International Council of Museums (ICOM)
Le Comité International pour la Documentation de l'ICOM (CIDOC)

CIDOC 2015, New Delhi
"Documenting Diversity – Collections, Catalogues & Context"

Session– 2. Techniques and methods of documentation

Vigyan Bhawan Annexe, 7th of September 2015.

Full paper presentation by Magnus Bognerud and Gro Benedikte Pedersen

RepCol (Representability in the Collections) - How to visualize an entire collection and the value of doing so

Authors

Magnus Bognerud (1983): Holds a MA. degree in museology from the University of Oslo (2012). Works as a documentation and digital collection management specialist at the National Museum of Art, Architecture and Design in Norway.

Gro Benedikte Pedersen (1973): Holds a Cand. philol. degree in art history from the University of Oslo (2002). Works as a professional coordinator for the Digital collection management at the National Museum of Art, Architecture and Design in Norway.

Abstract

Good metadata, both qualitative and quantitative, are fundamental in facilitating use and providing access to collections. Museums today are therefore obliged to provide consistent, rich and linked data about their collections. At the same time, museum professionals and art historians are always interested in having access to as much information as possible concerning an artwork, a cultural object or a whole collection. To meet these objectives, we always strive towards best practice, increased digitization and open access. However, how can we apply and reuse digitized material? In this paper, we will introduce two projects where we have investigated different solutions for the dissemination of art collections based on digitized back end content. Both are pilot projects by the Section for Digital Collection Management at the National Museum in Norway, realized in close collaboration with the Oslo-based company Bengler with funding provided by the Arts Council Norway.

Visualizing collections: how to and why

Our purpose as documentation specialists is among other things, to boldly go where no art historian has gone before. We explore how to extract the salience of our museum's digitized information. The incentive to digitize and open up the museum collections is there, and technological solutions are flourishing. So, how to engage with this material? How do we maximize the undoubted potential of open collection data, both for public and scholarly audiences? What are the needs and demands, and from whom? Exactly what are our long and short-term goals? What is even possible with the type of data that we have today? The field seems somewhat shrouded in uncertainty. Along with the evolving of our museum's digitizing processes, we have created concrete proposals towards answering such questions. This has resulted in the search for new visual grammars and the use of machine analysis methods. In this paper we focus on two of our recent projects: *Repcol* and *Principal Components* and their possible theoretical approaches.

Repcol

The first project, *Repcol*, started out as an idea to display our entire collection of old masters and 19th century and modern art in one image, one model, or one diagram. At the same time as two-dimensional representations were created from Tate Britain's dataset at the

OpenGLAM in 2013, Bengler suggested we create something similar, but in a three-dimension.[1] By using a very limited set of basic collection information, we wanted to make a visual example of representability in our collection. This resulted in the making of a visualization and search tool. *Repcol* is a visual imprint, unique and representative for the entire collection. It uses a simple visual grammar to translate core data of inventory number, production- and acquisition date and limited biographical data (name, gender, year of birth/death), into a navigable, three-dimensional figure. See fig.1 *Repcol*.

