglected artists like Jonas Suyderhoef, who may have played a far more important role in disseminating images through reproduction than has previously been acknowledged. This approach offers a crucial context for future studies of individual printmakers, while also demonstrating how network analysis can illuminate dimensions and scales of historical events that are otherwise difficult for art historians to conceptualize.

Supplementary Material

To this article supplementary material can be found at HeiDATA Dataverse Network https://heidata.uni-heidelberg.de/dvn/dv/dahjournal

Data, analysis, and visualization code (R-package)
Supplementary methodology details

Notes

1 David Landau and Peter W. Parshall, *The Renaissance Print, 1470-1550* (New Haven: Yale University Press, 1994), 260–283.
2 Riggs and Silver generally find support for continuing professionalization and centralization; Timothy Riggs and Larry Silver, *Graven Images: The Rise of Professional Printmakers in Antwerp and Haarlem, 1540-1640* (Chicago: Mary & Leigh Block Gallery, Northwestern University, 1993); however Orenstein has noted an increasing number of smaller print publishing houses flourishing in the Netherlands around mid-century: Nadine M. Orenstein, "The Shift from Antwerp: The Diversification of Print Publishers in the United Provinces Around 1600," *Block Points: The Annual Journal and Report of the Mary and Leigh Block Gallery* 5 (1995): 43–63; Nadine M. Orenstein, "Marketing Prints to the Dutch Republic: Novelty and the Print Publisher," *Journal of Medieval & Early Modern Studies* 28, no. 1 (1998): 141.
3 See for example, the strategies of Christopher Plantin; Karen L. Bowen, *Christopher Plantin and Engraved Book Illustrations in Sixteenth-Century Europe* (Cambridge: Cambridge University Press, 2008), 22–25, 54–55.
4 John Michael Montias, *Artists and Artisans in Delft: A Socio-Economic Study of the Seventeenth Century* (Princeton: Princeton University Press, 1982), ch. 7.
5 Even in the sixteenth century these products were usually not preserved; Jan van der Stock, *Printing Images in Antwerp: The Introduction of Printmaking in a City: Fifteenth Century to 1585*, Studies in Prints and Printmaking 2 (Rotterdam: Sound & Vision Interactive, 1998), 179–180.
6 For an effective demonstration of this technique, see Jan De Vries, "Art History," in *Art in History, History in Art: Studies in Seventeenth-Century Dutch Culture*, ed. David Freedberg and Jan De Vries (Santa Monica: Getty Center for the History of Art & the Humanities, 1991), 259–260.
7 Future work may involve applying statistical models of survivability to thinking about the loss of prints, possibly based on surviving archival evidence about the sizes of original print runs; see, for example, Karen L. Bowen and Dirk Imhof, "18,257 Impressions from a Plate," *Print Quarterly* 22, no. 3 (September 2005): 265–79.
8 Peter W. Parshall, "Prints as Objects of Consumption in Early Modern Europe," *Journal of Medieval & Early Modern Studies* 28, no. 1 (Winter 1998): 21.

9 On the BM Linked Open Data initiative, see Dominic Oldman et al., "Realizing Lessons of the Last 20 Years: A Manifesto for Data Provisioning and Aggregation Services for the Digital Humanities (A Position Paper)," *D-Lib Magazine* 20, no. 7/8 (July 2014), doi:10.1045/july2014-oldman. What documentation is available for the RKM data may be found at http://rijksmuseum.github.io/.
10 For the handling of unknown creators, see supplementary material.
11 These three roles encompass 97% of the artistic ties as described by the BM for the set of prints under consideration, and 77% of all ties described by the RKM. A more detailed discussion of the various types of production roles can be found in the supplementary material.
12 The interval of ten years is somewhat arbitrary. I have selected it because it strikes a good balance between granularity and generality, and also falls within the rough art historical convention of using "circa" to mean five years. See the supplementary material for a discussion of including only alive vs. deceased artists in the network.
13 Both the BM and RKM databases specify a start_date and end_date for artworks. Most artworks have only one date of production associated with them; in these cases, start_date and end_date are equal.
14 On directed versus undirected networks and network weighting, see supplementary material.
15 Freeman defines graph centrality by summing the difference of each individual nodes' degree centrality to that of the most-central node , normalized based on the size of the network: Linton C. Freeman, "Centrality in Social Networks: Conceptual Clarification," *Social Networks* 1, no. 3 (1978): 226, doi:10.1016/0378-8733(78)90021-7–237.
16 There are several other more subtle methods for measuring centralization, such as betweenness, closeness, or eigenvector centrality. The simple measure of degree centrality is useful for this particular question because we are interested in looking at the basic numbers of collaborators that these network participants had, rather than more complex 2nd-degree or nth-degree relationships whose meaning in this co-production network is more ambiguous.
17 Duncan J. Watts, *Small Worlds: The Dynamics of Networks Between Order and Randomness* (Princeton: Princeton Univ. Press, 1999), 53.
18 Such patterns are often found in cultural networks, as seen in Maximilian Schich et al., "A Network Framework of Cultural History," *Science* 345, no. 6196 (January 8, 2014): 558–62, doi:10.1126/science.1240064.
19 The model implemented here sets the connection chance, or fitness , of node such that: The exponent determines the skew of the probability distribution. In this context, the skew governs precisely how attractive well-connected individuals are to new entrants to the network, with a larger skew denoting a stronger attraction. A of provides a close fit for almost every network shown here. On this model for generating random networks, see Albert-László Barabási and Réka Albert, "Emergence of Scaling in Random Networks," *Science* 286, no. 5439 (October 15, 1999): 511, doi:10.1126/science.286.5439.509, and the R implementation in G. Csardi and T. Nepusz, "The Igraph Software Package for Complex Network Research," *InterJournal* Complex Systems (2006): 1695, http://igraph.org.
20 This sampling of years is somewhat arbitrary, but does provide a relatively even overview of the time period under consideration.
21 On the Goltzius studio, see Jan Piet Filedt Kok, "Hendrick Goltzius: Engraver, Designer, and Publisher 1582-1600," *Nederlands Kunsthistorisch Jaarboek* 42-43 (1991): 159–218; Huigen Leeflang, *Hendrick Goltzius 1558-1617: Drawings, Prints, and Paintings* (Amsterdam: Rijksmuseum, 2003), ch. 4.
22 Léna Widerkehr, "*Jacob Matham Goltzij Privignus*: Jacob Matham Graveur et Ses Rapports Avec Hendrick Goltzius." *Nederlands Kunsthistorisch Jaarboek* 42-43 (1991–1992): 219–60.
23 Nelke Barthelings, "Bernard Picart, a French Engraver in the Dutch Republic," in *Les échanges Artistiques Entre Les Anciens Pays-Bas et La France, 1482-1814*, ed. Gaëtane Maës and Jan Blanc (Turnhout, Belgium: Brepols, 2010), 49.

24 Future work is require to confirm this suggestion, particularly given that the BM and RKM databases return conflicting results for the early eighteenth century.
25 Johann Wussin, *Jonas Suyderhoef, son œuvre gravé, classé et décrit*, trans. Henri Hymans (Brussels: Labroue et Mertens, 1862), 11–12; F. W. H. Hollstein, *Dutch and Flemish Etchings, Engravings, and Woodcuts, ca. 1450-1700* (Amsterdam: M. Hertzberger, 1949–2001), 28:201–260.
26 Joris van Grieken, Ger Luijten, and Jan van der Stock, *Hieronymus Cock: The Renaissance in Print* (New Haven: Yale University Press, 2013), 30–36.
27 Nico van Hout and Paul Huvenne, *Copyright Rubens: Rubens en de grafiek* (Gand: Ludion, 2004); Ann Diels, *The Shadow of Rubens: Print Publishing in 17th-Century Antwerp: Prints by the History Painters Abraham Van Diepenbeeck, Cornelis Schut and Erasmus Quellinus II*, trans. Michael Hoyle and Irene Schaudies (Turnhout: Royal Library of Belgium, 2009).
28 On the downturn‘s effects on the Flemish art market, see Claartje Rasterhoff and Filip Vermeylen, "The Zeeland Connection: The Art Trade Between the Northern and Southern Netherlands During the Seventeenth Century," in *Moving Pictures: Intra-European Trade in Images, 16th-18th Centuries*, ed. Neil De Marchi and Sophie Raux (Turnhout: Brepols, 2014), 123–50.

Bibliography

Barabási, Albert-László, and Réka Albert. "Emergence of Scaling in Random Networks." *Science* 286, no. 5439 (October 15, 1999): 509–12. doi:10.1126/science.286.5439.509.

Barthelings, Nelke. "Bernard Picart, a French Engraver in the Dutch Republic." In *Les échanges Artistiques Entre Les Anciens Pays-Bas et La France, 1482-1814*, edited by Gaëtane Maës and Jan Blanc, 33–54. Turnhout, Belgium: Brepols, 2010.

Bowen, Karen L. *Christopher Plantin and Engraved Book Illustrations in Sixteenth-Century Europe*. Cambridge: Cambridge University Press, 2008.

Bowen, Karen L., and Dirk Imhof. "18,257 Impressions from a Plate." *Print Quarterly* 22, no. 3 (September 2005): 265–79.

Csardi, G., and T. Nepusz. "The Igraph Software Package for Complex Network Research." *InterJournal* Complex Systems (2006): 1695. http://igraph.org.

De Vries, Jan. "Art History." In *Art in History, History in Art: Studies in Seventeenth-Century Dutch Culture*, edited by David Freedberg and Jan De Vries, 249–71. Santa Monica: Getty Center for the History of Art & the Humanities, 1991.

Diels, Ann. *The Shadow of Rubens: Print Publishing in 17th-Century Antwerp: Prints by the History Painters Abraham Van Diepenbeeck, Cornelis Schut and Erasmus Quellinus II*. Translated by Michael Hoyle and Irene Schaudies. Turnhout: Royal Library of Belgium, 2009.

Freeman, Linton C. "Centrality in Social Networks: Conceptual Clarification." *Social Networks* 1, no. 3 (1978): 215–39. doi:10.1016/0378-8733(78)90021-7.

Hollstein, F. W. H. *Dutch and Flemish Etchings, Engravings, and Woodcuts, ca. 1450-1700*. Amsterdam: M. Hertzberger, 1949–2001.

Hout, Nico van, and Paul Huvenne. *Copyright Rubens: Rubens en de grafiek*. Gand: Ludion, 2004.

Kok, Jan Piet Filedt. "Hendrick Goltzius: Engraver, Designer, and Publisher 1582-1600." *Nederlands Kunsthistorisch Jaarboek* 42-43 (1991): 159–218.

Landau, David, and Peter W. Parshall. *The Renaissance Print, 1470-1550*. New Haven: Yale University Press, 1994.

Leeflang, Huigen. *Hendrick Goltzius 1558-1617: Drawings, Prints, and Paintings*. Amsterdam: Rijksmuseum, 2003.

Montias, John Michael. *Artists and Artisans in Delft: A Socio-Economic Study of the Seventeenth Century*. Princeton: Princeton University Press, 1982.

Oldman, Dominic, Martin de Doerr, Gerald de Jong, Barry Norton, and Thomas Wikman. "Realizing Lessons of the Last 20 Years: A Manifesto for Data Provisioning and Aggregation Services for the Digital Humanities (A Position Paper)." *D-Lib Magazine* 20, no. 7/8 (July 2014). doi:10.1045/july2014-oldman.

Orenstein, Nadine M. "Marketing Prints to the Dutch Republic: Novelty and the Print Publisher." *Journal of Medieval & Early Modern Studies* 28, no. 1 (1998): 141.

———. "The Shift from Antwerp: The Diversification of Print Publishers in the United Provinces Around 1600." *Block Points: The Annual Journal and Report of the Mary and Leigh Block Gallery* 5 (1995): 43–63.

Parshall, Peter W. "Prints as Objects of Consumption in Early Modern Europe." *Journal of Medieval & Early Modern Studies* 28, no. 1 (Winter 1998): 19.

Rasterhoff, Claartje, and Filip Vermeylen. "The Zeeland Connection: The Art Trade Between the Northern and Southern Netherlands During the Seventeenth Century." In *Moving Pictures: Intra-European Trade in Images, 16th-18th Centuries*, edited by Neil De Marchi and Sophie Raux, 123–50. Turnhout: Brepols, 2014.

Riggs, Timothy, and Larry Silver. *Graven Images: The Rise of Professional Printmakers in Antwerp and Haarlem, 1540-1640*. Chicago: Mary & Leigh Block Gallery, Northwestern University, 1993.

Schich, Maximilian, Chaoming Song, Yong-Yeol Ahn, Alexander Mirsky, Mauro Martino, Albert-László Barabási, and Dirk Helbing. "A Network Framework of Cultural History." *Science* 345, no. 6196 (January 8, 2014): 558–62. doi:10.1126/science.1240064.

van der Stock, Jan. *Printing Images in Antwerp: The Introduction of Printmaking in a City: Fifteenth Century to 1585*. Studies in Prints and Printmaking 2. Rotterdam:

Matthew Lincoln is a Data Research Specialist at the Getty Research Institute, where he focuses on data-driven research into the history of the art market. He earned his PhD in Art History at the University of Maryland, College Park in 2016, and has been a recipient of Kress and Getty Foundation grants for their summer institutes in digital art history.

Correspondence e-mail: mlincoln@getty.edu

Sound & Vision Interactive, 1998.

van Grieken, Joris, Ger Luijten, and Jan van der Stock. *Hieronymus Cock: The Renaissance in Print*. New Haven: Yale University Press, 2013.

Watts, Duncan J. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton: Princeton University Press, 1999.

Widerkehr, Léna. "*Jacob Matham Goltzij Privignus*: Jacob Matham Graveur et Ses Rapports Avec Hendrick Goltzius." *Nederlands Kunsthistorisch Jaarboek* 42-43 (1991–1992): 219–60.

Wussin, Johann. *Jonas Suyderhoef, son œuvre gravé, classé et décrit*. Translated by Henri Hymans. Brussels: Labroue et Mertens, 1862.

Interview

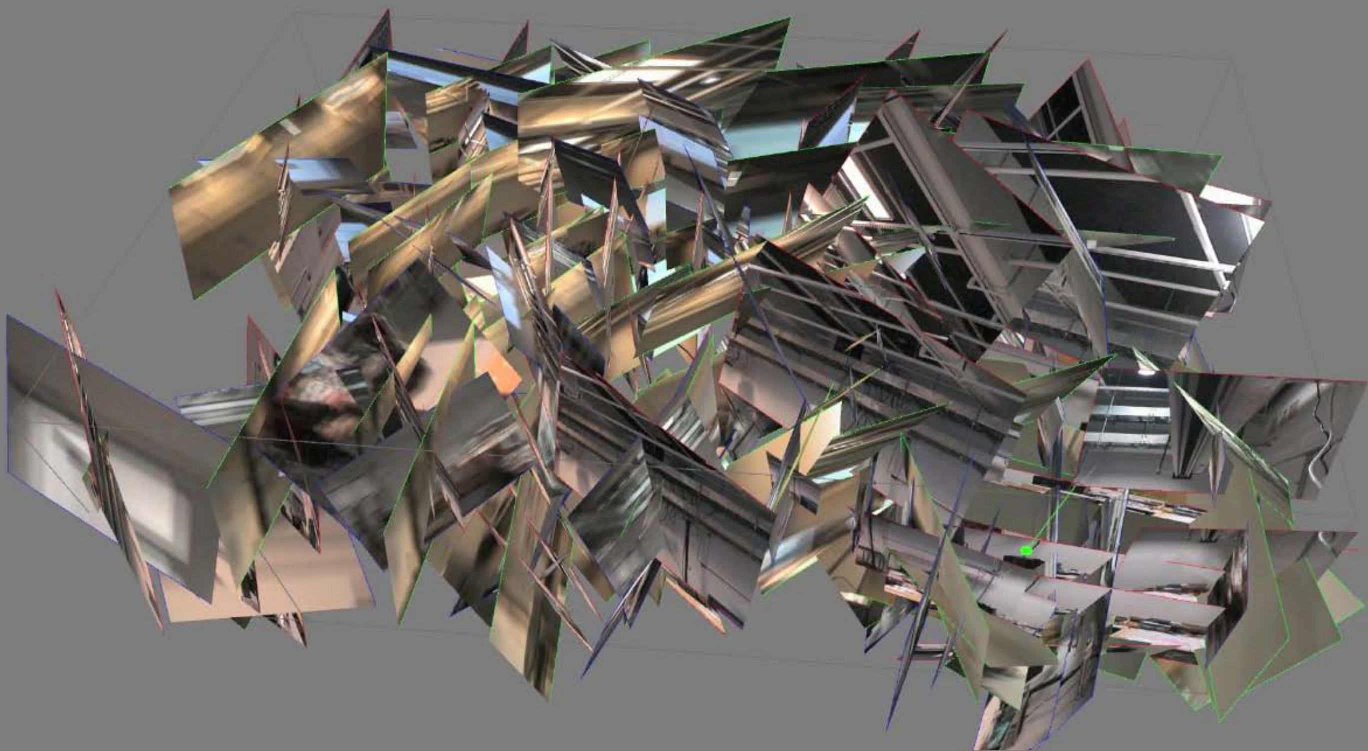

Video still from “AutoVision”, 2014, George Legrady, Danny Bazo, https://vimeo.com/111252770 (Date accessed: March 19, 2016).

In Conversation with George Legrady:

Experimenting with Meta Images. Artistic Approaches meet Computational Methods

Harald Klinke, Liska Surkemper

The computer is a valuable tool to get an overview of large datasets by creating visualizations as meta images. However, such visualizations that may seem so self-evident are never a one-to-one translation of the underlying data. Just like photography is no "pencil of nature" visualizations are no "pencil of data", but a transformation with many variables determining the result.

George Legrady is one of the pioneers who examine artistically the visual outcome of algorithms by creating new forms of visualizations and 3D installations. As the director of the Experimental Visualization Lab and Professor of Interactive Media at the University of California, Santa Barbara, he also introduces students of art, design, and computational engineering to concepts at the intersection between art and technology.

Liska Surkemper (LS) – George, you have been integrating computational methods in your artistic work since the mid-1980's. One may refer to you as an early adopter to the then new technology – what seems odd, considering you come from analogue photography. What was your initial motive as an artist to work with the computer so early?

George Legrady (GL) – The transition from working in the photographic medium, an optical-mechanical device, to expressing visualizations with a language-based processing machine seemed to me an inspirational progression, not as disconnected as many in the arts community at the time felt. I was already exposed to a digital instrument in childhood, the piano, and music training did involve the inscription and description of mathematically defined code – the music score – as a means of registering and replaying music.

I was actually surprised back in the early to mid-1980s that there were fewer visual artists captivated by the digital. Of course, a few artists from an earlier generation that saw parallels to conceptual art engaged with the medium. There was the "Software" exhibition in 1970 at the Jewish Museum NYC with Nicholas Negroponte and the art theorist Jack Burnham, and artists such as Stephen Willats in London, and a number of others. It may be that what held back the exploration was the lack of digital imaging technologies to produce images of a reasonable quality. I was perplexed that advances in imaging technologies took so long to enter the marketplace given the ubiquity of photographic representation in the culture-at-large. Electronic and algorithmic music composition, and audio technologies had already evolved in the 1970s, so imaging lagged significantly behind. I had to wait from 1981, when I was introduced to computer programming in the studio of the abstract painter Harold Cohen at UCSD, until 1986 when the AT & T Targa Truevision Image Capture Board came on the market to have the opportunity to explore ideas about photographic imaging in the digital realm.

Harald Klinke (HK) – In the 1980's you examined the conventions of the photographic image. How do those experiences have affected your work since then?

GL – It has greatly affected how I have approached computer-generated image construction. Keep in mind that I began working with computers at a time when software was scarce and therefore I had to learn how to write my own. Into that code production, I brought those questions that in the 1980s were directed at the veracity of the photographic image – the image being transparent onto the world but meanwhile hiding the ideological hand that has guided its framing. I have addressed the issues related to this topic in an article published in the late 1980s titled "Image, Language and Belief in Synthesis".[1] The article discusses issues of belief in the digital image, and the necessity of constructing all details of an experience, simulating nature by including random deviations.

HK – How much does the camera determine the look onto our world and what we think reality is?

I was perplexed that imaging technologies took so long to enter the market place.

GL – The optical component of the camera has not changed much in the last 6 centuries but much research is currently going on to mathematically simulate and go beyond the capabilities of conventional camera image capture systems. Depth perception is one area, and then creating something like the "Esper Photo Analysis machine" as seen in the Sci-Fi movie Blade Runner is another. In the film, the protagonist played by Harrison Ford, requests the machine to navigate inside a found photograph and redirect one's perspective view beyond the flattened depth of the standard 2D photographic image.

LS – Comparing what was possible thirty years ago and what is today: What was the most memorable or surprising technological development for you that has changed your approach or working method?

GL – Besides the amazing power of the "undo" button, the most significant transformation in the culture has been the exponential growth of communities of practitioners, media arts designers, architects, hacker cultures, etc. and the rapid transition to digitality in the culture-at-large. Prior to the introduction of the Mosaic web browser which opened the Pandora box of universal communication around 1993, there were few digital media practitioners, usually eccentric, intensively focused individuals who were more like inventors rather than users, as hardware and software had to be invented. 1992 saw the introduction of Quicktime, which allowed for time-based image and sound to be digitized. The Macromedia Director software introduced multi-linear branching, and interactivity into the sequencing of multimedia scenes.

LS – Talking about interactive media: what are current projects you are working on?

GL – I am working on a number of different projects in diverse fields such as computational research, media arts, data visualization, artworks for the gallery market, and reformulating early documentary photographic projects into digital archives as some now have historical relevance. The research component focuses on translating a photographer's decisions of framing and image composition into rules that can be computationally programmed to study what results a machine may deliver. "Swarm Vision", "AutoVision", "Exquisite Vision" have all been developed in my university lab working with my Ph.D. students to explore if the aesthetic decisions involved in photographic image composition could possibly be transferred to a computer. This work has been funded by a National Science Foundation Intelligence and Information Systems grant and the intent is to achieve both engineering and artistic results. We are very much at the beginning stage of teaching a computer how to take interesting photographs. The purpose of the effort is really to explore the question of to what degree can an artistic approach be inscribed into computational language. The media arts projects focus on creating interactive large-scale projection installations based on the engineering research just described. Many of these projects collect data while they are on exhibit. The analysis of that data is transformed into data visualization, as for instance, the public contributions in the "Pockets Full of Memories" exhibition and the data from the Seattle Public Library commission which has now gathered over 85 million datasets since its activation in September 2005.

LS – And what about your gallery work?

GL – The fine arts gallery work I am currently doing explores visual and cultural narratives using the lenticular imaging process, which I create at a relatively large scale. I am particularly interested in the cultural encoding that occurs with photographs over time. We see a photograph from today and we consider it for its information, whereas a photograph from thirty years ago has embedded within it a set of messages that it is from another era. There are semiotic, syntactic, and cultural content information that tells us so, and I am intrigued by the significance of why we consider the time displacement to be of such a critical concern when we view such images. Photographs taken at a particular moment in time may also have cultural and ethnographic relevance, allowing us to compare the then and the now.

I am particularly interested in the cultural encoding that occurs with photographs over time.

I have been digitizing and creating an online database of documentary photographs taken in the Canadian north in Cree Nation villages who have been negotiating over land rights since the early 1970s. The process of digitization and posting online has been recognized by research agencies as a form of cultural repatriation, and a National Science Foundation Arctic Social Science grant has made it possible for me to return to the Cree villages, to present the photographs I have taken some 45 years ago. This project has been done in collaboration with McGill University ethnographers, who have relations with the Cree cultural centers in the villages.

LS – Looking at the different kinds of projects you do: are you more interested in the aesthetic of the work, the examination of the digital method or is there another epistemic impact you want to foster? Or is it a combination of those things?

GL – I am interested in the combination of the two – the aesthetic and also what you call the epistemic as this may imply a few things. As you point out, there is the examination of methodologies of how one works with digital construction of information/

knowledge, and there is also the larger question of how digital processes are reformulating how we see the world. Our understanding of things change over time, for instance nature in the 19th century was understood as a dangerous unknown. In the early 20th century, nature became a resource, and now we sample, digitize, reconstruct, simulate, and reinvent nature. Jean Baudrillard has raised the question of to what degree do we today truly experience nature. The digital is a way of understanding the world, it is a specific kind of filter, allowing us greater control as we can numerically process the analog world, but it also reformulates our understanding of the world, and creates distance. To use Vilém Flusser's discussion about the technical image, we "create, process and store symbols." If we look through the camera, it is "to pursue new possibilities" to produce information.[2]

HK – Since you mention the term image: what is the difference for you between visualization and (artistic) image?

GL – I would say that visualization is an expression of information (data) where the visual result is an outcome of how the data has been processed and has influenced the shape and form of the visualization.

The artistic image is somehow the reverse: an image is created through a process or through a technique that is an outcome of a combination of intended expression, procedure and chance events. For instance, how a pigment may be applied, or how an unintended movement or disruption may impact on the image composition at the moment the camera's shutter is released.

The process of digitization and posting online has been recognized by research agencies as a form of cultural repatriation.

HK – Your work also includes experimenting with algorithmic generated visualizations. What kind of challenges do you encounter by operating with Big Image Data?

GL – In contrast to the more precise filters I imagine the Social Sciences proceed (I may be wrong), the projects I do with large data, for instance, "Pockets Full of Memories" and the Seattle Public Library, have a broader range of data expression and are therefore noisy. There are a lot of outliers, data outside of the expected range. This in itself is of course an interesting topic to explore and some of my students who have worked on the Seattle project have studied how the classification system in place will


PLEASE, MOVE THE MOUSE OVER THE MAP
AND CLICK ON AN OBJECT TO SEE ITS DATA.
INSTRUCTIONS
Contribute an object in your
the archive. Simply place it in the
station in the middle of the gallery and
the instructions using the touch screen.
Use this computer to view
www.pocketsfullofmemories.com where you
can read comments about other visitor's
objects and add a story to your own entry.

Conceived as an installation on the topic of the archive and memory, “**Pockets Full of Memories**” was exhibited on the main floor of the Centre Pompidou from April 10 to September 3, 2001. During this time, 20000 visitors came to view the installation and contributed over 3000 objects in their possession, digitally scanning and describing them. This information was stored in a database and organized by an algorithm that positioned objects of similar value near each other in a two-dimensional map. The map of objects was projected in the gallery space and also accessible online at www.pocketsfullofmemories.com where individuals in the gallery and at home could review the objects and add comments and stories to any of the them.

The archive of objects consists of objects that museum visitors carried with them, for instance, such common items as phones, keys, toys, clothing, personal documents, currency, reading material, and others. The size of the scanning box was the only limiting factor that determined what could be added to the archive. Surprisingly, the database includes an unusual number of scanned heads, hands and feet, extending the archive from simply being a collection of objects to encoding it with the corporeal presence of the contributors.

The 2D map on the projected screen in the gallery consists of 384 objects selected from the total database by the Kohonen self-organizing map algorithm. The ordering of the objects are based on the ways that the audience described them through the touchscreen questionnaire. The map of objects continuously organized itself until the end of the exhibition and the order of the final map is a consequence of all the contributions from the duration of the exhibition. This phenomenon is called emergence as the order is not determined beforehand but emerges through the large number of local interactions on the map. This is why the system can be called ‘self-organizing‘. Accessibility on the internet has provided a means by which to extend the dialogue for visitors, as the internet audience has the opportunity to add comments and stories to any object, and from anywhere in the world. Many visitors who have traveled from other geographical areas have used this as a means to make contact with friends and family back home who then have added their own responses.

Produced in collaboration with Dr. Timo Honkela, Media Lab, University of Art and Design Helsinki, (Kohonen self-organizing neural-net algorithm); C3 Center for Culture and Communication, Budapest (touchscreen data collection, hardware and software); Projekttriangle, Stuttgart, (design and visual identity); Dr. Brigitte Steinheider, Fraunhofer Institute of Research, Stuttgart / University of Oklahoma, Tulsa (questionniare and data analysis); Andreas Schlegel, (visualization programming); CREATE lab, UC Santa Barbara, (web software development). With the financial assistance of The Daniel Langlois Foundation for Art, Science, and Technology, Montreal, Canada, the Centre Georges Pompidou, and the Office of Research, UC Santa Barbara. The collected data can be viewed today at http://tango.mat.ucsb.edu/pfom/databrowser.php

Source: www.georgelegrady.com

Above: Detail of the Wall of Images, “Pockets Full Of Memories”, 2001, Comissioned by the Centre Pompidou, Paris, George Legrady.

Below: Installation view of “Pockets Full of Memories” (extended), 2005, at Cornerhouse Gallery, Manchester, George Legrady.

We are very much at the beginning stage of teaching a computer how to take interesting photographs.

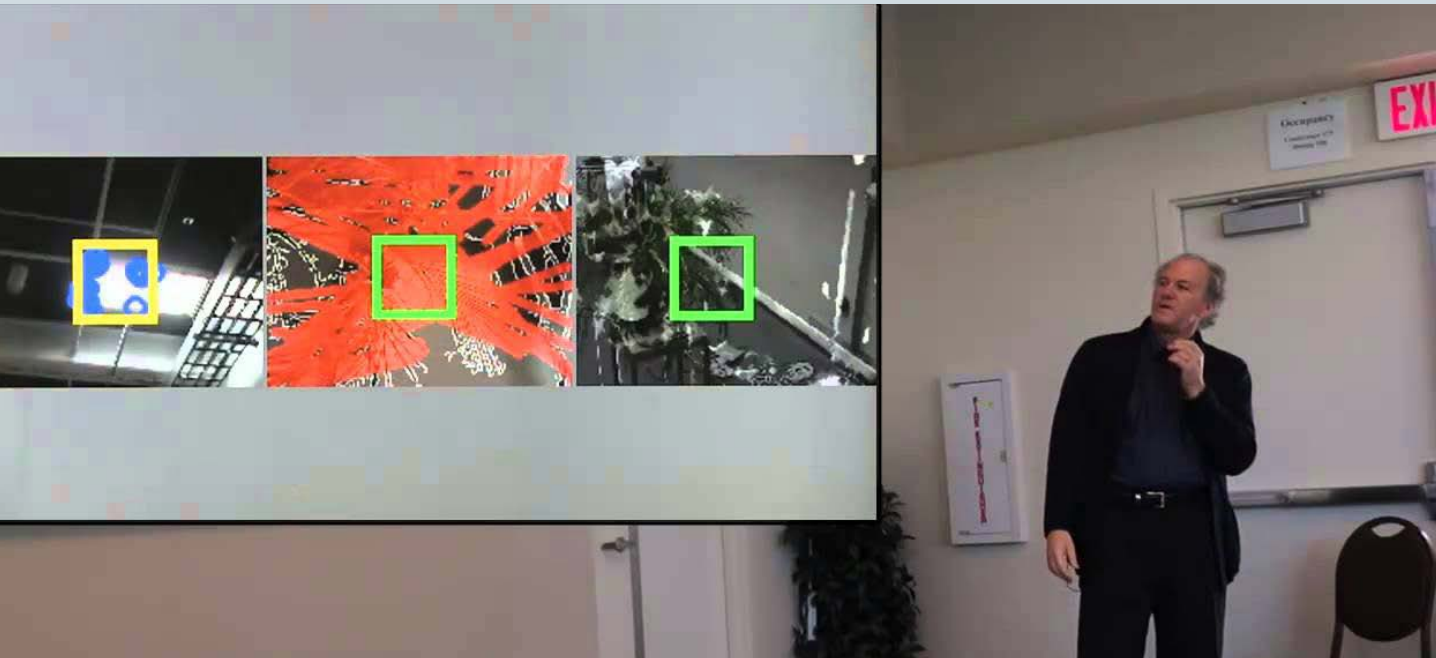

George Legrady on “Swarm Vision“ at the spatial@ucsb Lightning Talk on Feb. 25, 2014. Video still: Center for Spatial Studies, UCSB 2014.

Sony PTZ cameras on rails as part of Swarm Vision, 2013, George Legrady, Marco Pinter, Danny Bazo.

“**Swarm Vision**“ explores the translation of rules of human photographic behavior to machine language. Initiated by research in autonomous swarm robotic camera behavior, “SwarmVision” is an installation consisting of multiple Pan-Tilt-Zoom cameras on rails positioned above spectators in an exhibition space, where each camera behaves autonomously based on selected rules of computer vision that simulate aspects of how human vision functions. Each of the cameras are programmed to detect visual information of interest based on separate algorithms, and each negotiates with the other two, influencing what subject matter to study in a collective way.

Viewers can perceive both individual robotic camera behaviors (microcosmic) and their relationships to each other (macrocosmic) on 2 large screens. Visual fragments of spectators who enter the viewing space populate the images, leaving an imprint of their presence that become erased over time as the stream of new images replace the older ones.

Source: www.georgelegrady.com

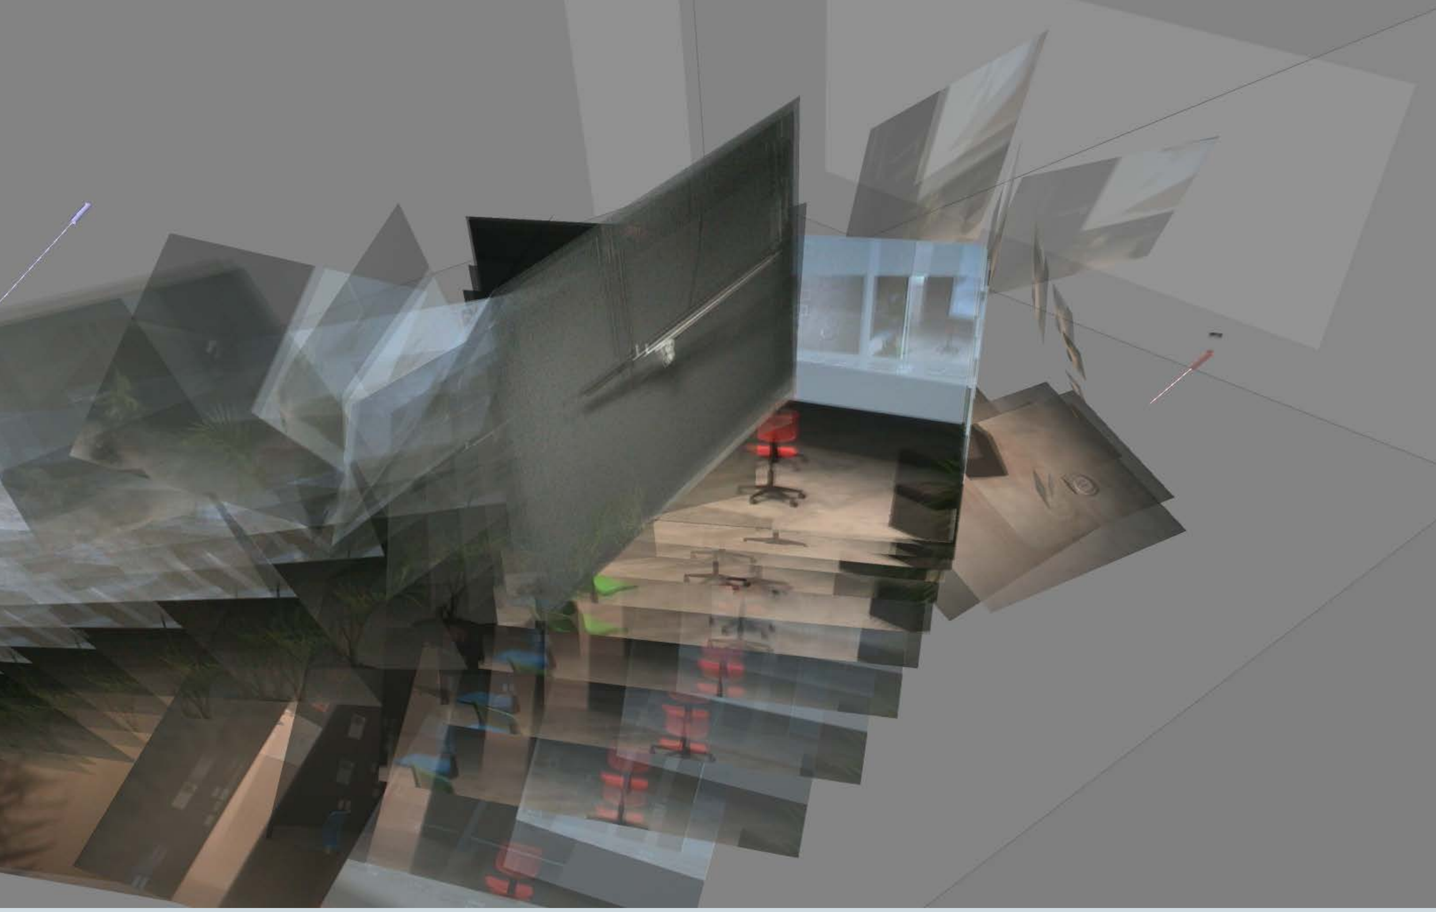

Examples of different 3D Scene Overviews, “Swarm Vision”, 2013, George Legrady, Danny Bazo, Marco Pinter.

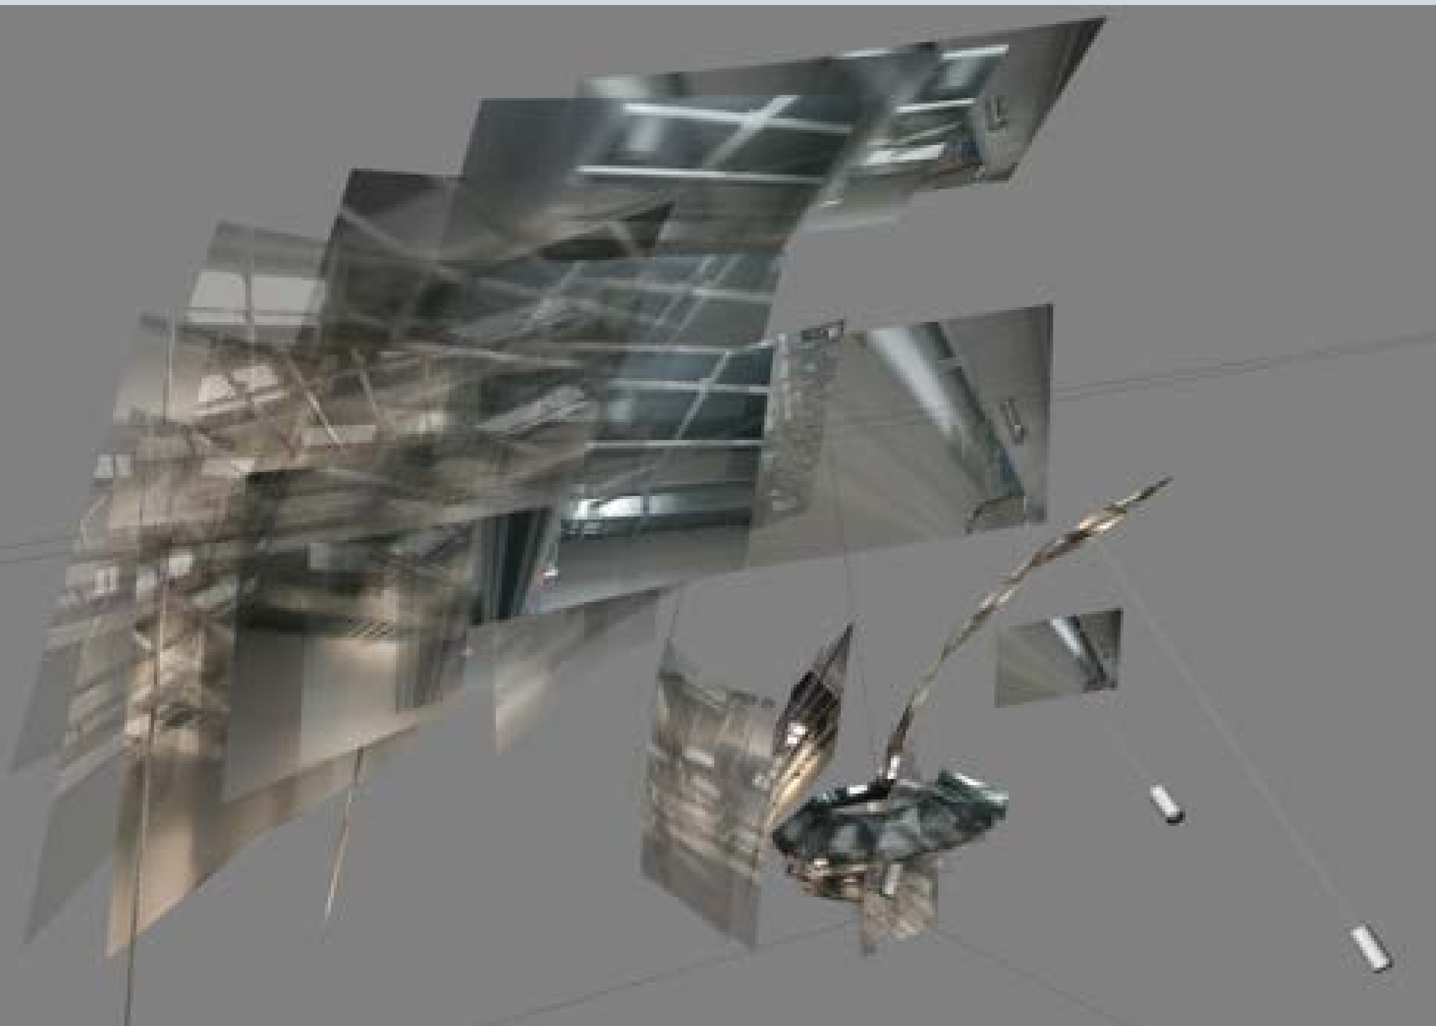

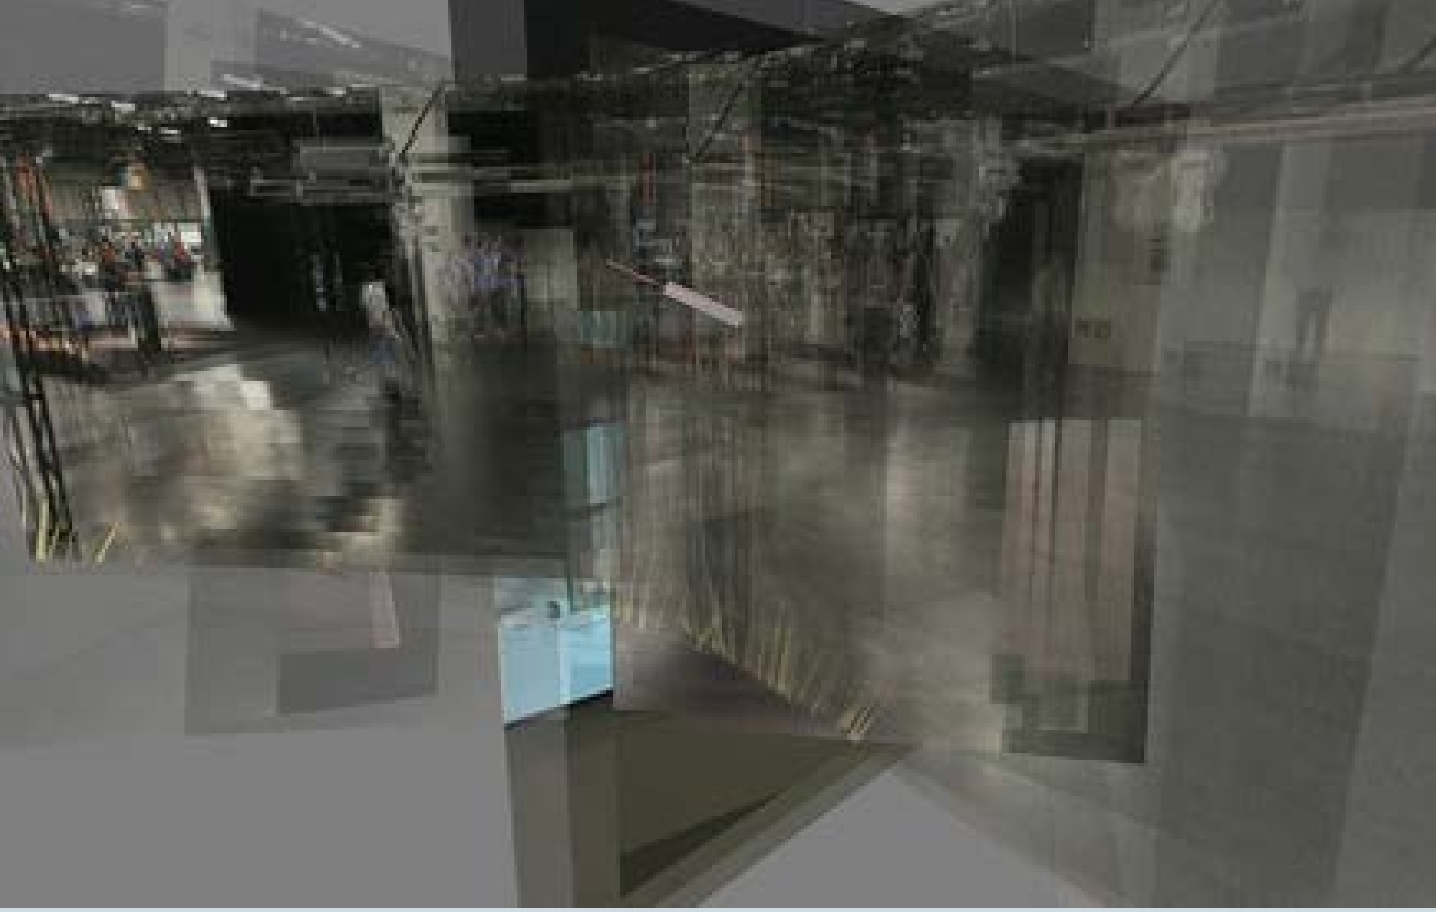

Visualization Screens: In the installation, four visualizations are featured on two screens/projections. The first screen features what each of the three cameras "see" - a depiction of what their vision algorithms are currently processing. The second screen shows an overview in a 3D reconstruction of the environment featuring a live video stream of the location of the cameras, and of the images they generate. Each camera continuously produces 10 still frames per second, and fills the 3D space with up to a hundred images per camera resulting in a volumetric form of layered stacked photographs that continuously changes as images fade away. The images' sizes and locations are determined by the locations and poses of the cameras, as well as their focal planes and focus locations at a given moment. The 4th visualization features the sum of activities situating all generated images and the three camera locations within a reduced virtual 3D spatial reconstruction of the exhibition space.

Contributions: Danny Bazo has a background in Film Studies, Engineering, and Robotics. His contributions include building most of the custom hardware and software development. Marco Pinter's background is in dance performance and kinetic artworks. His contribution to the project also includes his expertise in live video technology, robotics, and telepresence. George Legrady is project manager and brings conceptual directions based on his background in photography, conceptual art, and interactive digital media installations.

Source: www.georgelegrady.com

generate misclassified data entry, or system errors in the electronic transcription of data.

HK – You once said about your "Pockets Full of Memories" (2001-2007) installation that "any artwork that functions to gather data creates through necessity another artwork, consisting of the analysis of the collected data."[3] And in the installation "Swarm Vision" (2013) robot cameras simulate the photographer's gaze and perception by artificial intelligence. Is the computer substituting the artist in the long run? Or is the computer just a new kind of intelligent companion?