Fig.1 *Repcol*

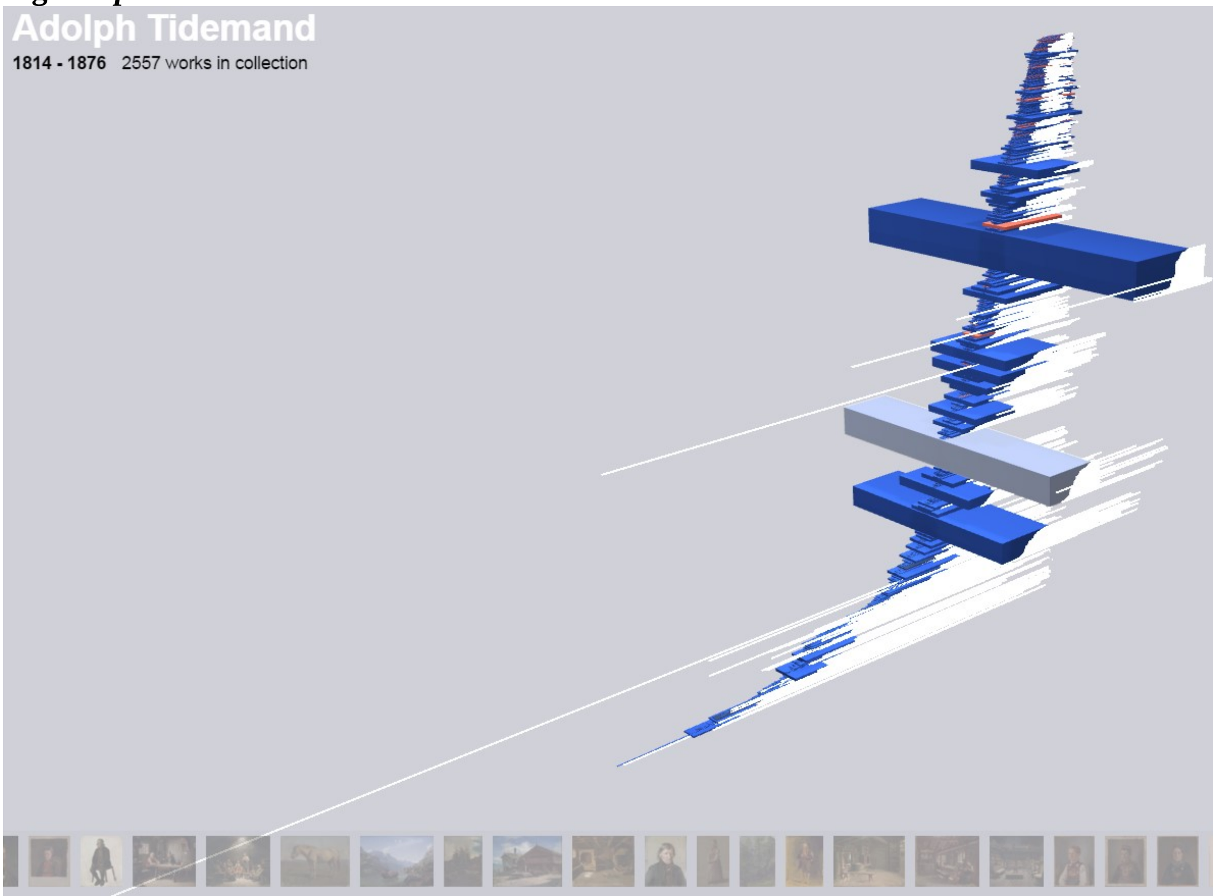

(Illustration source: The National museum of Art, Architecture and Design/ Bengler).

Our collaborators constructed the three dimensional shape with 'building blocks' that each represent one artist in the collection. The oldest artist in our collection, Andrea di Bartolo (1389-1428), is represented by the lowermost block, and the youngest, Marthe Karen Kampen (b.1986), is represented by the upper. The number of works collected, together with the artist's lifespan, define the size of each block. Along the sides, thin white lines represent the artworks created by the artist. The lines start at its creation date, and end when our museum acquired it. The color blue represents male artists and red represents females. You can flip the figure, zoom in and out, and move up and down the blocks step by step. [2] Fig.2 *Repcol*, *detail*.