GL – In each of the artworks mentioned, they have been specifically designed to collect data so that the data could then be analyzed after the fact. This was an original approach in 2001 but of course all businesses do this today, in many instance, the services/sales they produce are just an alibi to collect the data as that is what is determining decision-making today. The bestseller "Big Data" from Mayer-Schönberger and Cukier shows how data collection, data correlation results in value, control and competitive advantage. In my case, I was purely driven by curiosity and feedback, the desire to get a sense of how the public understands the project, to what degree they will explore the boundaries, and what I will learn in the process that I would not have thought of otherwise.

LS – Speaking of learning: as we mentioned earlier, you are professor at UC Santa Barbara. Is the history of visual media part of your teaching?

GL – I am teaching in a practice-based program with a significant engineering foundation. One of the five core courses is titled "Art & Technology" and the course' content is determined by the faculty's expertise. We unfortunately do not have a "History of Visual Media" course. I used to focus on artistic practice and methodology and introduce examples from contemporary art and digital media art. The other core courses include Music & Technology, Digital Signal Processing, Multimedia Engineering, and Computer Graphics.

LS – What seminars are you offering right now?

GL – I have three major courses: In the fall I teach "Arts & Engineering / Science Research" seminar-type course which looks at how different disciplines approach research methodologies. We visit engineering and science labs each week, and ask the scientists to describe to us how they get from analysis of data to discovery, what their research methodology is, and what the process is by which results are achieved. We then discuss artistic methodologies and make comparisons. We ask the scientists to what degree does aesthetics play a role in the process of their discovery and representation? By aesthetics I mean decisions and observations based on the senses, an insight, a perception, etc.

In the winter I teach a studio course titled "Visualizing Information" which

Given that critical decisions are made at the conceptual, aesthetic and software design level, the ideal is to be conversant in both.

is very intense as it covers much in ten weeks: A production course on techniques of 1) data mining, 2) data aggregation, and 3) visualization in the java based Processing environment. Knowledge acquired include 1) how to identify and retrieve significant data from a dataset with MySQL, 2) develop skills in the fundamentals of visual language through programming, 3) visualize abstract data to reveal patterns and relationships, 4) normalize data to enhance legibility and coherence, and 5) implement interactivity within 3D volumetric visualization.

In the spring I currently teach a production studio course titled "Optical/Motion Computational Processes" with a focus on motion-capture and depth sensing using the Kinect or Asus Xtion sensor which students can use to create a work based on movement sensing and feedback systems through presence of spectators. All three courses integrate knowledge and methods from both the arts and engineering.

LS – What type of seminars work well and where do you run into problems? As your curriculum is open for students from different disciplines, it probably acquires interdisciplinary work skills, is that true?

GL – The MAT program includes students with a broad range of backgrounds such as computer scientists, engineers, physicists, electronic composers, audio technologists, graphic designers, architects and media artists. Each student arrives with a set of expertise, but is also challenged to acquire new skills. The intended goal is that students come in with one or more expertise and leave transformed and hybridized.

HK – You are also the director of the Experimental Visualization Lab in the Media Arts & Technology. What role does the software play in your curriculum?

GL – The Media Arts & Technology program is an arts-engineering program and computation is at the core of what we do, how we engage through the design of software and hardware. The Arts side of the faculty includes electronic composers, virtual architects, a systemics-based artist and I am an image-maker. Our engineer colleagues are specialists in haptics (touch), gesture recognition, Computer Vision, Augmented Reality and Computer Graphics. So computer programming is at the core of what we all do, and how we hybridize between our disciplines.

The challenge is to achieve our goal of interdisciplinary hybridity.

HK – What software do you use with the students?

GL – In my datavis course, we begin with exploring data collected through my Seattle Central Library artwork with MySQL queries, then follow visualizing using the java-based Processing language created by Casey Reas and Ben Fry. We correlate data from diverse sources (such as Apple, New York Times, Instagram, etc.) with JSON which is a data interchange format. My audio colleagues may use Python, the computer scientists use C++, architecture designers use Mathematica.

HK – Being a user of software seems to make the artist a second author besides the creator of the software. Do you think that is true? Or let's put it differently: How relevant is the relation artist/programmer to you?

GL – Given that critical decisions are made at the conceptual, aesthetic and software design level, the ideal is to be conversant in both. At this stage of my practice, I am privileged to collaborate with my students who develop much of the software for my projects. We work very closely through an interactive feedback and iterative process so that the evolution of the code production is guided by much interaction.

HK –Do you think that being an artist today requires to become an IT-expert in order to regain a complete freedom of expression?

GL – That was a big topic of debate in the 1990s –should the driver of the car know how the engine was built. It all depends on how one purposes the technologies. Traditional artists may use digital tools such as Photoshop to enhance their photographs, or explore painting methods whereas media artists or artists-engineer hybrids will create their own software tools as instruments for specific projects. Media artists are interested in investigating new possibilities through inventing new software. Those with a theoretical focus create software to research the impact of technologies on our way of understanding the world, how to interact with the world, and what new technological representations can be achieved that have not previously existed. As mentioned earlier, this way of thinking takes us back to the 1980s and Jean Baudrillard's discussion of Simulacra – to what degree do we situate our understanding to interactions with the world, as contrasted to interactions mediated through technologies? And which one seems "natural" as the conventions of realities transform to increased technological mediation?

LS – What are the consequences for the academic curriculum? Being at the intersection of arts and technology, where do you see challenges in teaching?

GL – The challenge is to achieve our goal of interdisciplinary hybridity. It was difficult to achieve when the program began fifteen years ago as our students were set in their discipline-specific ways, and also the culture had yet to embrace interdisciplinarity to the degree we are witnessing today. There was a turning point around 2006 and we have been very successful since.

LS – In a similar way, we see interdisciplinary hybridity to be also the future of digital art history.

GL – You may be suggesting a creative approach to art history where software development becomes part of the analysis of data (historical and visual) and used to describe things in a creative (possibly non-linear way. There is nonetheless a separation between the practice of creation and the practice of reflection on the creation process. The first engages with the expressive process of representations and the latter is an analytical study of how that process takes place. I have had a theorist mention to me that their work begins when the artist's work is complete, and some artists have concerns that the theorists transform the intent of the work, but one cannot control how a work is perceived. The third model is the hybrid collaborative interaction between producer and analyst.

HK – Hence, a dialog of artists and art historians is critical – as this interview shows. We thank you very much for this inspiring conversation.

Notes

1 http://www.mat.ucsb.edu/~g.legrady/glWeb/publications/publ_art/textimage.html
2 Vilém Flusser: Towards a Philosophy of Photography, Göttingen, 1984 p.25 and 27.
3 Making Visible the Invisble. Iker Gil interviews interactive media artist George Legrady, in: MAS Context: Issue 7, Iker Gil, Chicago http://www.mascontext.com/issues/7-information-fall-10/making-visible-the-invisible/, Date accessed: 19 March 2016.

George Legrady is Chair of the Media Arts & Technology PhD program at the University of California, Santa Barbara, director of the Experimental Visualization Lab, and professor of digital media in the College of Engineering and the College of Humanities and Fine Arts. He is an internationally published scholar and exhibiting media artist, a pioneer in the field of interactive digital media arts.

His current research engages with data visualization, robotic computational integrated photography, and digital visual ethnography. He has received awards from Creative Capital Foundation; the Daniel Langlois Foundation for the Arts, Science and Technology; the Canada Council; the National Endowment for the Arts, and the National Science Foundation, and this year, a lifetime achievement prestigious Guggenheim fellowship.

Correspondence e-mail: legrady@arts.ucsb.edu

Case Studies

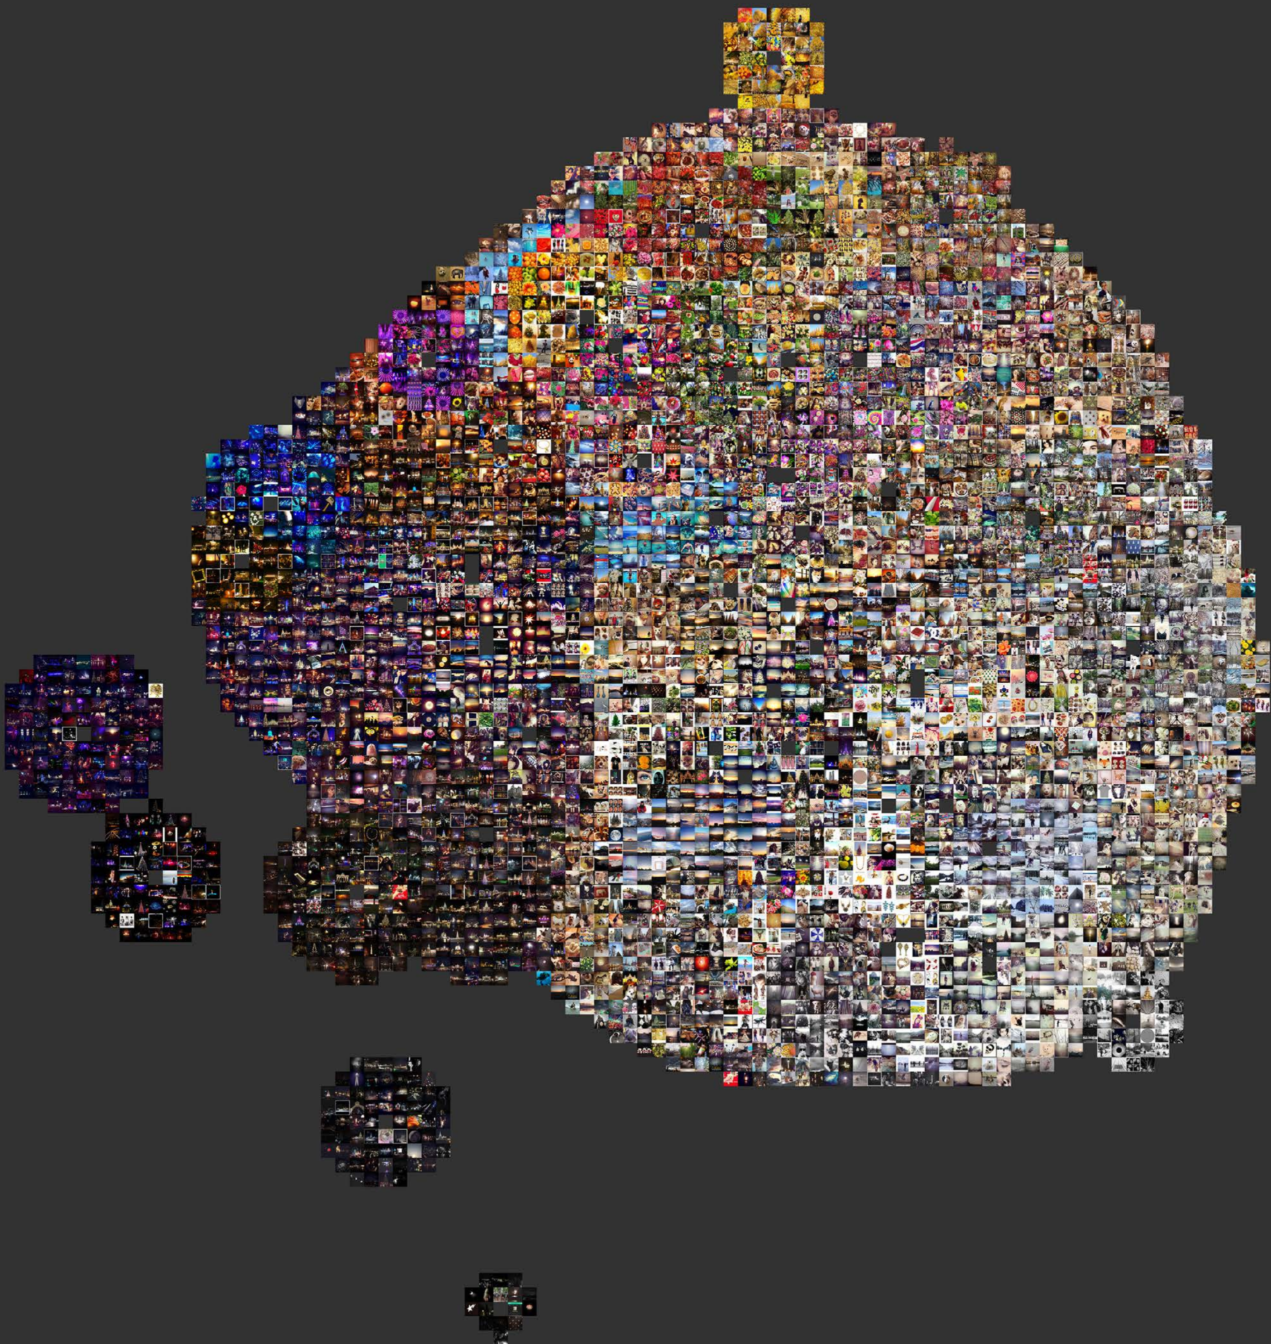

Fig. 1. Growing Entourage with 50 clusters of Instagram photos machine-tagged under the heading ‘nature’.

Invited Article

Direct visualization techniques for the analysis of image data: the slice histogram and the growing entourage plot

Damon Crockett

Abstract: The following is a description of two novel techniques for the direct visualization of image data. Direct visualizations of image data make use of the images in their original visible format. The first technique, the slice histogram, arranges slices of images as histograms, organized by both visual and non-visual variables. The second technique, the growing entourage plot, organizes high-dimensional clusters of images on a 2D canvas by projection. Both techniques are designed for exploratory analysis of image datasets.

Keywords: direct visualization, images, exploratory data analysis

Introduction[1]

In the following, I discuss two novel methods of direct visualization: (1) the slice histogram; and (2) the growing entourage plot. 'Direct visualization' is a term coined by Lev Manovich (2010) to describe visualizations of image or video data that make use of the data in its original visible format. The general form of direct visualization for static images is the 'image plot'—an arrangement of images on a single digital canvas.[2] The techniques described here are instances of this general type.

The active form of direct visualization for the digital humanities is media visualization (Manovich 2012) — direct visualization of cultural media collections—and the Software Studies Lab at UCSD has been largely responsible for its promotion and development. Several early forms of media visualization are cited in (Manovich 2010), and Manovich and collaborators have produced many examples of media visualizations, including Mapping Time, a chronologically sorted montage of every Time magazine cover from 1923 to 2009;[3] The Shape of Science, a visualization of over 10,000 pages each of the magazines Science and Popular Science;[4] Manga Style Space, a visualization containing over 1 million Manga pages sorted by their visual features, meaning simply their visual properties, basic things like average hue or brightness, or more sophisticated measurements like entropy;[5] and Phototrails, a series of image plots, each containing thousands of Instagram photos sorted by basic color features.[6]

In these and in many other projects, Manovich has demonstrated the power of this method for the digital humanities, and the tools I describe here are both extensions of this general method and, perhaps, attempts to show direct visualization relevant not only for cultural media but for all image data and to point in the direction of a new and purely perceptual paradigm for image data analysis.[7]

Direct Visualization

Direct visualizations are special cases of glyph visualizations, a class of statistical plots that present data points as glyphs—icons that carry information by way of their non-relational characteristics, things like size, shape, color, etc.

The dominant form of relational data visualization is the scatterplot—a collection of points plotted along (usually) two axes that reveals the relationship between the variables mapped to the axes. Because traditional scatter points have no non-relational characteristics—they are in fact point locations and not objects at all—they carry information only by their spatial positions. Traditional scatterplots, then, can present only as many informational dimensions as there are plotting axes. If, however, each scatter point is made a glyph with n non-relational characteristics, the dimensionality of the visualization is increased by n. Direct visualizations can therefore be understood as limit cases of glyph visualization, because they preserve, strictly speaking, all of the visual information in the dataset.

But the mere presence of the information in the visualization does not guarantee its being readable for the viewing subject. Important for direct visualizations are particular choices about sorting or otherwise organizing the images on the canvas in order to reveal patterns. Our lab has pioneered a suite of techniques for organizing images as plot elements on large digital canvases. We use sorted rectangular montages (Image Montage),[8] Cartesian scatterplots, and perhaps most identifiably, polar scatterplots (Image Plot).[9]

Image Histogram

Recently, we've pioneered a new technique, one that has roots in our Selfie City project:[10] the image histogram. The most basic form of the image histogram simply gives the distribution of a single variable—like average brightness or hue, time of day, geolocation, etc.—-and uses the images themselves as plot elements. Like the image montage, the image histogram gives every image its own place in the plot, and like the image plot, the image histogram uses an axis to organize data points. We have found this combination of characteristics to be very useful in presenting, e.g., temporal patterns in image data.

But because image histograms are direct visualizations, we needn't settle for the presentation of a single variable. The images themselves are right there on display, and although the particular choice of histogram (i.e., what gets assigned to the x-axis) dictates the horizontal sorting, the vertical sorting is still up for grabs and can reveal additional patterns in the data. We can therefore think of each histogram bin as a columnar montage that can be sorted as many times as we like. In Figure 2, we can see the results of multiple vertical sorts.

Slice Histogram

My own work with the lab began around the time we started making image histograms, and I confronted a problem: I wanted to make direct visualizations that reveal, with great clarity, the color properties of image datasets. Images wear their colors on their sleeves, of course, and so direct visualizations do carry color information—all of it, in fact—but again, particular choices about sorting can matter a lot.

Unfortunately, even our very best sorting choices for image histograms do not yield crystal clear presentations of color. And the problem is that images typically contain lots of different colors. When you sort them by color, how do you do it? Do you take the mean hue of the whole image? The mode? Do you look at all color dimensions, or just one? How you answer these questions will depend on your goals, of course, but the questions are less important if the images have very low standard deviation of their color properties. The more uniformly-colored the images, the easier it is to sort them by color. But we can't simply enforce uniformity in our data. What we can do is plot slices of images instead of whole images. This general idea is at work in various computer vision algorithms, and it makes good sense: images typically capture scenes, and scenes have parts, so we should at the very least be looking at those parts, whatever else we do. In computer vision (and human vision), the selection of parts can be very sophisticated (in human vision, this is roughly the function of visual attention), but it is computationally expensive. We need a way to get better color visibility without sacrificing computational speed.

I developed a technique that is fast and yields excellent color visibility (Figure 3; close-up in Figure 4). Each image is sliced into some number of equal-sized parts, color properties are extracted from the parts, and those parts are then plotted as an image histogram. The entire process, carried out with 1 million slices, can take less than an hour.

The number of parts will depend on the kind of images we use. In my own work, I set a criterion value for average standard deviation of hue (~0.1) and then find the minimum number of parts needed to meet the criterion. This approach ensures both

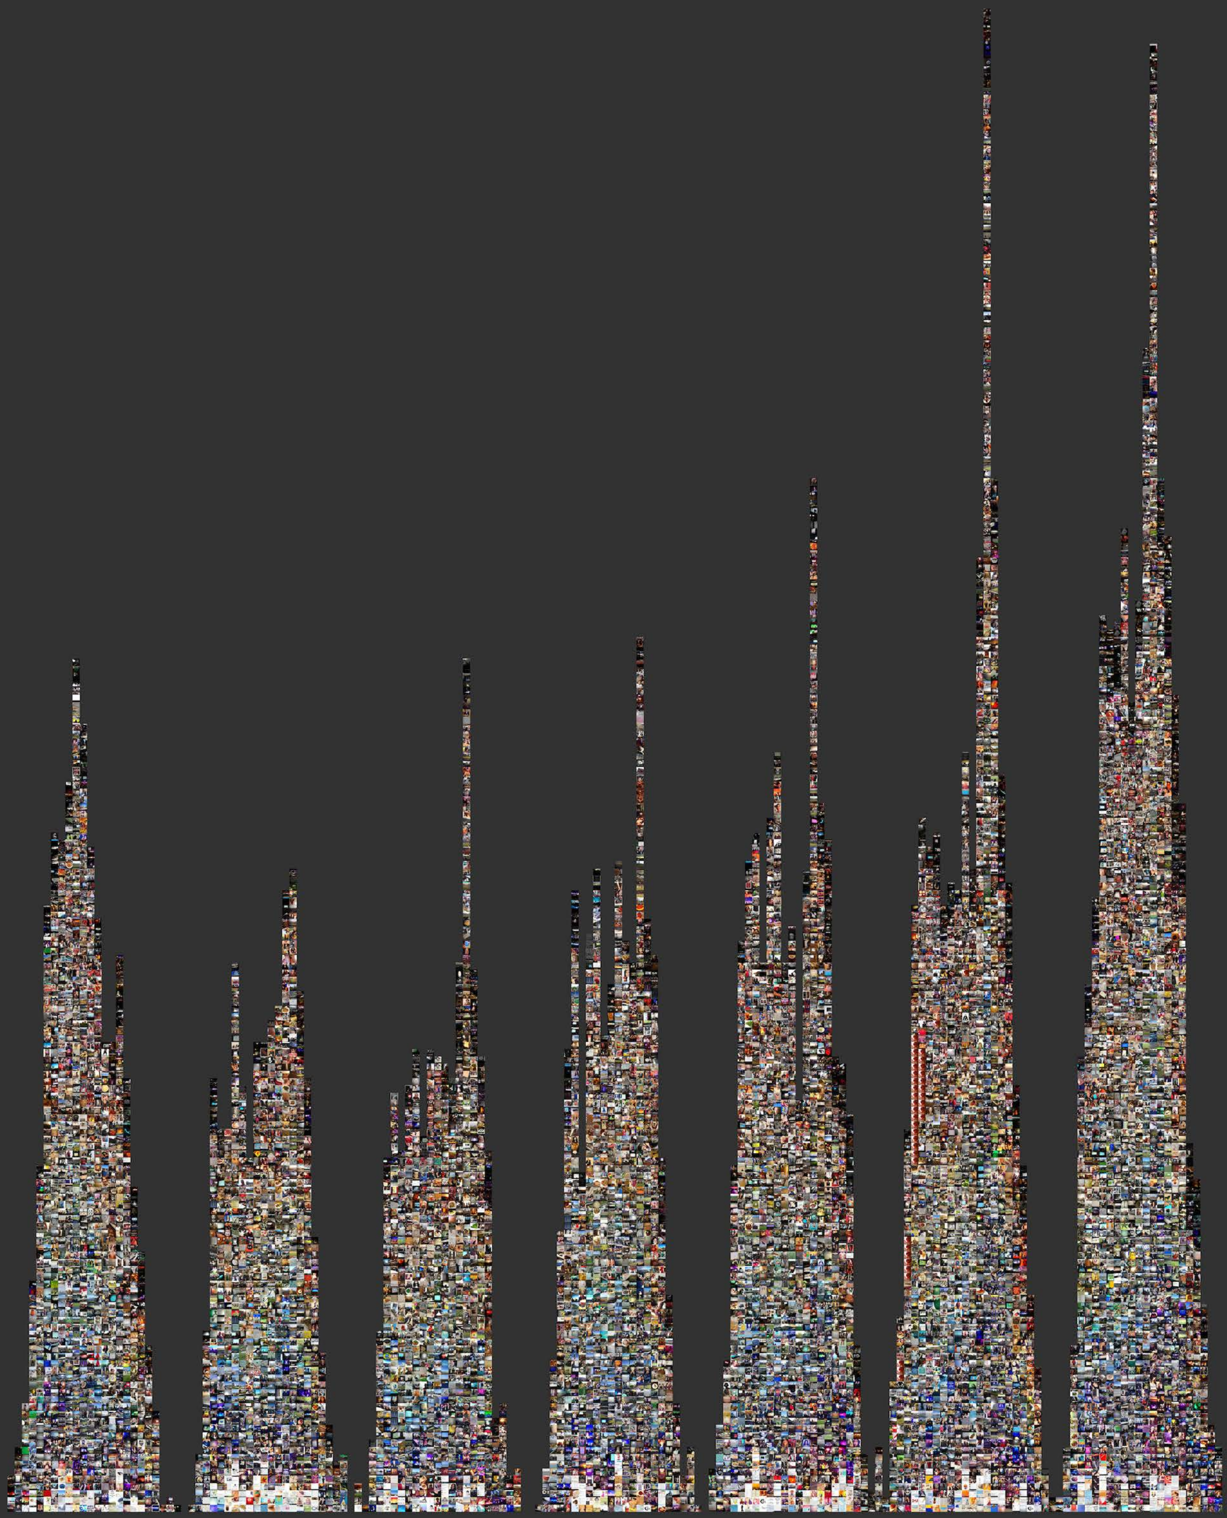

Fig 2. Image Histograms binned by hours over a week, sorted vertically by both brightness and hue.

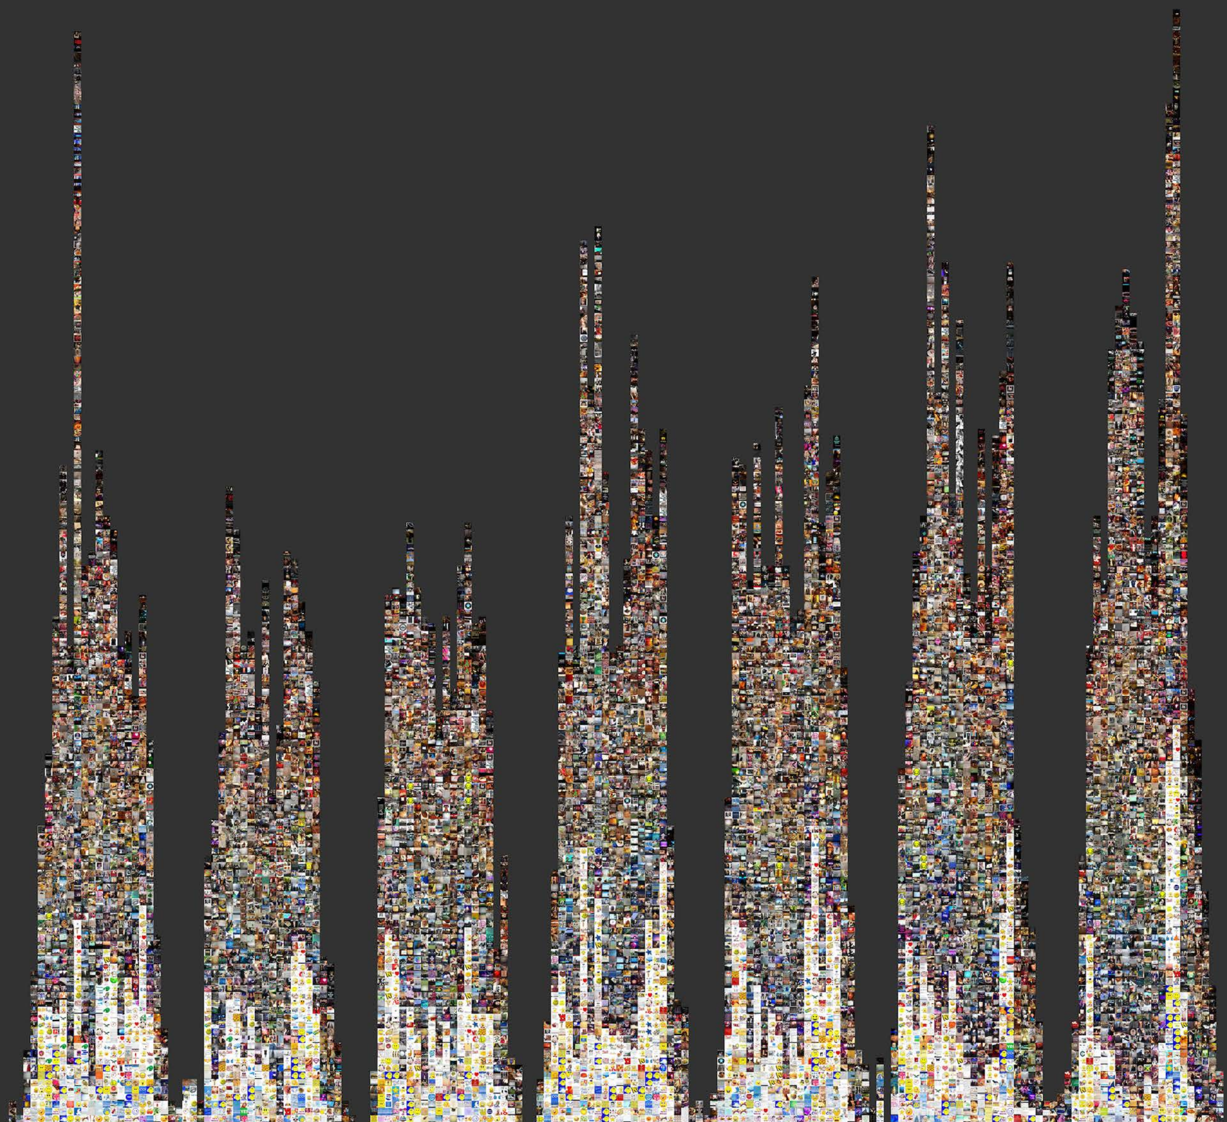

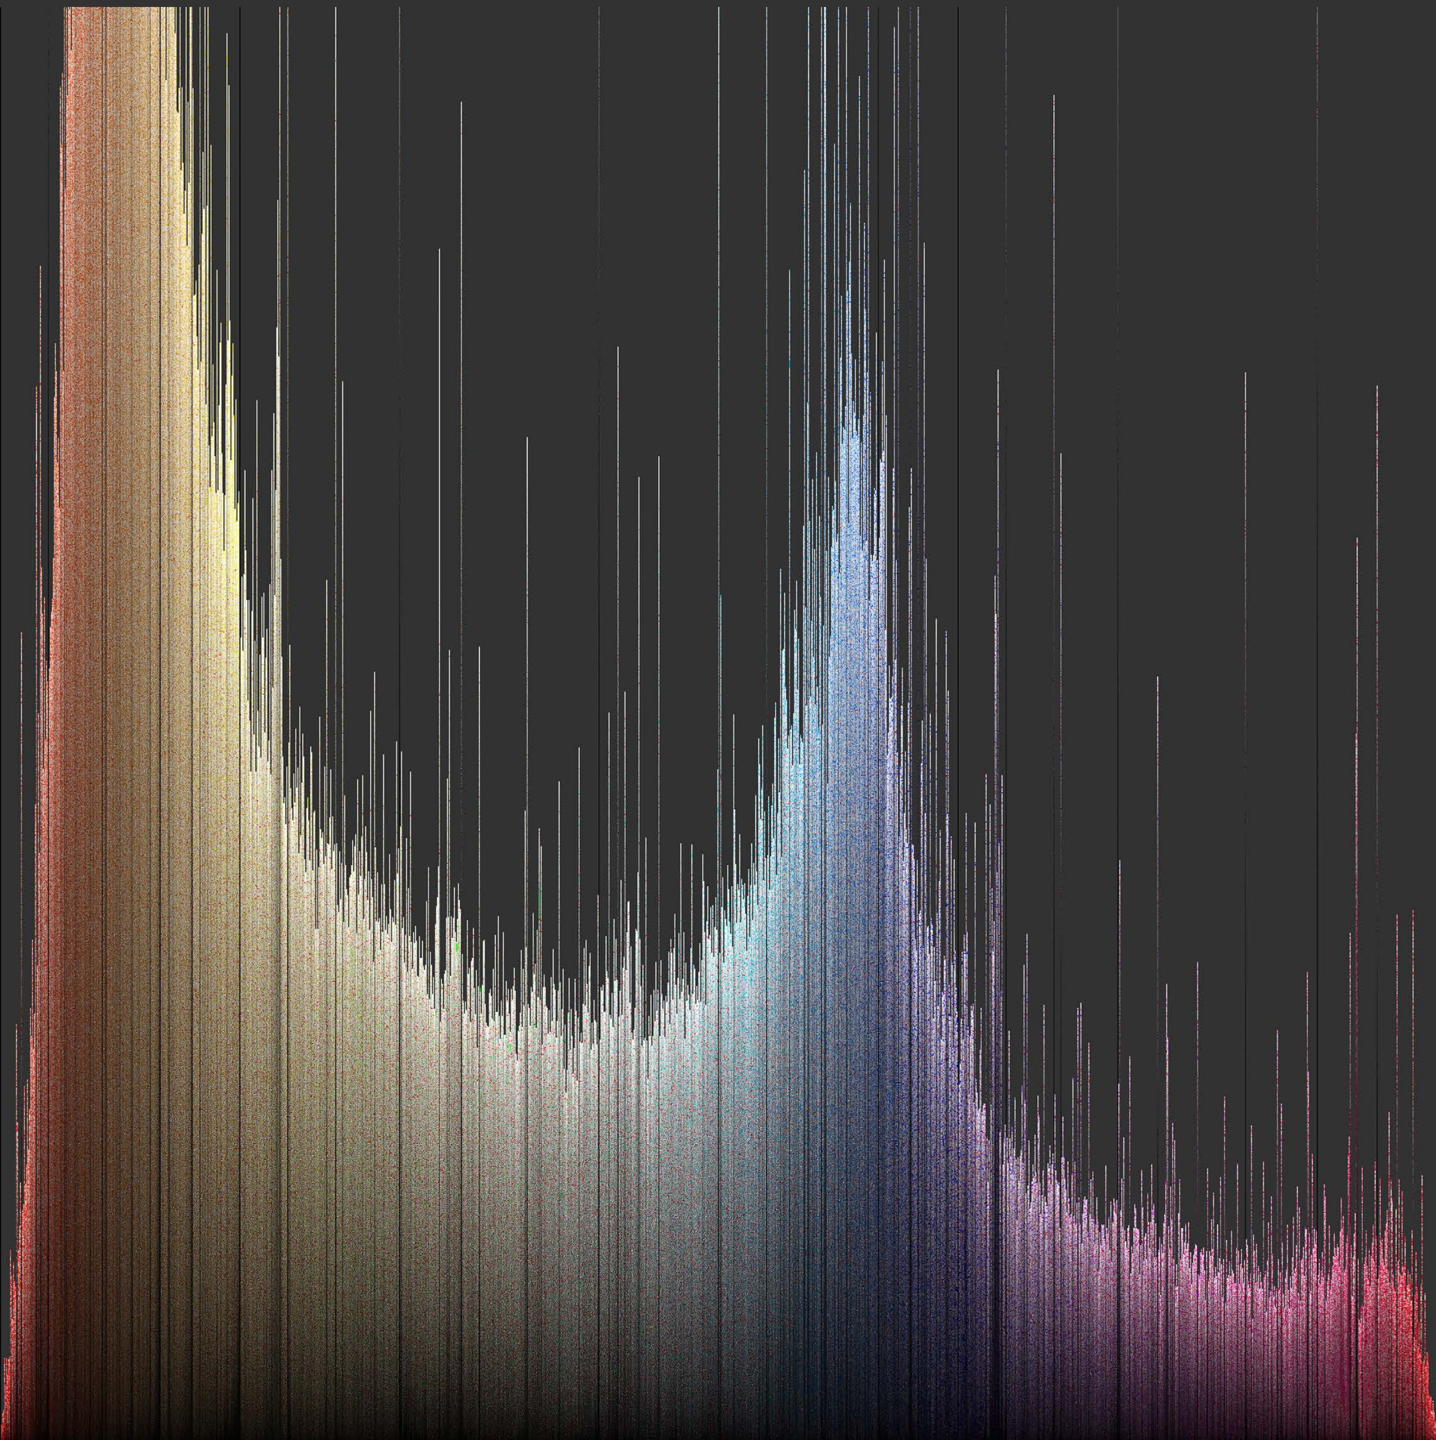

Fig 3. Hue histogram sorted vertically by brightness.

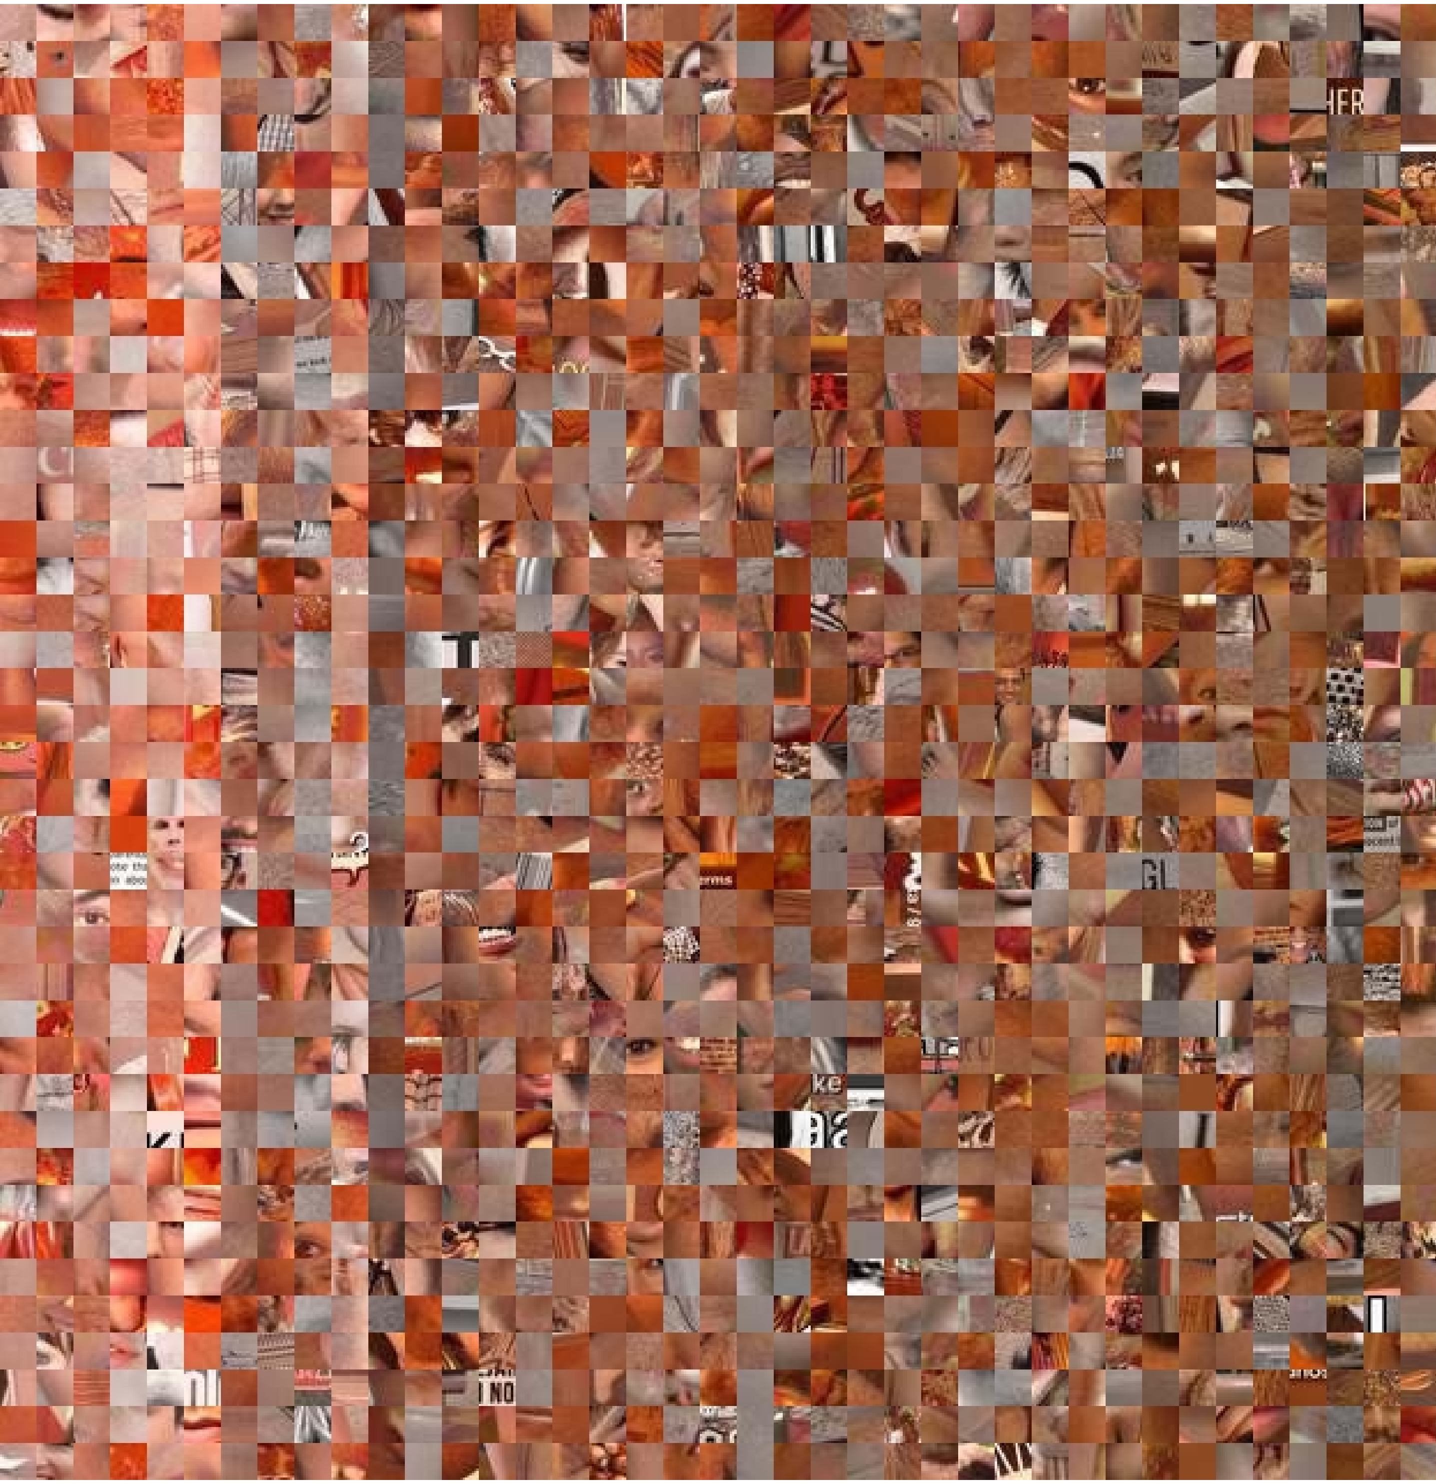

Fig 4. Close-up of slice histogram left.

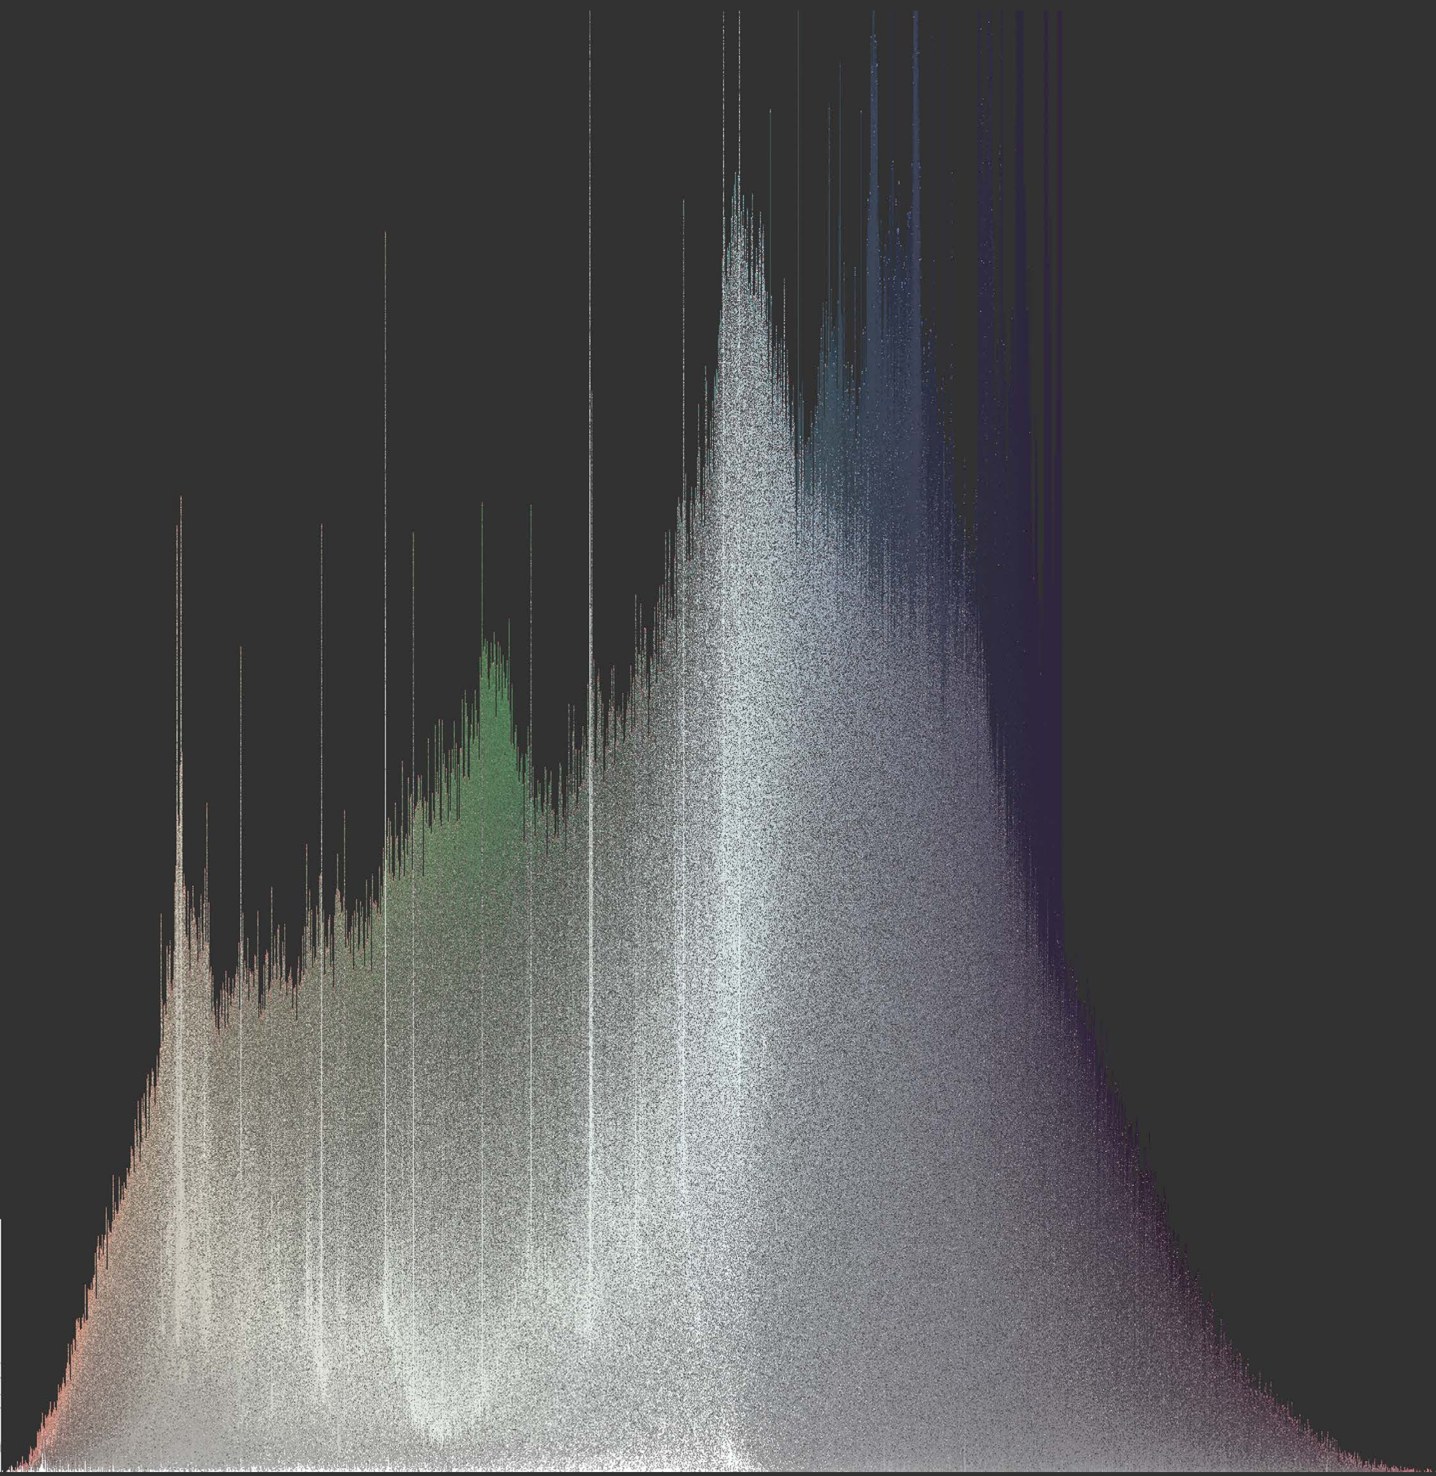

Fig 5. One million slices of a satellite image from downtown San Diego, arranged as a hue histogram sorted vertically by saturation.

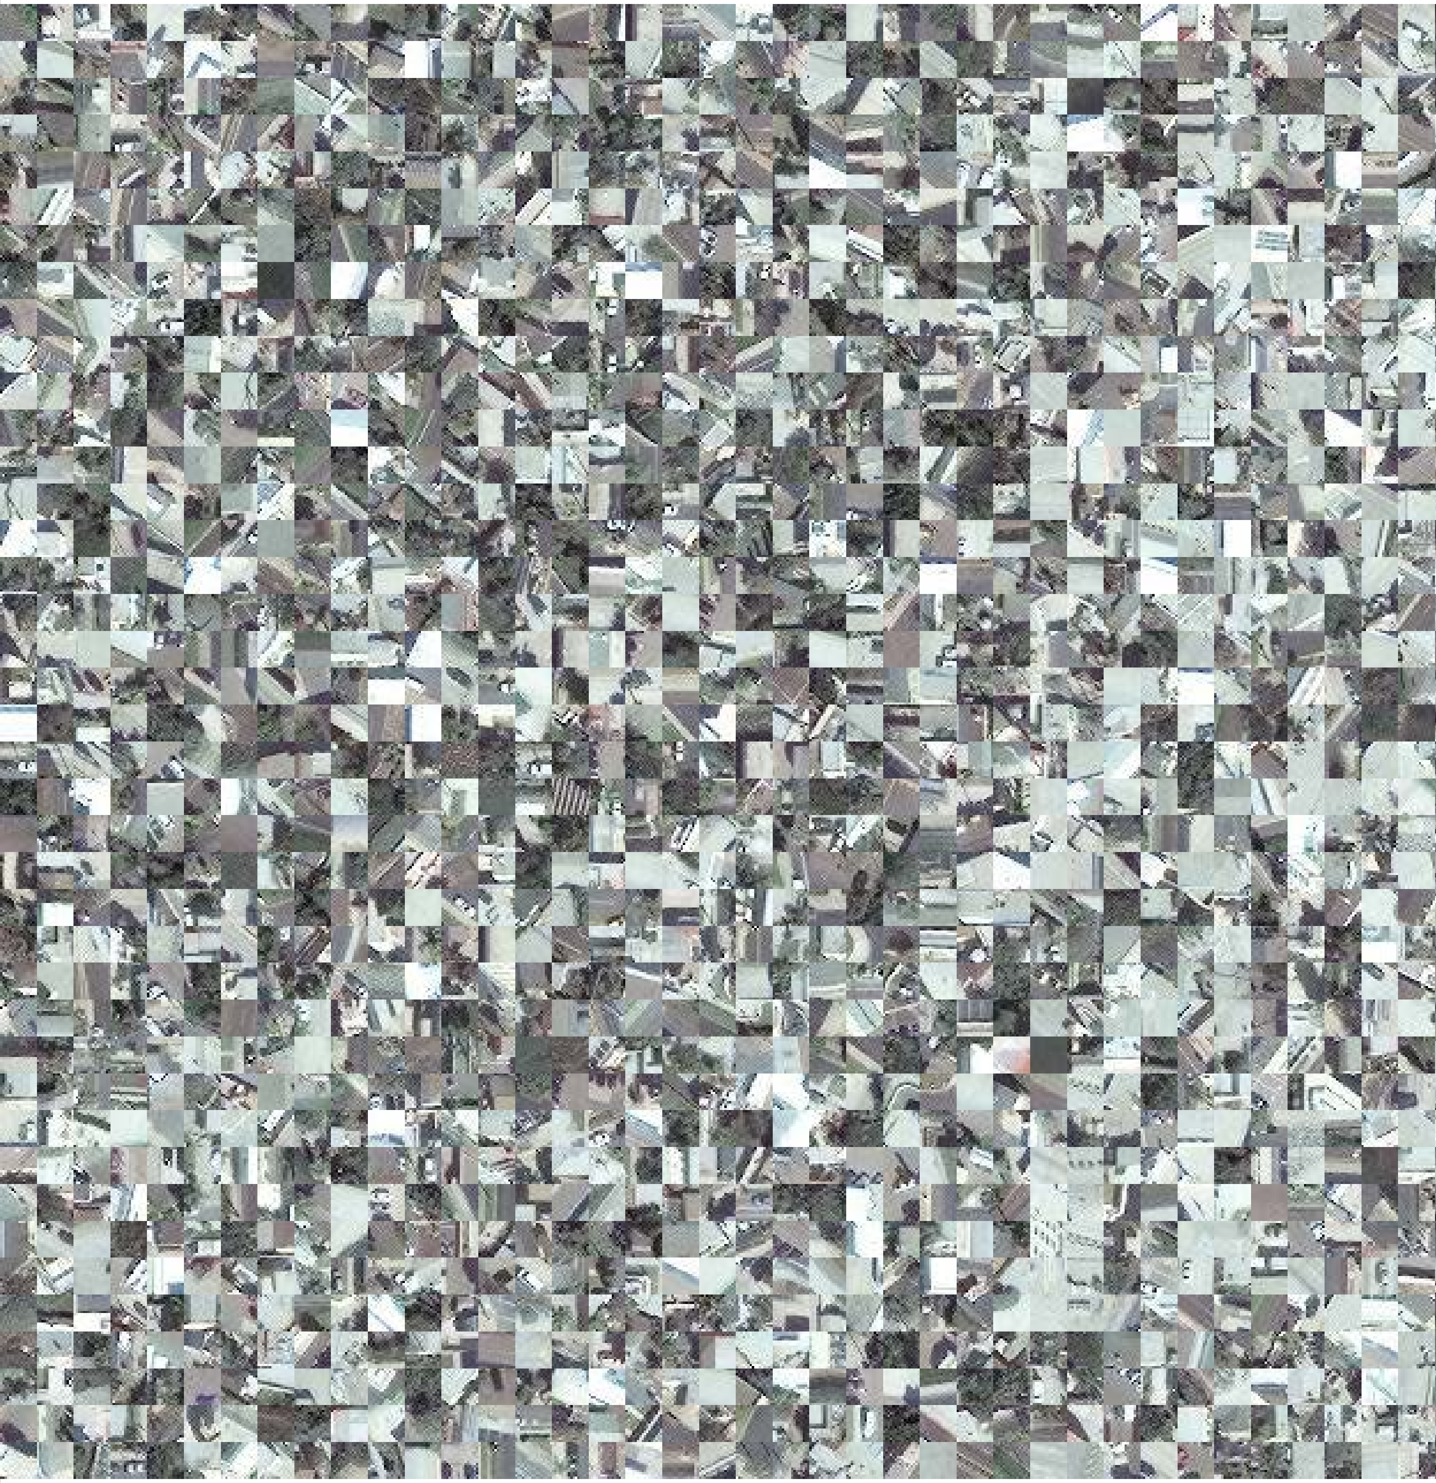

Fig 6. Close-up of slice histogram left.

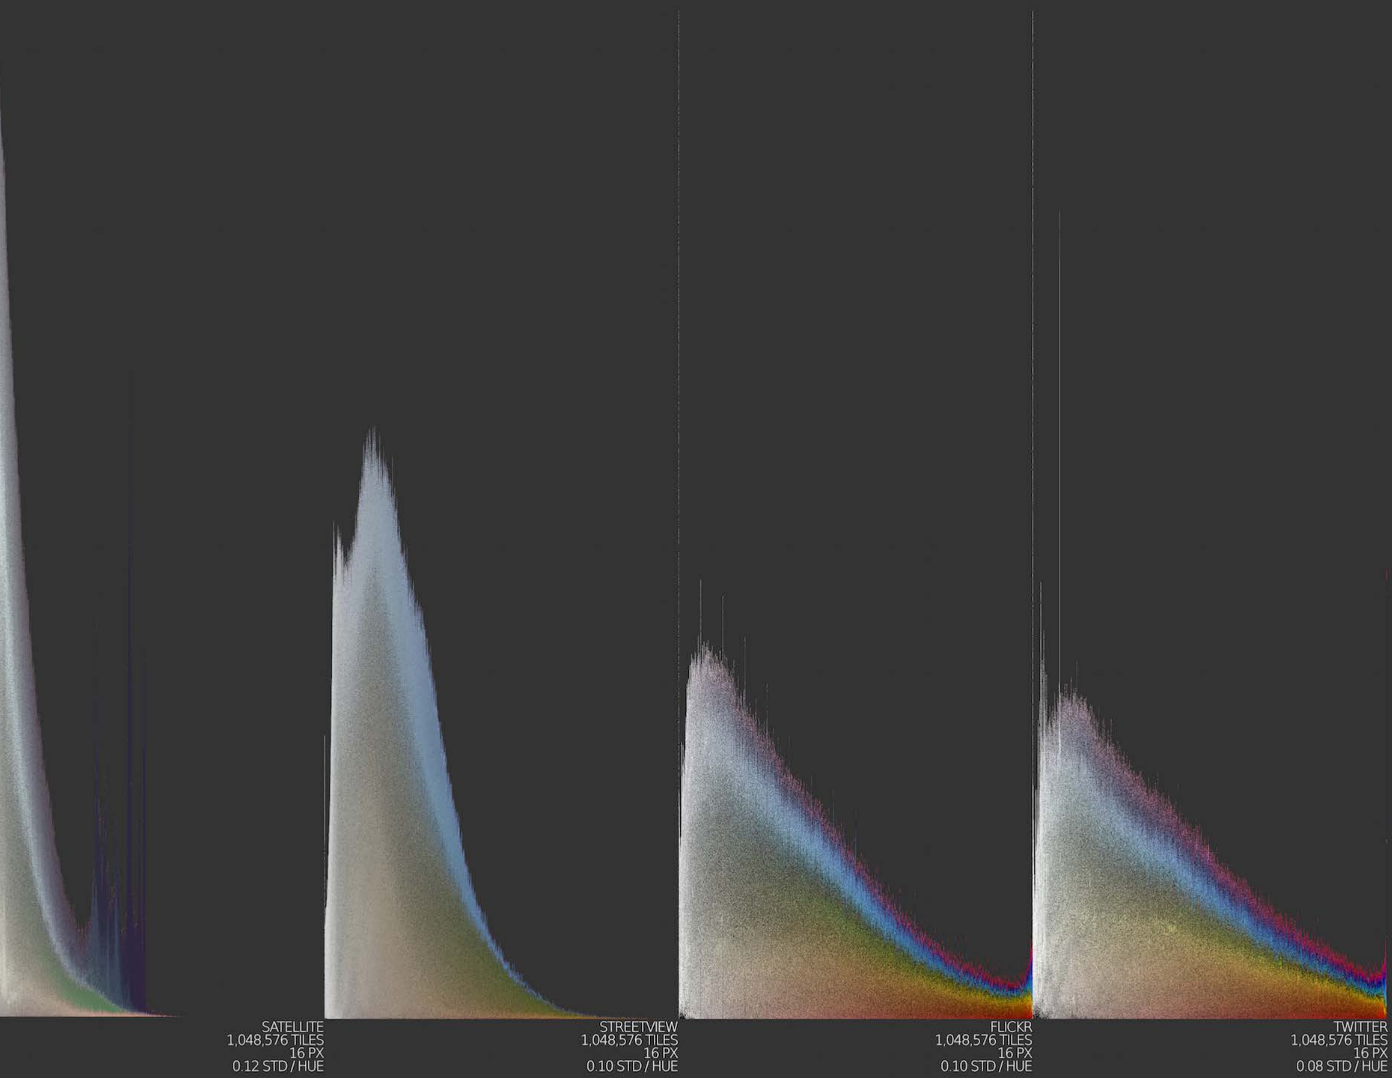

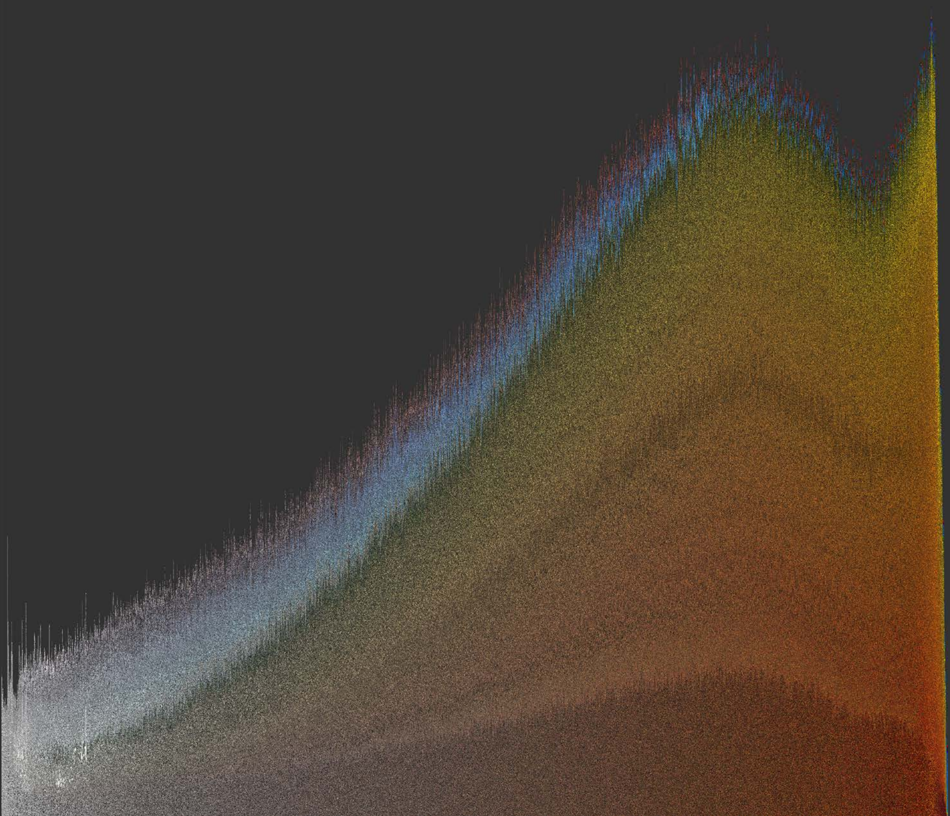

Fig 7. Slice histograms for four different sources of image data: satellite, Google Streetview, Flickr, and Twitter.

Fig 8. Slice histogram of Flickr images autotagged 'autumn' from New England.

Fig. 9. Growing entourage with 50 clusters of Instagram photos machine-tagged under the heading 'food/drinks/meals'.

Fig. 10. Close-up of plot above. Empty grid squares are centroid locations.

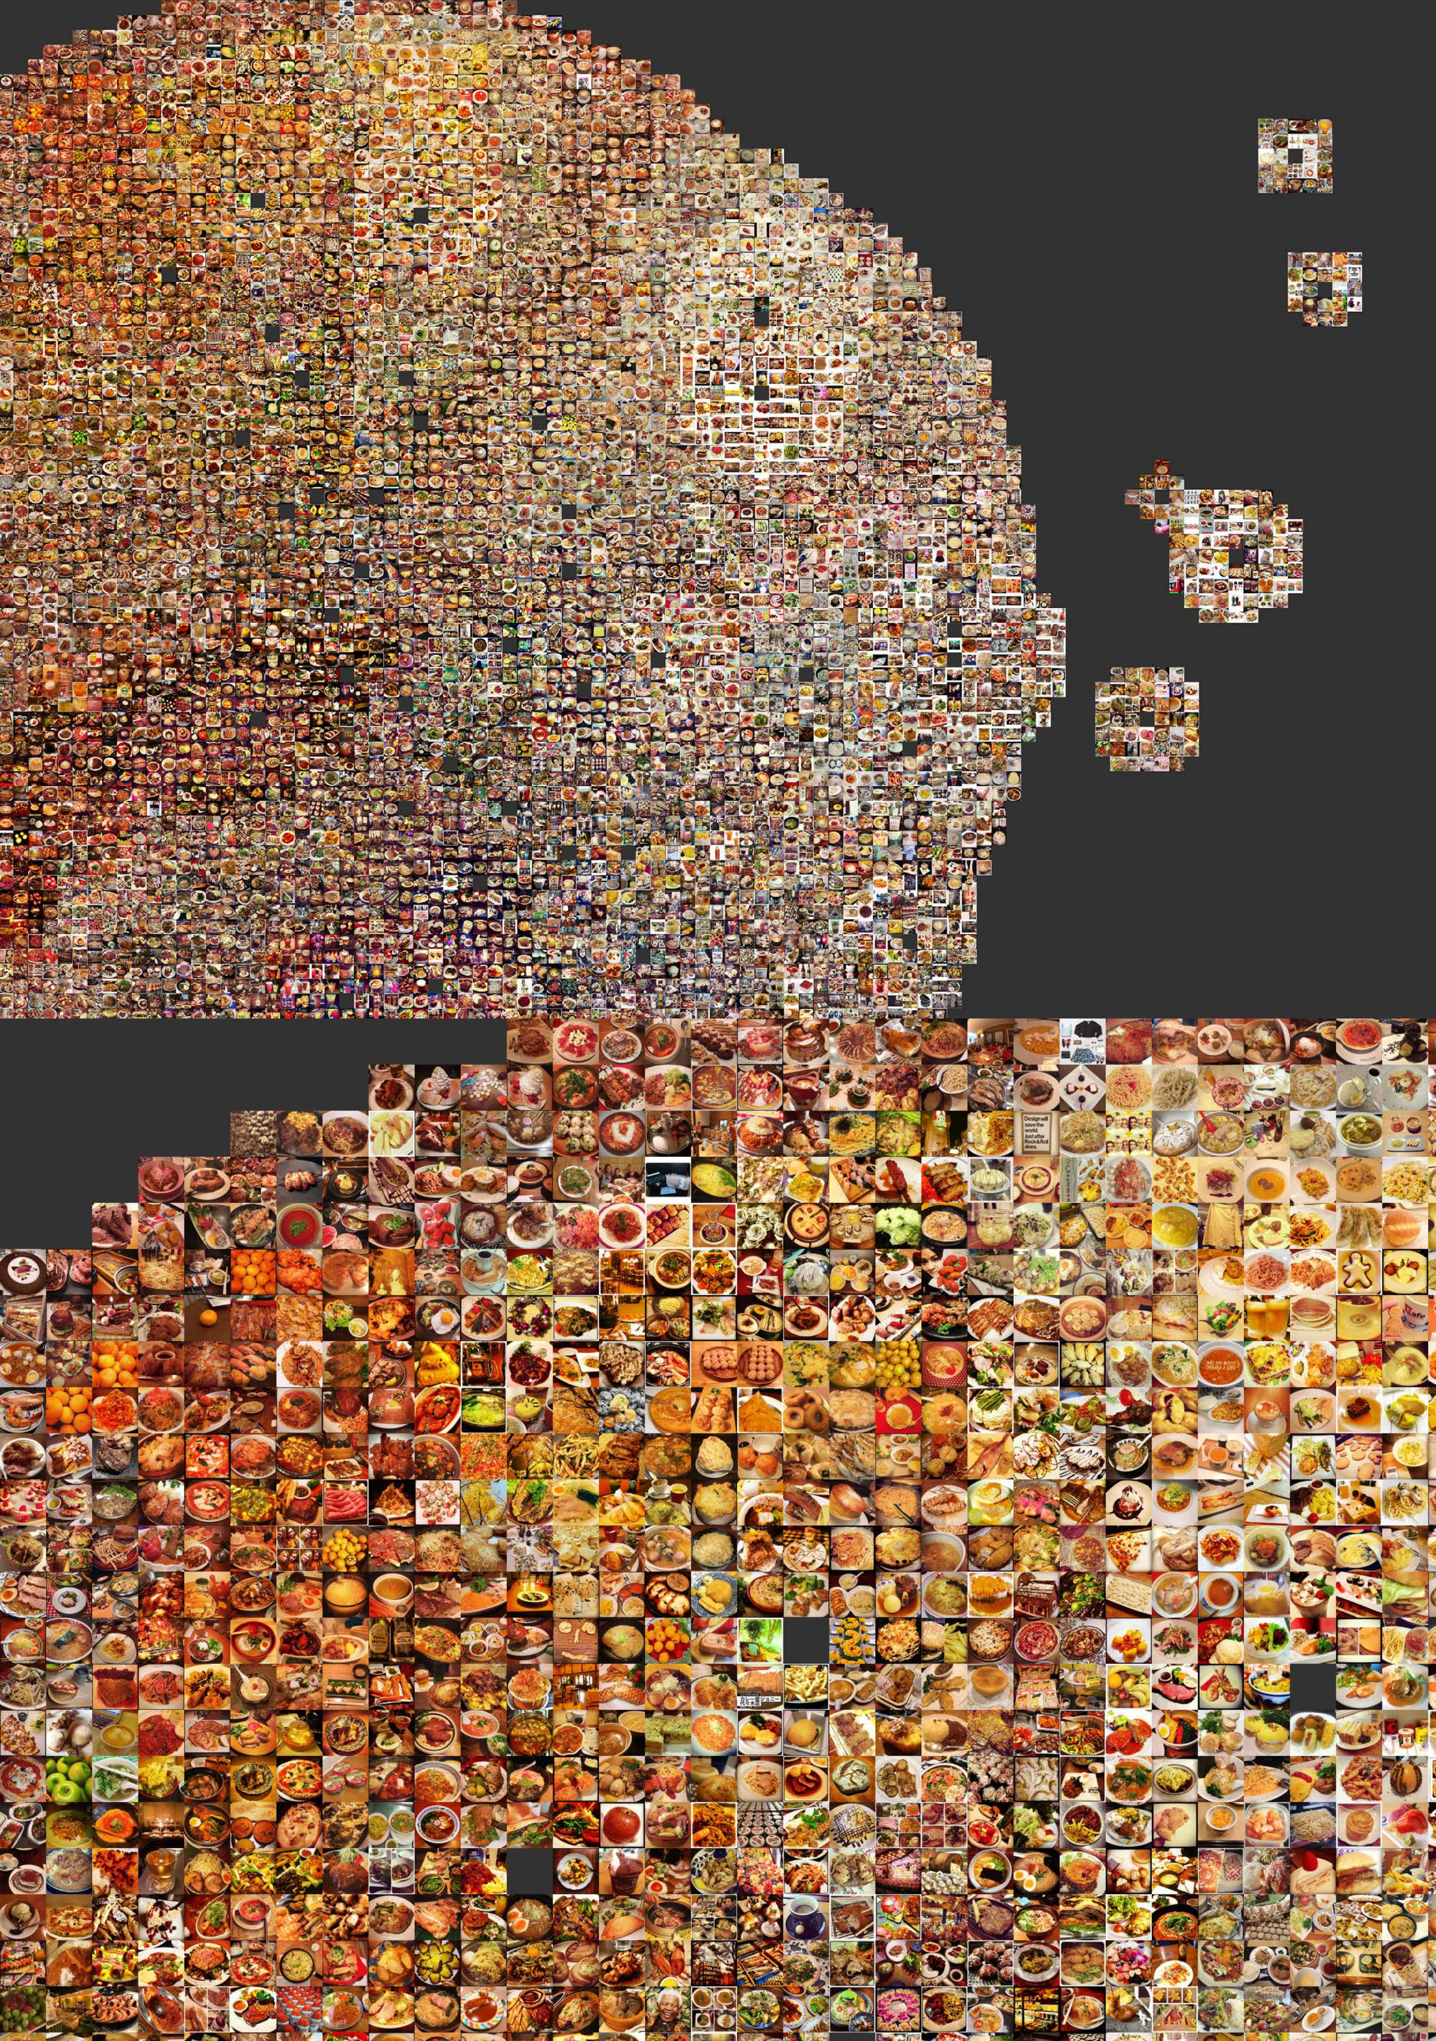

that the plots will make a smooth (and consistent) presentation of color and that they will carry as much object content as is possible for that particular source of image data at that particular criterion value (this as opposed to slicing all types of image data into the same number of parts). As you might imagine, the more 'zoomed out' the original image data, the clearer the object content at any particular criterion value. Satellite photos give excellent results, for example (Figures 5 and 6).

The object content is an important product of this technique. Because the plot is, like all direct visualizations, composed of the images themselves, it still carries much of the information it would have carried had we used whole images—that is, unless we set the criterion value so low that our slices are uniform fields of color (even single pixels!). This would yield maximal color visibility, of course, but we'd lose all the other information. Choosing a criterion value, then, is choosing the balance between color visibility and object information. As we lower the criterion value, we look at progressively smaller parts of scenes, and gradually, they lose their structure. Just where we choose to stop this degradation will depend on our analytical goals.

It is important to note that both the image histogram and the slicing technique are very general use tools, and admit of great variation in their application. And of course, the tools have no particular analytical power without some intelligent selection of data. For example, we might want to compare the color signatures of different sources of image data for the same region (Figure 7).

Or, we might want to visualize the colors of a particular concept, like 'autumn' (Figure 8).

Additionally, in making slice visualizations, it's not essential that the visualization take the form of a histogram. What is essential is that the plot elements have low standard deviation of whichever visual properties we're interested in, and that there be some method of sorting that groups together similar elements. That's it. We can transform these histograms into any sort of plot, or montage, or map we like. In the following section, I'll discuss my efforts to expand the space of possible forms of direct visualizations can take.

Growing Entourage Plot

Since 2007, our lab has been visualizing large collections of cultural images. These visualizations have used either metadata variables such as date and location, or basic image features such as hue, saturation, brightness, number of lines, and texture. In particular, sorting by hue, saturation and brightness turned out to be very useful for quick exploration of large image collections. More recently, however, we've expanded the scope of our analysis to include presence and characteristics of faces[11] and now a wide range of object and scene contents.[12]

Being able to use the latest computer vision techniques for the analysis of image content is very exciting, but it also brings new challenges. For example, how can we effectively visualize and explore the results of machine classification into many object categories? In this section, I'd like to discuss one particular method we've developed. This method visualizes high-dimensional image clusters using two dimensions. I call this method the 'growing entourage plot'.

The plots in this section are drawn from our current collaboration with Miriam Redi[13] on the clustering and visualization of large collections of Instagram images. Miriam extracted over 1000 image features that include image content (objects and scenes), photo style composition, style, texture, color and other characteristics. She then computed clusters of images using subsets of these features.[14]

The Challenge Of Visualizing High-Dimensional Data

In the field of information visualization, a great deal of energy is spent on the problem of how to present high-dimensional data on 2D canvases. There are at least three broad categories of solution: (1) preserve all features; (2) preserve some by selection; and (3) preserve some by redefinition. I'll discuss each of these in turn.

The first way of solution is simply to try visualizing everything. Such visualizations can be difficult both to design and to read. Direct visualizations are (or can be) officially of this type, but additional choices about sorting can make for big differences in readability. The effect of sorting on direct visualization is so important, in fact, that any feature not used for sorting is essentially invisible to the viewer.

The second way of solution is the one we've used most often: select some subset of features and use them for sorting. We might, for example, sort our high-dimensional image data by only brightness and hue. This is powerful and useful, but it does make invisible very complex sorts of similarity relations between images.

The third way of solution reduces the dimensionality of the data by defining a new, compact feature space. Principal Components Analysis (PCA) is the standard here, although there are others.[15] We can present images in, say, a 1000-feature hyperspace by projecting them to two dimensions. This approach has the advantage that similarity relations between images in 1000D feature space are preserved as best as possible during the projection to 2D, meaning that our visualizations can reveal very complex sorts of similarity between images. This way of solution is quite popular and is used by our lab.[16]

Visualizing High-Dimensional Clusters

But I'd like here to talk about a different approach to dealing with dimensionality, one that is quite common in data analysis but whose use in information visualization is less common: clustering. Dimensionality reduction algorithms are powerful and useful but suffer major data loss in most cases. You simply can't preserve all the complexity of 1000D relations after projecting to 2D. Clustering algorithms, however (e.g., k-means)[17], preserve a greater share of the relational data, because they find groups of data points in your original feature space (or whichever subset you choose).

Now of course, there is still the problem of visualizing these clusters. This is particularly difficult for traditional sorts of statistical visualization, because their plot elements carry information only by their spatial positions, and the human visual system can parse a maximum of 3 spatial dimensions. Thus, if we want to see these clusters in their 'natural habitat', so to speak, we're probably out of luck. Additionally, seeing clusters of points, even in a 2D or 3D space, is not particularly illuminating, since clusters, unlike classifier outputs, are not 'classes' at all and have no conventional meaning or significance from the outset. We derive the meaning or significance of a cluster of points in a high-dimensional space from the feature values of those points, and in order to see those features in a plot, we'll have again to confront the problem of presenting high-dimensional data in 2D (or 3D). For these reasons, you don't see many cluster plots (what you'll see sometimes is PCA scatterplots with cluster memberships coded by color).

Direct visualizations are at an advantage here. Because direct visualizations use images as plot elements, a simple presentation of each cluster is actually quite illuminating, because the elements that make up the cluster are not points but images, and images, unlike the points in a scatterplot, carry information independently of their spatial positions. So even if our image data is very high-dimensional, we can present the image clusters in a simple way, ignoring the vexed matter of their spatial positions—perhaps by so many square montages—-and we nonetheless achieve some meaningful characterization of the clusters.

Fig. 11. Growing Entourage with wide grid, resulting in relatively isolated circular clusters.

Fig. 12. Growing Entourage using same data as above, but with a tighter grid. Some clusters are isolated, some have clumped with neighbors.

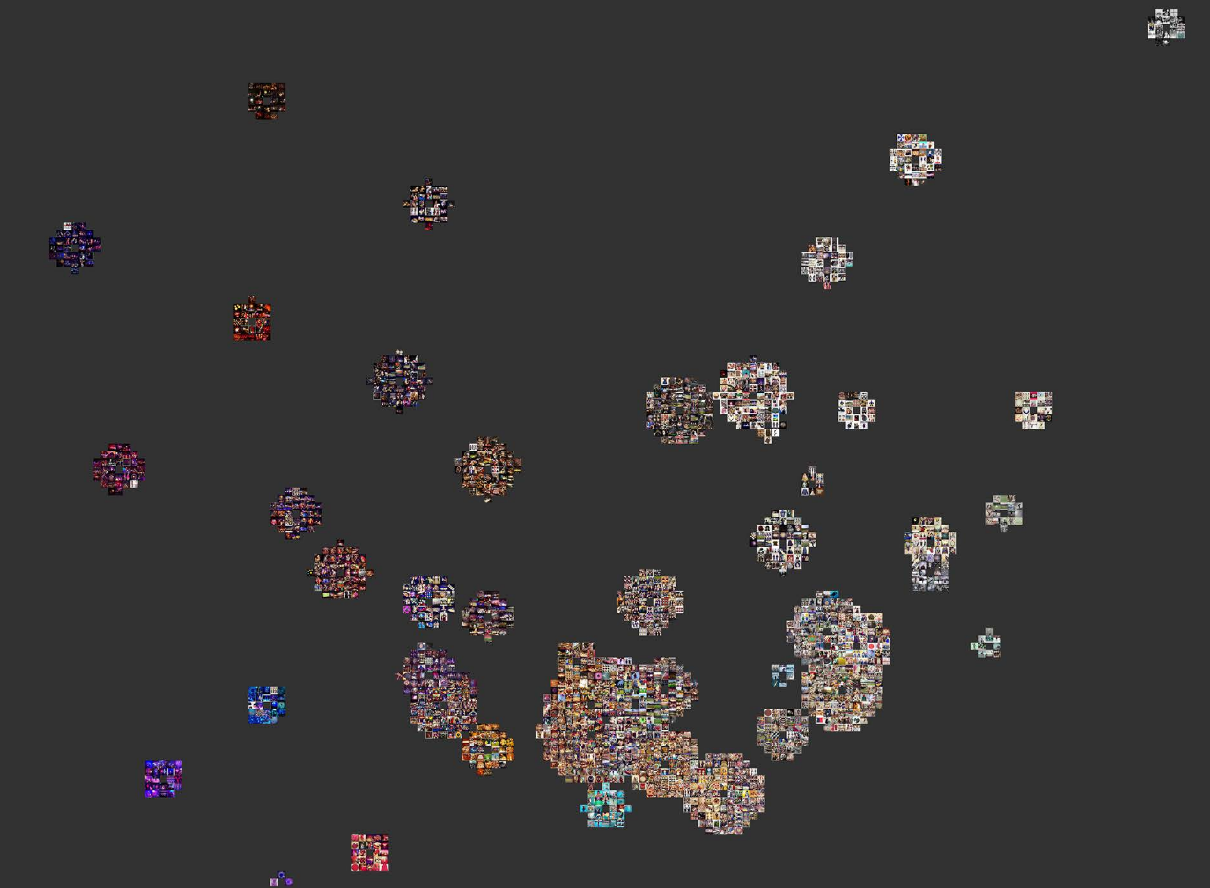

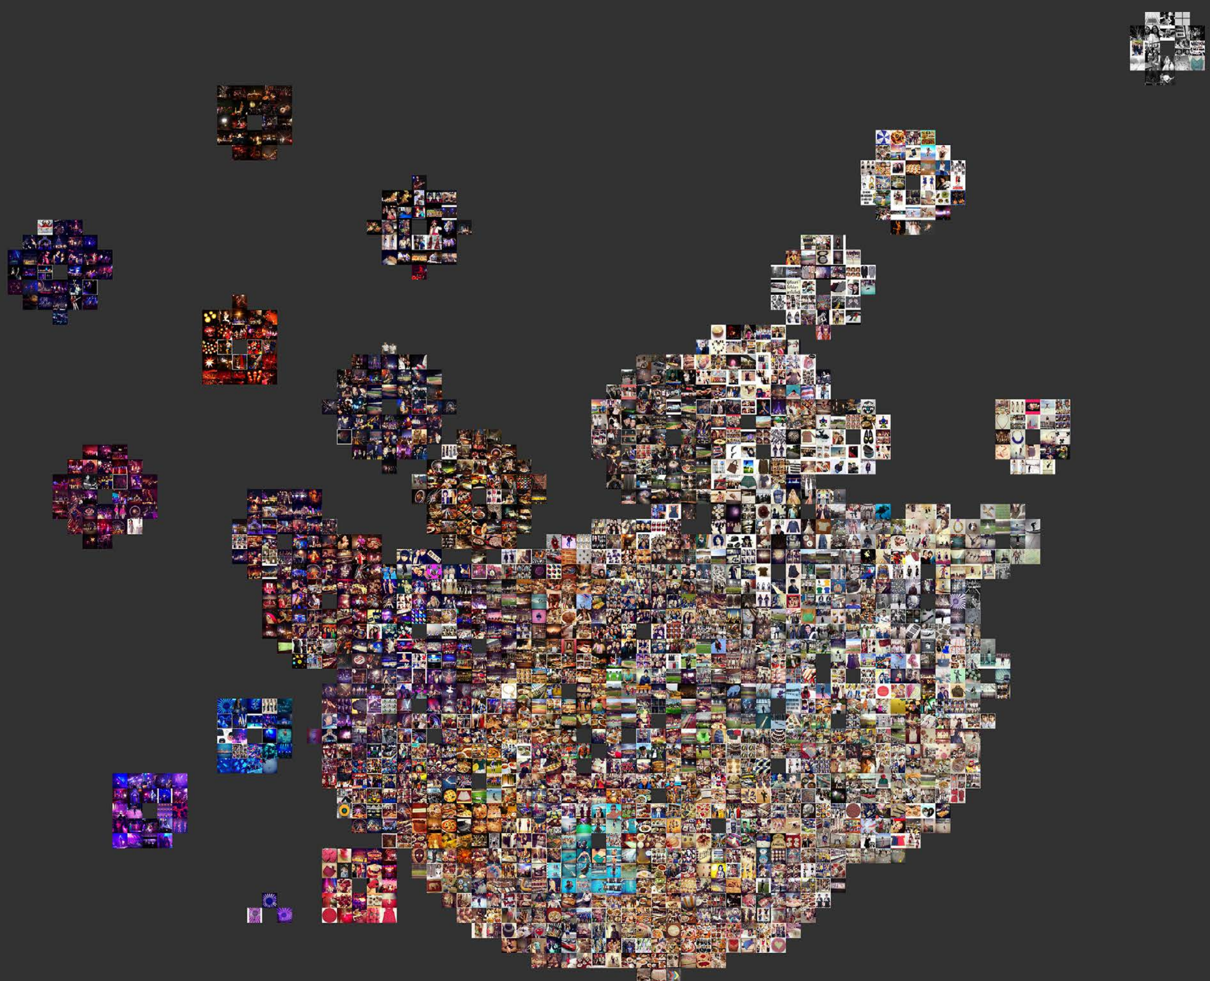

Growing Entourage Plot

The question now is how exactly to arrange these clusters on a 2D canvas. As I've said before, a simple presentation of each cluster is helpful. We could, for example, make square montages of each cluster and just leaf through them. But we might want more than this - we might want also to see the relations among the clusters. And now we confront a familiar problem: we have a set of data points in n-D, where n > 2, and we want to see them in 2D. The 'points' are now clusters, but the shape of the problem is exactly the same as before. The growing entourage plot is my solution to this problem. It projects cluster centroids—the middles of clusters—-to 2D and builds clusters around them by turn-taking and semantic priority.

We begin with high-dimensional image data and then use an algorithm like k-means to find k clusters in the original feature space. Each cluster has a centroid, given as a point location in the original feature space. We project the centroids to 2D using a dimensionality reduction algorithm, like PCA or t-SNE. We then bin these coordinates to a grid (making sure that no two centroids have the same grid location, which is typically not difficult). Now we have complex similarity relations among cluster centroids, and it remains only to build these clusters of images on the grid at the 2D centroid locations.

Every image in a given cluster is ranked according to its Euclidean distance[18] (in the original feature space) from the centroid. We can think of the centroid as the 'leader' of an 'entourage', and each image in the cluster is a member of the entourage. The closer they are to the centroid, by the aforementioned ranking, the closer they get to 'stand' near the centroid. Each cluster takes turns adding members of its 'entourage', starting with those closest to the leader. Each added member stands in the open grid space nearest its leader. Local conflicts between entourages are settled by this principle, since added members must occupy open grid squares.

This means that the look of the plot will depend on how we generate the original grid. We might end up with an array of circular clusters in 2D (Fig. 11), or we might end up with one large clump of images, with high-ranking members bunched up around their leaders and lower-ranking members scattered in nearby territories (Fig. 12).

This is not, of course, the only way to present clusters on a 2D canvas. It is, however, probably the best way to preserve as much of the complexity of intercluster relations as is possible in 2D. Additionally, it preserves similarity relations among images in the original feature space, something we lose by pure projection methods. Finally, it preserves intracluster relations by giving the semantically closest entourage members the privileged locations nearest their leaders.

Notes

1 This work is funded by a grant from the Frontiers of Innovation Scholars Program at UCSD. The code used to generate the visualizations described here can be found at https://github.com/damoncrockett/DataVisualizationTools/blob/master/learning.py (for the Image Histogram); and https://github.com/damoncrockett/SSI/blob/master/growing_entourage_plot/growing_entourage_general.ipynb (for the Growing Entourage Plot). These are not software packages but simple Python scripts. Please feel free to contact me at damoncrockett@gmail.com for implementation details.

2 I here use 'image plot' as a general term; it also refers to a software package developed by Lev Manovich and collaborators. See http://lab.softwarestudies.com/p/imageplot.html.

3 Jeremy Douglass and Lev Manovich, 2009. See http://www.flickr.com/photos/culturevis/4038907270/in/set-72157624959121129/.

4 William Huber, Lev Manovich, and Tara Zapel, 2010. See http://www.flickr.com/photos/culturevis/sets/72157623862293839/.

5 Lev Manovich and Jeremy Douglass, 2010. See http://www.flickr.com/photos/culturevis/4497385883/in/set-72157624959121129/.

6 Nadav Hochman, Lev Manovich, and Jay Chow, 2013. See http://phototrails.net/.

7 Content-based image retrieval systems in computer science often produce direct visualizations, typically involving small sets of images sorted by visual similarity. These systems are designed to support effective exploration and overview of image datasets (see Smeulders et al 2000 for review). The 'new paradigm' I point to here is meant to be considerably more extensive than this, although there is some clear overlap in our goals and methods.

8 See http://rsbweb.nih.gov/ij/plugins/image-montage/index.html.

9 See http://lab.softwarestudies.com/p/imageplot.html.

10 http://selfiecity.net.

11 For example, Selfie City: http://selfiecity.net.

12 For example, we've used deep learning image classification to analyze the contents of one million Twitter images (Yazdani and Manovich, 2015).

13 http://www.visionresearchwitch.com/.

14 Big thanks to the object detection team at Flickr for the object and scene tags. Their work is described here: http://code.flickr.net/2014/10/20/introducing-flickr-park-or-bird/.

15 For example, t-distributed stochastic neighbor embedding, or t-SNE (Van Der Maaten and Hinton 2008). Both PCA and t-SNE try to minimize the loss of relational information amongst points in high-dimensional space during projection to a lower-dimensional space. PCA does this by identifying axes along which the data points vary the most; t-SNE does this by privileging local relations between points over longer-range relationships.

16 See, e.g., https://www.flickr.com/photos/culturevis/albums/72157637904898314/with/10977282326/. The visualizations mentioned in note 7 are of this type.

17 K-means clustering partitions a set of data points into k clusters, for some choice of k. The algorithm attempts to minimize the intracluster distances between points.

18 Euclidean distance is the 'straight-line' distance between any two points in a space.

Bibliography

Manovich, L. (2011). What is visualization? Visual Studies, Vol. 26, no. 1, 36-49. Available at http://manovich.net/index.php/projects/what-is-visualization.

Manovich, L. (2011). Media visualization: Visual techniques for exploring large media collections. The International Encyclopedia of Media Studies Vol. VI: Media Studies Futures, edited by Kelly Gates (Chichester: Wiley Blackwell). Available at http://manovich.net/index.php/projects/media-visualization-visual-techniques-for-exploring-large-media-collections.

Smeulders, A. W., Worring, M., Santini, S., Gupta, A., & Jain, R. (2000). Content-based image retrieval at the end of the early years. IEEE Transactions on Transactions on Pattern Analysis and Machine Intelligence, 22(12), 1349-1380.

Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. Journal of Machine Learning Research, 9 (2579-2605), 85.

Yazdani, M., & Manovich, L. (2015). Predicting social trends from non-photographic images on Twitter, 2015 IEEE International Conference on Big Data, 1653-1660. Available at http://manovich.net/index.php/projects/predicting-social-trends

Fig. 13. Close-up of Fig. 1. Empty grid squares are centroid locations.

Damon Crockett is an information visualization researcher for the Software Studies Lab at the Qualcomm Institute, UCSD. He is a recent Ph.D. in Philosophy and Cognitive Science from UCSD. His doctoral dissertation is an extended argument about the information-carrying function of human visual perception. His current research concerns the role of human perceptual judgment in information visualization.

Correspondence: damoncrockett@gmail.com

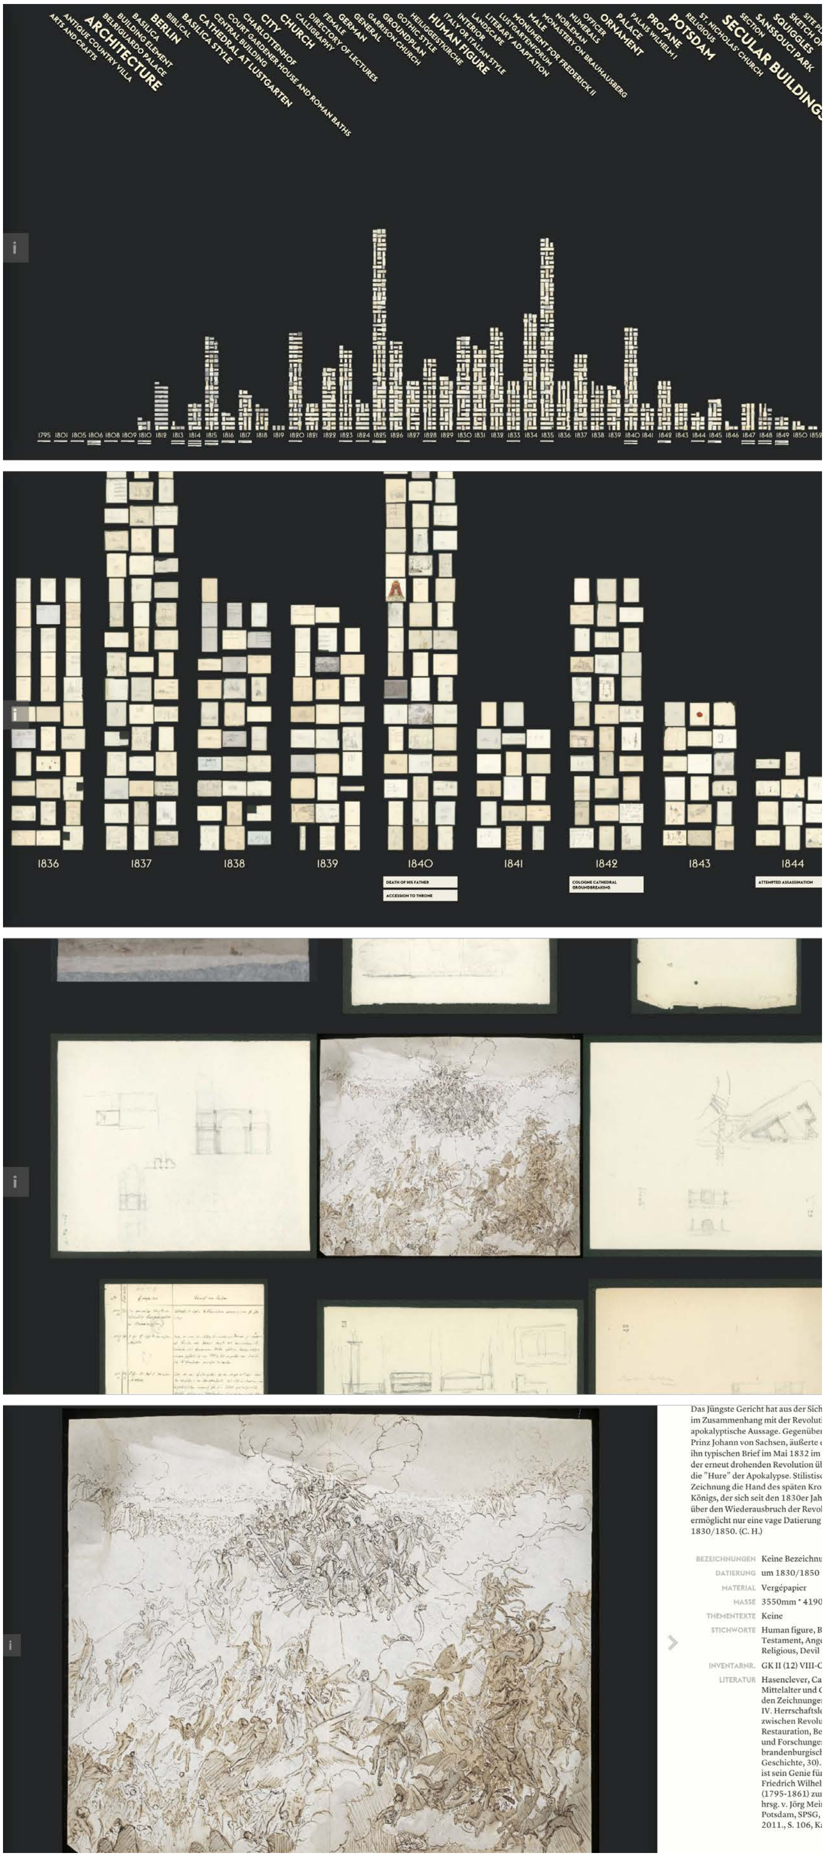

Figure 1: From top to bottom: Gradually zooming from an initial bird's-eye view of the visualization to a mid-point view showing several sketches in temporal proximity to close-up views of one individual drawing.

Invited Article

Linking structure, texture and context in a visualization of historical drawings by Frederick William IV (1795-1861)

Katrin Glinka, Christopher Pietsch, Carsten Dilba, Marian Dörk

Abstract: In this article we present a case study on digital representation of the art historical research and metadata brought together for a scientific collection catalogue by the Prussian Palaces and Gardens Foundation Berlin-Brandenburg. The resulting interface aims at linking the structure and texture of a collection of drawings by Frederick William IV of Prussia (1795–1861) with additional contextual information. The article describes the context of the larger research project and presents the resulting visualization and interaction techniques specifically designed for dynamic exploration along time and subjects.

Keywords: information visualization, metadata, zoomable user interface, direct visualization, case study, inventory catalogue, digital art history

Link: https://uclab.fh-potsdam.de/fw4/en

Introduction[1]

»Every idle moment [...] he draws on paper; sketches for great historic pieces, [...] persons and things that he has seen while travelling, mythical beings and allegorical matters. He even paints heaven and hell; and quite often biblical things.«

Johann Friedrich Herbart (1810)[2]

According to its definition, a museum does not only serve as a (semi-) public space for education, indulgement or even enjoyment. Its self-perception and objective is deeply rooted in an obligation to conduct research and ensure conservation[3]. Thus, institutional exhibition and publication activities target a broad range of different activities and audiences. For the non-scholarly public, the most visible has been so far the physical exhibition in a museum or gallery, often accompanied by an exhibition catalogue. Adding to that, the scholarly and expert public seeks to be informed on the research conducted in the specialised departments of a collecting institution on a more scientific level. Among these publishing activities that are aimed at an expert public are inventory or collection catalogues that provide domain experts and researchers with a fully developed art historical examination of a set of artefacts, which are often thematically focused on well-defined parts of a collection. The still predominant medium of publication for both non-scholarly exhibition catalogues and scientific collection catalogues is the printed book format. Museums and other collecting institutions have meanwhile understood that the digitization of their collections is an inevitable need in order to provide reproductions of objects and artworks for research, make them accessible via e.g. databases with web interfaces, and secure the conservation of sources and material. Correspondingly, the Prussian Palaces and Gardens Foundation Berlin-Brandenburg (SPSG) has started to employ digital forms of publication and is working on digitizing their collection. However, as is the case for many collecting institutions engaged with digitization efforts, it is still an open question how to make the newly digitized collections available for open exploration and visual analysis.

In this context, the objective of our overall research project "VIKUS–Visualizing Cultural Collections" is to examine the potentials of visualization techniques when applied to, and developed for, digitized cultural collections. Given the fact that the SPSG manages and administers several historical buildings, palaces, gardens, vast collections of paintings, furniture, sculptures, porcelain, drawings, and other historical objects (that have not all been entirely digitized, yet), the first step of our research project was to identify a suitable subset from the range of collections that could serve as the first case study. Accordingly, the existing digital resources had to be analyzed together with professionals from various areas of the foundation. We conducted a first co-creation workshop[4] in order to identify promising collection areas and aspects that could suit our aim to conceive a dynamic visualization using the existing digital sources and material provided by the SPSG. After these early stages of the research project, we decided to use a fully developed digital inventory catalogue of the drawings by Frederick William IV of Prussia. The decision was mainly influenced by the ambition to explore the potential of visualization as a tool that does not only allow for an overview (e.g. analytical visualizations of metadata[5]), but also serves an exploratory gateway to the collection by combining overview and detail (on a visual as well as textual level) while also integrating contextual observations and scientific findings. The inventory catalogue had already been published digitally, but in a static format, comprising high-resolution digital copies of 1492 sheets of drawings by the King alongside the corresponding metadata and a full art historical analysis, indexing, and interpretation in several object-related and thematic texts.

Aspirations

Being the first use case developed in the framework of the VIKUS research project, the visualization of drawings by Frederick William

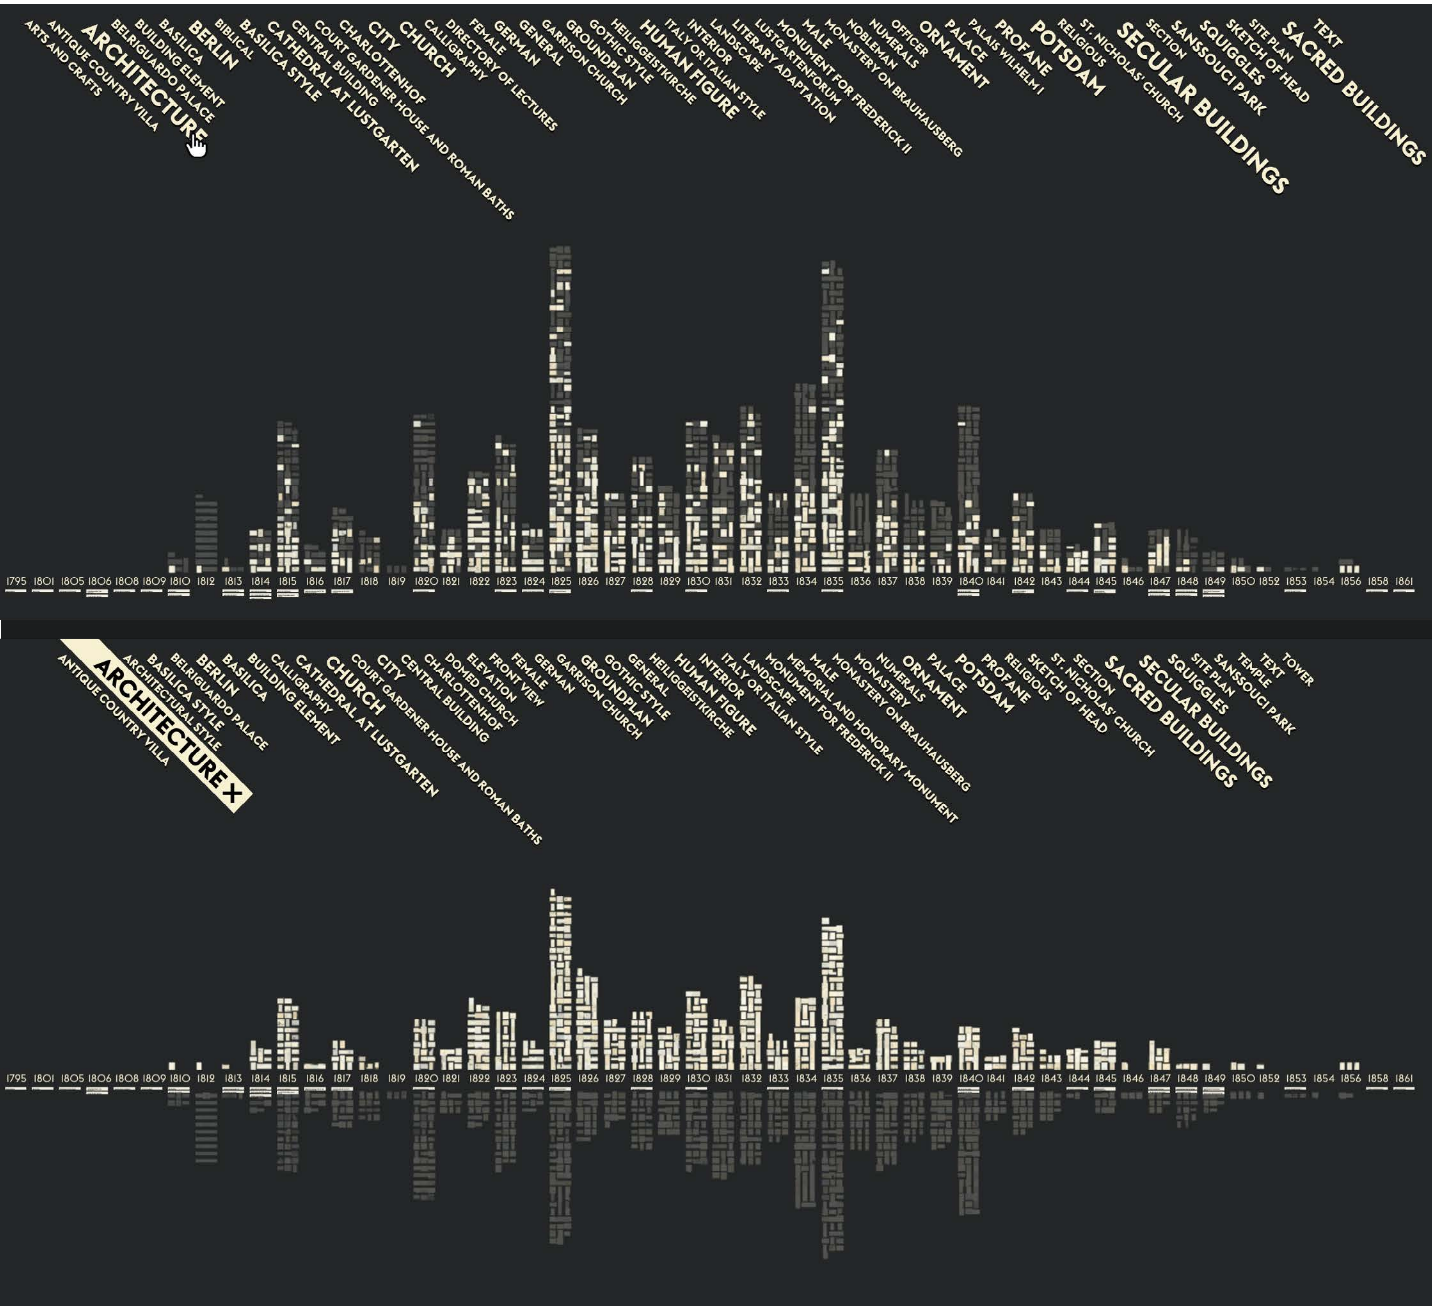

Figure 2: Hovering over a keyword (top) and selecting the keyword “Architecture” (bottom).

IV serves as an object of study that is, on a more general note, aimed at investigating the potentials inherent to visual exploration of digitized collections. In the scope of our research, we develop and evaluate graphical interfaces that are aimed at enabling interactive examination of cultural objects. We thereby practically contribute to the overall challenge that archives, museums, libraries and other collecting institutions are now facing, namely the need to improve the accessibility of their digitized inventories while also providing new modes of engagement with digitized artifacts. While conducting applied research and working in close collaboration with cultural institutions - in this case the SPSG - we also wish to engage in the ongoing discourses that are prevailing in disciplines like computer sciences and visualization research, (digital) humanities, (digital) art history, museology, as well as design research and adjoining disciplines.

This decidedly interdisciplinary approach offers the possibility to build upon and draw from previous disciplinary and interdisciplinary research. At the same time, it also poses the challenge to translate between different academic cultures, methodologies and assumptions. To put this in concrete terms, digitizing a collection does not implicate that this process is nested in a thoroughly digital approach. As Johanna Drucker has pointed out, there is a distinction between *digitized* art history and *digital* art history[6]. This in our view important distinction also served as a vantage point for Benjamin Zweig in his "attempt to help define [...] what »digital art history« is."[7] Reflecting the course of the project, we were challenged to transfer or rather translate a *digitized* art historical source into a *digital* art historical source. More specifically said, although the source for our visualization already is a digitally published inventory catalogue, its internal logic, framework, and structure still resembles non-digital ways of working and researching in art history. Thus, we had to identify properties, metadata, and structures together with our project partner, and in turn communicate the logic and functionality of data-driven visualizations. Our self-proclaimed goal consequently led to interdisciplinary co-creation, continuous integration and discussion of approaches and priorities. Another feature that developed throughout the process was the aim to widen the audience that might engage with the collection of drawings.

As pointed out earlier (and emphasized in interviews with our cooperation partner), the inventory catalogue is mainly published for a scientific expert public and is distinctly conceived within the research department (Department Palaces and Collections) of the foundation. The King's drawings are predominantly consulted by researchers in the domain of art history and architecture as well as by experts on Prussian history. Only on the occasion of the King's 150th day of death in 2011, the drawings were made publicly accessible in an exhibition at the Roman Baths in Sanssouci Park.[8] With the visualization, we built upon the scientific content of the inventory catalogue and enhanced it with contextual information within the visualization, aiming at making the collection of drawings explorable also by a non-expert audience. We are investigating if and how novel visualizations of cultural collections and the corresponding art historical research can be employed to also encourage innovative ideas in the field of cultural mediation and be used to communicate scientific findings to a broader public.

Drawing on data

Frederick William IV of Prussia (1795 – 1861) left a collection of drawings behind. They bear witness to historical events such as wars and revolutions, literary influences or personal obsessions with the devil. Numerous sheets reveal the planning eye of the King in the form of architectural visions and dreamy drafts. So far, 1492 sheets of drawings produced by the King have been fully accessed. The existing inventory catalogue that was published by the SPSG online in 2013 comprised the high-resolution digital copies of the drawings and sketches and corresponding metadata. Not all, but most of the following fields were available for each sheet: an image id, an inventory number derived from the internally used database, a description of the sheet, an art historical commentary or interpretation, labelling (e.g. hand-written notations by the King or markings by later researchers), watermarks, measurements of the sheet, time and year, material, title(s) of corresponding thematic text(s), list of corresponding secondary literature, and a hierarchical index-based list of descriptors. The descriptors were worked out by the group of art historians during the analysis and research process alongside a controlled vocabulary and were developed in two parallel strings: a thematic order and a topographical order.[9] In accordance with our aspiration to make the art historical sources explorable by a non-expert public, we also included the content of the exhibition catalogue10 into our choice of data. The texts from the exhibition catalogue – in contrast to the object description and interpretation – were written and published for a broader audience and offered background information on the King's life, prevailing interests, and historical context. Although our aim was to only use existing digitized material and leave the art historical content and structure of the data intact, we decided to manually extract the information from the exhibition catalogue and request additional information supplied by historians of the SPSG in order to create a custom timeline structure. Nonetheless, the extent of work that had to be put into the manual gathering and structuring of the additional timeline data was still moderate. Hence, for the most part we only used the existing data derived from the art historical research and identified those properties that could be used to offer dynamic insights. By focusing and limiting the data mainly to existing parameters and content, we wish to illustrate how *digitized* art historical research can be reused and adapted when developing a *digital* and data-driven exploratory gateway.

Interaction and Design

The visualization is conceived as a dynamic canvas arranging the King's drawings by their creation year, linked with contextual information, and made accessible through interactive filtering and zooming capabilities. On the one hand, the interface mechanics are inspired by zoomable interfaces11 and more recent applications to cultural collections.[12] On the other hand, the design is based on the recurring wish from our collaborators to be able to see and explore the collection along temporal and thematic aspects while not abstracting the individual drawings into aggregated shapes. Thus, the overall aim for the interface is to reveal structures of temporal and topical distributions in the collection that invite the viewer to explore the collection and provide seamless access to the rich textures of individual drawings in high resolution. The interface has three main parts: 1) the zoomable canvas containing the scans of the drawings, 2) the index-based list descriptors or "keywords", representing the main subjects and places associated with the drawings, and 3) contextual information for each drawing and several time periods. In the following we briefly describe these three parts.

The drawings are positioned according to their year of creation in horizontally arranged columns. As some sketches do not have specific dates or years, but rather estimated time ranges, the median of their estimated range is used for the positioning. Within the columns, sketches are sorted vertically based on their complexity, i.e., the number of motives, themes, and places that are associated with each sketch. When the interface is launched (see Figure 1, top), the initial view offers a bird's-eye view on the complete set of sketches with each of the 1492 images being displayed at a relatively small size of about 20 mm on a 13" laptop display. Similar to the image plots promoted by Lev Manovich and colleagues,[13] it is almost indistinguishable shades and shapes that can be differentiated from this perspective. In order to gain a better sense of the actual artifacts, it is necessary to move closer to the individual items. For this purpose, the canvas serves as a continuously zoomable space, allowing for the gradual increase of detail for particular segments of the arrangement of images. Zoom operations can be carried out either through the mouse wheel or by performing scrolling or zoom gestures on touchpads and touch-enabled displays. By zooming into particular groups of images, the thumbnails continuously grow into larger images with a higher resolution. By clicking and dragging, it is possible to pan the

Figure 3: The timeline provides historical context at varying levels of detail.

canvas in all directions. The viewer can examine sketches in the same year by performing vertical movements and shift in time between different year columns by moving horizontally. Besides zooming and panning, it is also possible to select individual sketches (by clicking or tapping on them) in order to immediately focus on this particular sketch and display it in high resolution. Once an image is in focus, the neighboring images fade out and contextual information becomes available on the right side of the display. In the same visual context of the canvas, it is now possible to further zoom into the image of the drawing to closely inspect the texture of the paper and even the grain of the pencil strokes (see Figure 1, bottom).

In the top part of the interface, the most prolific subjects and places are displayed as a horizontal list of keywords that are derived from the index-based list of descriptors. The list contains between 30 and 50 words and phrases, depending on the size of the screen. Akin to word clouds, the font sizes represent the relative frequency of sketches per subject. The keywords are sorted alphabetically to allow viewers to quickly locate a subject they might be looking for. Hovering over a keyword highlights all the sketches that are associated with this subject (see Figure 3, top). Ironically, the resulting aesthetic does remind of urban skylines, however, not those envisioned by Frederick William IV in Prussian times. When selecting a keyword, all images not associated with this word move to the bottom of the timeline and are displayed with a lower opacity. The drawings that match the keyword selection remain above the timeline. As more and more tags are clicked, the selection of drawings gets more constrained, resulting in smaller image columns. Clicking on a selected keyword again cancels its selection. Changing the selection of keywords also changes the display of the remaining keywords. As some keywords may not be associated with the drawings in a given selection, they will be hidden, making space for more related keywords to be revealed, thus increasing the level of accuracy or specificity in the descriptors' content.

In addition to the general contextualization of the drawings provided by their temporal arrangement and the keyword visualization, the interface features two additional levels of in-depth information about the collection. On the one hand, a zoomable timeline just below the years presents biographical and historical information in the temporal context of the drawings. On the other hand, the metadata, descriptions, and art historical interpretation for each drawing contained in the collection is displayed in the single-image view in a text panel. In addition, the information panel contains links to in-depth articles on specific topics pertinent to the collection. The detailed events in the timeline are shown as soon as the user zooms into a particular period (see Figure 3). The timeline contains 40 events related to Frederick William

IV's life as well as historical events related to the political developments of the time. These events are positioned right below the yearly columns in order to facilitate the establishment of a connection between the personal and historical developments and the sketches that the King produced at the time.

Technological implementation

In accordance with our aim to make the collection of drawings accessible and explorable by a broad audience, it was paramount for us to publish the visualization in a web interface. Given the size of the collection and its digital images, we had to develop an approach that would allow us to display the whole of the collection on a zoomable canvas, while at the same time be highly responsive and reduce the loading time. In order to be able to reconcile these objectives we developed new technological approaches. Knowing that waiting time while loading a website is crucial to the perceived quality of the user experience[14] it was one of our pronounced goals to instantly be able to not only display the first zoom level of the visualization, but to also allow for dynamic interaction right from the start. For this first state of display (the bird's-eye view) we employed a progressive loading approach which instantly loads a data layer with low level of detail. This first layer includes the metadata and thumbnail of each sheet. This preliminary data is visualized in form of the list of keywords, the image plot, and the timeline. While the user is able to gain a first overview and filter the visualization by keywords, a second data layer is being transferred in the background, which holds a high-resolution version of each image. Each loading progress is communicated by a progress bar, which has been identified as being a good approach for reducing the perceived waiting time for the user.[15] As soon as the thumbnail resolution is exceeded during the zoom interaction, the next level of detail is displayed with a version of the image with higher resolution. In order to be able to provide a fluid user experience during the zooming and reloading of images, a fast and efficient way of streaming the data was built. Conventionally, a website loads each image separately in a TCP request, which would in our case inevitably flood the user's connection and browser. To bypass this limitation, we clustered the images into chunks of data blocks. These data blocks have an approximate size of 7mb and hold up to 150 images. Those chunks are then streamed with a GZIP compression from a Content Delivery Network, which provides the fastest transfer speed, to the client's browser. After a chunk is received by the browser, the contained images are then pumped into the visualization as a detailed version of each entry. The data chunks are kept under a size of 10mb, which is the file size limit of cacheable objects in today's browsers. In order to bring the desired visual features together, we combined the web visualization library D3.js with the web graphics library pixi.

js. The list of index keywords that are used for filtering the visualization, as well as the timeline with the additional contextual annotations, are rendered in HTML with CSS3 animations and the area holding the images is rendered in a WebGL canvas. Consequently, we were able to combine the benefits of HTML5 features, such as interactive user interfaces and high performance CSS animations, with the performance of modern GPUs via WebGL.

Qualitative insights from quantitative views

The collection subset used for the case study consists of approx. 1500 data records, or about 20% of the drawings executed by Frederick William IV which are preserved in the collections of the Prussian Palaces and Gardens Foundation. The first version of the online catalogue, launched by the SPSG in 2013, did not cover all subjects evident in the King's drawings to the same extend. Instead, specific compilations – e.g. drawings of Charlottenhof Palace within Park Sanssouci – were published almost entirely, while other series and subjects were only represented by singular sheets. In reference to the significance and gain of knowledge expected from the actual visualization, one must take into account that its data set is basically identical with the 2013 catalogue and will only successively be complemented with sheets missing so far. The temporal arrangement and thematic filter functionality, complemented by supplementary biographical and historical information, creates highly interesting curve progressions, showing which subjects were most relevant for the king in a specific time or maybe readopted in later phases. A most significant example for this observation is provided by the sheets covering Charlottenhof Palace. A present to Frederick William IV by his father in 1825, Charlottenhof was supposed to be complemented with new structures and become transformed into a mediterranean villa. The curve progression along the timeline clearly shows the intensity with which the crown prince worked on his later executed designs. Frederick William IV reacted quite similarly to other biographical or political events. His intensive planning on the restructuring of his study and his apartment within the Berlin city palace has to be seen in close regard to his marriage, celebrated only one year before. And Frederick William's constant changes, improvements and alternative proposals for the construction of Babelsberg Palace, carried out for his brother in several phases, will also impressively become visible along the timeline. However, this requires the recognition of all 6,900 records to obtain exact results. Then it will be possible to see how the revolutionary events of the year 1848 left deep scars on the psyche of the King, for example visible in the frequent depictions of demons and devils within his drawings of the time. His journeys to Italy will probably become just as evident on the timeline, showing how the king processed his impressions. It is especially here, that highly interesting insights may be expected for both the expert user and the interested layperson. Especially art and architecture historians will appreciate the ability to combine different index terms with each other. Presumably, it will be possible to obtain a more precise dating when combining findings on individual drawings executed on the same sheet. For example, some of the king's drawings depict the German piked helmet, only used by the Prussian Army since 1843. However, methodically it has to be remarked that the arrangement of the drawings along a timeline already represents the result of an art historical approach, taking into account the individual fixed datings and other biographical and historical corner points. Researchers will therefore have to be aware of circular reasoning; and yet, this new visualization may provide an effective and powerful tool to refine and verify previous findings, datings and examinations.

Conclusion

With this first use case we were able to apply principles and concepts from visualisation research and interface design to a digitized art historical source and illustrated the potential of visualizations when implemented in an art historical context. The concept and functionalities were developed in an interdisciplinary co-creation process, thus ensuring the result to be well grounded on the respective disciplines. During the next phase of our project, we will carry out thorough user-tests and additional empirical research in order to be able to validate some of the functionalities and investigate the potential of our model when used for art historical research as well as the applicability for a broader audience. Regarding the technology, a new framework for displaying collections of large image data has been successfully developed, opening up the potential to scale the number of images up to approximately 8000 images per canvas. The ability to fluidly combine a distant view with a detailed examination of the high-resolution images has already raised considerable interest among collecting institutions. We are now considering generalizing this specific visualization into a framework that may be applicable to other visual collections with e.g. a broader variety of visual qualities and higher contrast between the shapes and shades of images.

Notes

1 Acknowledgements: We wish to thank the Editors of the *International Journal for Digital Art History* for the invitation to present our case study. We would like to acknowledge the German Federal Ministry of Education and Research (BMBF) for their generous funding of our research project *VIKUS–Visualising Cultural Collections*. We thank our project partners SPSG and Programmfabrik GmbH for the productive cooperation. We wish to thank Matthias Graf for his work on a previous version of the visualization and Dr. Jörg Meiner for his scientific consulting. We are grateful

for valuable ideas, continuous feedback, and support from our colleagues Sebastian Meier, Till Nagel, Stephanie Neumann und Jan-Erik Stange.

2 Herbartische Reliquien, ein Supplement zu Herbart's Sämmtlichen Werken, eds. v. [Tuiskon] Ziller, Leipzig 1871, p. 201. Cited after »Unglaublich ist sein Genie fürs Zeichnen« Friedrich Wilhelm IV. von Preußen (1795-1861) zum 150. Todestag, published on behalf of the Prussian Palaces and Gardens Foundation Berlin-Brandenburg by Jörg Meiner 2011, 8. (Translation by the authors)

3 cf. ICOM Code of Ethics for Museums, 2013

4 For a description of the framework for such a participatory workshop with collection maintainers and contributors cf. Chen, K., Dörk, M., & Dade-Robertson, M. (2014). Exploring the Promises and Potentials of Visual Archive Interfaces. In iConference 2014 Proceedings (735–741). doi:10.9776/14348

5 see e.g. the visualization of metadata of the DDB that provides experimental overviews of the temporal and spatial distribution of objects and the associated topics, people, and organizations, using the metadata of more than 7 million cultural heritage objects aggregated by the German Digital Library (DDB) Bernhardt et al (2014).: http://infovis.fh-potsdam.de/ddb/index_en.html.

6 Johanna Drucker, "Is There a "Digital" Art History?" In: *Visual Resources: An International Journal of Documentation*, Vol. 29, No. 1-2 (2013), 7.

7 Benjamin Zweig, "Forgotten Genealogies: Brief Reflections on the History of Digital Art History." In: *International Journal for Digital Art History*, Issue 1 (2015), 39.

8 The exhibition »Friedrich Wilhelm IV. von Preußen (1795-1861) zum 150. Todestag« was open to the public from 7 May until 31 July, 2011 and was accompanied by an exhibition catalogue.

9 The digital copies, metadata, and texts used in the visualization were created by Dr. Jörg Meiner (project lead), Dr. Catharina Hasenclever (project lead 2006-2008), Antje Adler, Astrid Fritsche, Klaus Dorst, Stefan Gehlen, Dr. Gabriele Horn, Dr. Andreas Meinecke, Dr. Gerd-H. Zuchold, Dr. Rolf H. Johannsen, Prof. Dr. Harry Falk, Dr. Rolf Th. Senn and Dr. Sepp-Gustav Gröschel. Editor: Dr. Carsten Dilba.

10 »Unglaublich ist sein Genie fürs Zeichnen« Friedrich Wilhelm IV. von Preußen (1795-1861) zum 150. Todestag, published on behalf of the Prussian Palaces and Gardens Foundation Berlin-Brandenburg by Jörg Meiner (2011)

11 Bederson, B. and Hollan, J. (1994). Pad++: a zooming graphical interface for exploring alternate interface physics. In UIST 1994: Symposium on User Interface Software and Technology, 17–26. ACM.

12 Hochman, N. and Manovich, L. (2013). Zooming into an instagram city: Reading the local through social media. First Monday, 18(7).

13 Manovich, L. (2015). Data science and digital art history. International Journal for Digital Art History, (1): 13–35.

14 cf. eg. Antonides et al. "Consumer Perception and Evaluation of Waiting Time: A Field Experiment". In: Journal of Consumer Psychology, 12(3) (2001); Egger et al. "Waiting times in quality of experience for web based services", *QoMEX*, IEEE (2012)

15 cf. Myers "The importance of percent-done progress indicators for computer-human interfaces". Presented at the Computer Human Interactions, Vol. 16, ACM (1985)

Bibliography

Antonides, G., Verhoef, P. C., & van Aalst, M. "Consumer Perception and Evaluation of Waiting Time: A Field Experiment". In: *Journal of Consumer Psychology*, 12(3) (2001), 193–202. doi:10.1207/S15327663JCP1203_02

Bederson, B. and Hollan, J. "Pad++: a zooming graphical interface for exploring alternate interface physics." In: *UIST 1994: Symposium on User Interface Software and Technology.* ACM (1994), 17–26.

Bernhardt, C., Credico, G., Pietsch, C., Dörk, M. Deutsche Digitale Bibliothek visualisiert, http://infovis.fh-potsdam.de/ddb/ (2014) (accessed March 3, 2016).
Chen, K., Dörk, M., & Dade-Robertson, M. "Exploring the Promises and Potentials of Visual Archive Interfaces". In: *iConference 2014 Proceedings.* doi:10.9776/14348 (2014), 735–741.
Drucker, Johanna. "Is There a "Digital" Art History?" In: Visual Resources: An International Journal of Documentation 29, No. 1-2 (2013), 5–13.
Egger, S., Hossfeld, T., Schatz, R., & Fiedler, M. "Waiting times in quality of experience for web based services". Presented at the 2012 Fourth International Workshop on Quality of Multimedia Experience (QoMEX), IEEE (2012), 86–96. doi:10.1109/QoMEX.2012.6263888
Hochman, N. and Manovich, L. "Zooming into an instagram city: Reading the local through social media". In: *First Monday*, 18(7) (2013).
Manovich, L. "Data science and digital art history". In: *International Journal for Digital Art History*, Issue 1 (2015),13–35. DOI: http://dx.doi.org/10.11588/dah.2015.1.21631
Meiner, Jörg (ed. on behalf of the Prussian Palaces and Gardens Foundation Berlin-Brandenburg). »*Unglaublich ist sein Genie fürs Zeichnen*« Friedrich Wilhelm IV. *von Preußen (1795-1861) zum 150. Todestag.* Potsdam, 2011.
Myers, B. A. "The importance of percent-done progress indicators for computer-human interfaces". Presented at the Computer Human Interactions, Vol. 16, ACM (1985), 11–17. doi:10.1145/1165385.317459
Ziller, T. (ed.). *Herbartische Reliquien: Ein Supplement zu Herbart's Sämmtlichen Werken.* Leipzig, 1871
Zweig, Benjamin. "Forgotten Genealogies: Brief Reflections on the History of Digital Art History." In: *International Journal for Digital Art History*, Issue 1 (2015), 38–49. DOI: http://dx.doi.org/10.11588/dah.2015.1.21633

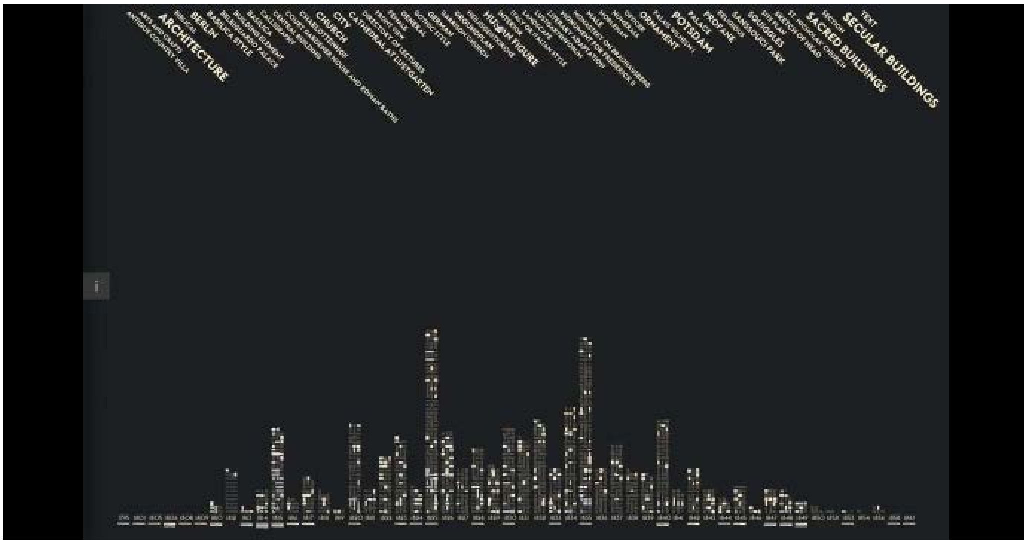

The research project »VIKUS - Visualizing Cultural Collections« (Visualisierung kultureller Sammlungen, 2014 - 2017) brings together researchers from various fields such as information visualization, computer science, interface design, and the humanities to develop and evaluate graphical interfaces aimed at enabling interactive examination of cultural objects. The project liaises closely with cultural institutions (e.g., museums, libraries and foundations) as well as with developers of media databases. The VIKUS team approaches the area of digital cultural heritage by combining technological possibilities with cultural considerations in order to develop visualizations and interfaces that open up interesting and useful perspectives on digitized collections. Thereby, novel interaction techniques and representations are designed and evaluated for their suitability for different scenarios.

More information about the Urban Complexity Lab: uclab.fh-potsdam.de

Katrin Glinka is a research associate and lecturer at the Potsdam University of Applied Sciences. She combines approaches from art history, sociology and museum studies with an interest in digital cultural heritage and visualization research. She studied cultural sciences with a focus on art theory, visual culture, sociology and philosophy and holds an M. A. degree from Leuphana University Lüneburg. Since 2014 she has been working on her doctoral thesis on digitization and visualization in the cultural field and their means and potentials for curation, critical and interventionist approaches and visitor orientation in museums.

Correspondence e-mail: glinka@fh-potsdam.de

Christopher Pietsch is a Berlin-based interaction designer. His passion for interfaces goes beyond the digital layer as he tries to connect the physical with the digital world. He studied computer science at the HTW Berlin and interaction design at the University of Applied Sciences Potsdam and holds a bachelor's degree for his thesis on brain-computer-interfaces. As a freelance information and interaction designer he now explores novel types of visualization metaphors (www.chrispie.com).

Correspondence e-mail: cpietsch@gmail.com

Carsten Dilba, Ph. D., is scientific editor at the Prussian Palaces and Gardens Foundation Berlin-Brandenburg, responsible for both printed an online publications. He studied at the Universities of Bonn, Leicester, Vienna and the University College London, subsequently working for the Fraunhofer IAIS (netzspannung.org). He published on mediaeval and 18th century art history.

Correspondence e-mail: c.dilba@spsg.de

Marian Dörk is a research professor for Information Visualization at the Institute for Urban Futures of the Potsdam University of Applied Sciences. During his PhD at University of Calgary and his postdoctorate at Newcastle University he designed and studied novel visualization techniques in particular with regard to their potential for exploratory information practices. He leads a 3-year research project on visualizing cultural collections (VIKUS) and since January 2015 he has been co-directing the Urban Complexity Lab, a research space at the intersection between information visualization and urban transformation.

Correspondence e-mail: doerk@fh-potsdam.de

Workshops

Computing Art: A Summer School for Digital Art History

Peter Bell
Heidelberg Collaboratory for Image Processing (HCI)

The summer school for digital art history[1] in 2015 was organized by the Computer Vision Group of the Interdisciplinary Center for Scientific Computing (IWR) in Heidelberg with generous support by the HGS MathComp, the Heidelberg Academy of Science and Humanities and German working group "Digital Art History".[2]

Digital art history lies at the intersection of art history and the digital humanities and benefits from digital methods, thus requiring a truly interdisciplinary approach. In short, this means using digital infrastructures and tools, but also critically reflecting upon their usage and understanding the technical background (quoted from "Memorandum zur digitalen Kunstgeschichte in der Lehre").[3] Digital art history has to mediate its research practice between the original objects and increasingly mimetic digital representations. It also has to embed these representations in the semantic and stylistic contexts of the artwork. In order to overcome these challenges, immediate access to the visual information of the digital representation is required, as well as access to virtual research environments that will help to depict semantic relations.

Starting with the digital representation of an artwork, the summer school demonstrated ways of processing the artwork visually, iconographically and contextually. The presentations focused on new methods of annotation and image analysis. Up until now nearly all steps of digital image investigation were left up to various experts: content processing was the task of highly qualified art historians supported by student assistants, whereas the creation of new digital records only concerned IT specialists or ambitious, self-educated scholars. New research strategies, such as crowd sourcing, machine learning and computer vision, provide the specialist and the layperson with algorithms/ artificial intelligence that enables effective mass processing of image data. In addition, new collaborations and interdisciplinary links are forming between the classical scientific fields. The digital image is not only described, but also analyzed in terms of content and compared with others, utilizing computational support. These processes must be further developed and critically monitored. The aim of the summer school was to introduce young scholars (master's students and PhD candidates) to the field of digital art history, in its full breadth, as represented by experts from computer science and art history. There was time for detailed talks, live demonstrations and extensive discussions.

After an introduction by the event's hosts, Georg Schelbert (Berlin), a representative of the working group "Digital Art History", presented the group and its aims and gave a short introduction of the topic. First, Björn Ommer and Peter Bell (Heidelberg) introduced various strategies of automatic image understanding and showed several approaches to computer vision for image analysis and image comparison. This included a theoretical overview in the shared methodical framework, e.g. Gestalt theory, and a practical introduction to their newly developed prototype for image search, where specific objects in different Sachsenspiegel-manuscripts where retrieved.

The second day began with a discussion of the obstacle of image rights. The many different modes of image usage has lead to great insecurity regarding permissions, which often hinders the development of digital humanity projects in art history. Consequently, Lisa Dieckmann (Cologne) who made the introduction and Maria Effinger (Heidelberg), as well as the lawyer Felix M. Michl (Heidelberg), led an open discussion image rights versus open content. Dieckmann and Michl discussed, for example, the pending cases of the Prometheus Bildarchiv[4] and the Reiß Engelhorn Museums. The conclusion of this session was that art historians should organize a change of copyrights to encourage the research and publication of art.

Clemens Schefels (Munich) and Daniel Kondermann (Heidelberg) went on to discuss the potentials and limits of crowd sourcing and presented applications for the segmentation and indexing of large records. Schefels showed that the project Artigo[5] is still exploring new ideas in the analysis of data and new gamification instruments. Kondermann showed how spam and incorrect annotation can be reduced by crowd driven quality management inside the Pallas Ludens system. Both talks explored a variety of options to the traditional workflow. Critical remarks in the discussion showed that art historians still need time before they can trust crowdsourcing.

Independent of the method used for capturing image information, a user-friendly, innovative, semantically highly cross-linked and sufficiently standardized database is necessary. Thorsten Wübbena (Frankfurt a. M./ Paris) and Thomas Hänsli (Zürich) gave a report on current developments and delivered some insight into some of their own projects like the Swiss database of the Archives of CIAM powered by the ConedaKOR system from Frankfurt. A very vivid discussion arose afterwards on the topic of standardization. The speakers believe in the use of identification-numbers for artworks like the GND-standard, whereas some participants criticize this as a positivistic fallback.

The subsequent workshops gave participants the opportunity to explore subject areas in greater detail: Piotr Kuroczyński (Darmstadt/Marburg) and Jan-Eric Lutteroth (München) from the digital 3D reconstruction research group introduced the topic of 3D reconstruction, which is especially relevant for architectural history. In the context of the digital image in art history, 3D reconstruction was only briefly mentioned, however this was already appreciated by many of the young scholars as a welcome point of orientation.

At the same time, Holger Simon (Cologne) presented occupational areas and job markets, which focus on expertise in digital art history. Museums and other cultural institutions are in increasing need of digital applications and virtual/augmented reality. Simon described new concepts and underlined these with examples from social media, websites and merchandizing of international museums. Jens-Martin Loebel (Bayreuth) and Heinz-Günter Kuper (Berlin) presented options for image annotation and image cross-linking with multimedia research environments using their own application HyperImage and simple XML-examples. This talk was extended by Matthias Arnold and Violetta Jantzen (Heidelberg) from the Cluster Asia and Europe in a Global Context, who presented HyperImage results in the form of huge Japanese picture scrolls, in which every object is segmented and annotated.

At the end of the event, participants gathered for a final discussion, which was followed up by a meeting of the research group digital art history. It is important to note that the summer school was not just centered on the invited talks. A core component was also the questions and contributions of students who were highly motivated enriched the discussions with their own experience in digital art history and even some voluntary posters. We hope that the *Computing Art Summer School* can be a pilot for a sequence of following, more specialized events and that it helps to connect the growing community of digital art history.

Notes

1 https://hci.iwr.uni-heidelberg.de/CompVis_Summerschool2015
2 http://www.digitale-kunstgeschichte.de/wiki/Hauptseite
3 http://www.digitale-kunstgeschichte.de/wiki/Erkl%C3%A4rung_zur_Digitalen_Kunstgeschichte_in_der_Lehre
4 http://prometheus-bildarchiv.de/
5 http://www.artigo.org/

Graduate Workshop on Digital Tools for Art Historians: The Visualizing Venice Summer Program "The Biennale and the City"

Caroline Bruzelius for the Wired! group
Duke University

The Wired! group at Duke University[1] was created in 2009 with a commitment to teaching and training scholars and students with the digital visualization and storytelling technologies appropriate for the study of art, architecture, and urbanism. Our primary goal has been generate baseline competence and digital confidence, building an understanding of the scholarly, teaching and public outreach capacities of digital tools, and enabling art historians to engage in productive collaborations and conversations with other disciplines, including computer scientists and engineers. We envision and work towards a future where scholars in Art History and Visual Studies can not only create maps, models, and displays that narrate research questions, but also have enough knowledge to join with urban history and art museums in developing interactive displays that can engage the public in a more meaningful experience.

We realized almost immediately in 2009, however, that most Art History Departments are neither equipped to support technical training, nor able (or even willing) to integrate emerging technologies with teaching and research. Most humanist scholars in universities, research centers or museums have difficulty identifying which technologies are most appropriate for art historical questions; there is little access to expertise or to laboratories with the requisite software. To address this challenge, and with the fervor of recent converts, we started offering two-week summer workshops in 2009 that present participants with a "buffet menu" of low-cost or free and open-source digital tools and a forum for discussing each tool's relevance for the practice of art and architectural historical research. The birth of Visualizing Venice in 2010,[2] initiated a collaboration with colleagues at the Architectural University of Venice (IUAV), the Engineering Department at the University of Padua, and Venice International University (VIU) to create a digital laboratory for on-site training that focuses on the history of art, architecture and urbanism in Venice. Workshop topics have included the cistern system (2012), the Ghetto (2013), The City and the Lagoon (2014), and The Biennale (2015). The participants integrate the collection of on-site research data (for example, scanning, 3D modeling, geo-referencing and archival information) into multimedia presentations. It has always been an important principle of our approach to show how multiple technologies can be combined (for example, video editing with photogrammetric 3D models) to address research questions and generate effective presentations for public outreach.

One of the most significant outcomes of our on-going work has been the realization that modeling and mapping can engage art historians in broadly-framed studies of systems that exist in dynamic relationships to each other. The type of new questions that emerged from this approach were powerfully visible in 2012 in our workshop on mapping and modeling the Venetian cistern system and its implications for urban and architectural planning, and the workshop on the City and the Lagoon (2014). The latter has been developed into a spectacular exhibit created by the Visualizing Venice consortium and curated by Donatella Calabi inaugurated on September 26th, 2015, at the Palazzo Ducale in Venice.[3]

As is often the case with new enterprises, funding has been a challenge. VIU worked with us to create our Venice laboratory and covered the expenses of hardware and software. Starting in 2012 the Gladys K. Delmas Foundation supported the cost of tuition. In 2015 a grant provided by the Getty Foundation covered the expenses of travel and residency in Venice for participants and faculty. The new support from the Getty Foundation has transformed our program, enabling participants to come from long distances (China, Australia, Mexico) and shorter ones (Poland, Slovakia, Germany, France and Italy itself) to join the workshop, greatly enhancing both the quality and number of applicants and participants.

Notes

1 http://www.dukewired.org/
2 http://www.visualizingvenice.org/
3 http://palazzoducale.visitmuve.it/it/mostre/mostre-in-corso/mostra-acqua-e-cibo/2015/04/7052/storie-della-laguna-e-della-citta/

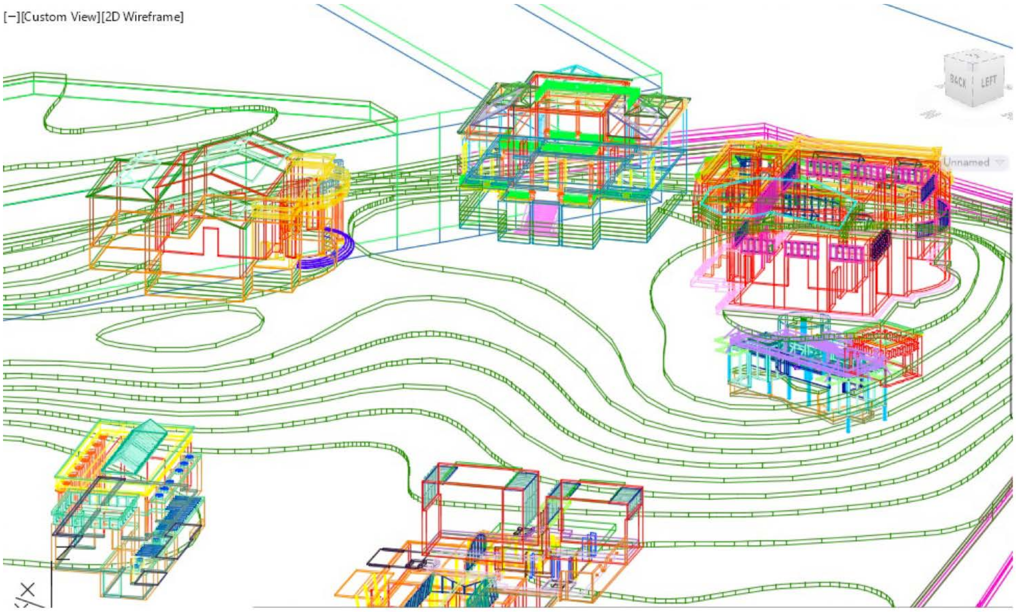

Fig. 1: Image of Biennale Pavilions by Mirka Dalla Libera, IUAV