Fig.2 ***Repcol, detail*****.**

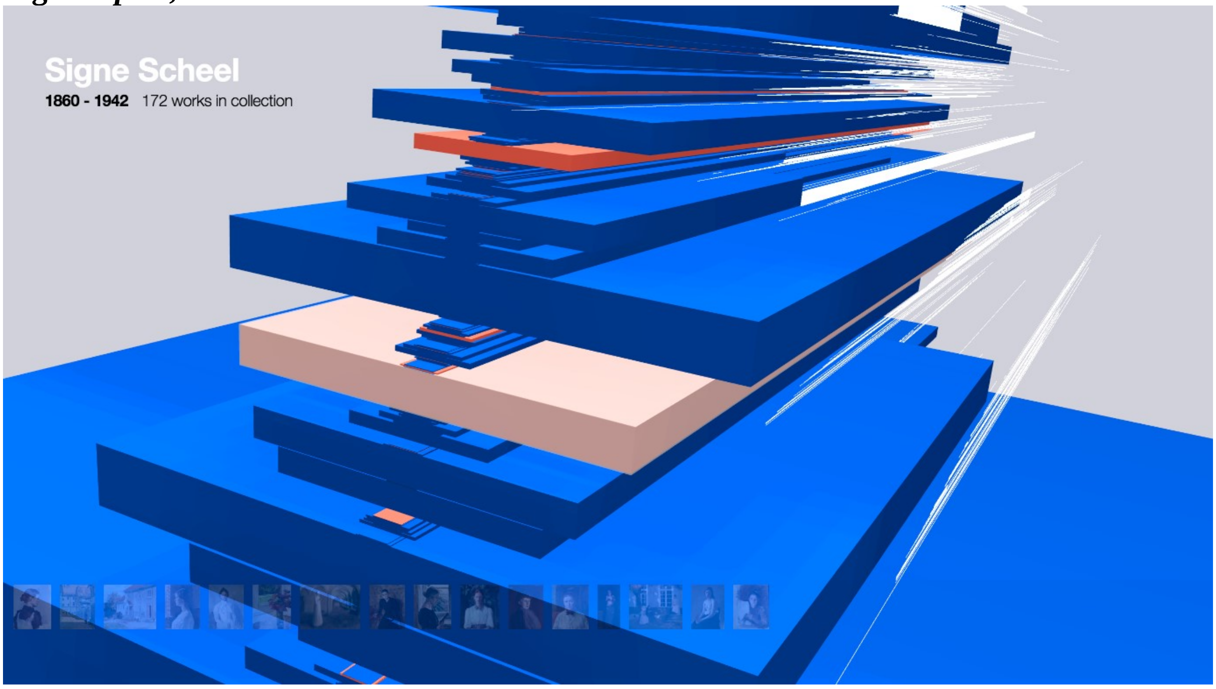

(Illustration source: The National museum of Art, Architecture and Design/ Bengler).

As we connected photo files and browsing possibilities to the figure, *Repcol* became a new and experimental prototype for querying and exploring the collection both at micro and macro levels.[3]

In configuring the dataset that would form this visualization, and in the aftermath of its completion, we uncovered several interesting aspects of our collection data as a whole. We were able to review the structure of our database from a new angle, and discovered both limits and hidden possibilities. The flaws in our dataset were suddenly exposed and the amount of work lacking in qualitative registration became painfully apparent. Nevertheless, with the simplest dataset, we made a great tool for both art historical and cataloguing purposes.

An aim for the project was to provide means for the public to engage with our data in novel ways and perhaps also encourage museum staff in their daily digitization work. One of the key results though, was that we found ourselves inspired. The prototype is great, fun and useful and the project gave us a deeper insight into the published metadata. Until the machine analysis and this display of our material, we considered our collection metadata to be rather good. In the last years, we have emphasized the standardization of catalogue content and believed that the data had improved more than it really had. Even when reducing the parameters that *Repcol* would include, the Bengler-team and we still had to perform considerable data cleaning before having a more or less complete set. Over the last 180 years of collecting and cataloguing, diverse strategies within different professional fields of art history and different institutional eras have informed the cataloguing and standardization of content in varied ways. Hence, it was more challenging than we had foreseen to establish a sound and structured set of core metadata that could represent the collections as a whole. Therefore, one of the main takeaways from the project was gaining an appreciation of the difficulties of creating good tools and experiences, for both researchers and casual browsers, when little or inconsistent metadata is attached to each work. To us this was another example of how essential rich and consistent metadata is in facilitating the use of and access to art collections.

Knowing this, we still do not have the time or resources to register and crosscheck all old and newly catalogued information at once. This is a Sisyphean and ever evolving task. In the meantime, we choose to publish limited data sets. With regard to this, we were left with the question: Are there other strategies that we could employ to enrich our dataset without de-prioritizing other important museum tasks? We ended up on testing out methods of 'principal component analysis'.

Principal components

This second project is still a work in progress. After having investigated representative volumes within our catalogued text in *Repcol*, we have now turned to look for ways to define representative quantities within our photo files. We are currently working with machine vision and so called deep-learning techniques that have developed from the general field of machine learning. The discipline of machine learning has experienced rapid progress over the last few years, and algorithms are rapidly becoming more capable at classifying images - in the past year reaching near human skill within certain domains. The gains within the field have been driven by access to larger data sets and more affordable computations resulting from parallel computing - enabled by the graphic processing units (GPUs) initially developed for computer entertainment.[4] Overall, we have easier access to a more fertile ground for training artificial neural networks (ANN) than what was the case in the last decades of the 20th century. For the processing of images, as broadly delineated in the Economist's article *Rise of the machines*, from May 9, 2015: ‘Early neural networks were limited to dozens or hundreds of neurons, usually organized as a single layer. The latest, used by the likes of Google, can simulate billions. With that many ersatz neurons available, researchers can afford to take another cue from the brain and organize them in distinct, hierarchical layers (see diagram). It is this use of interlinked layers that puts the “deep” into deep learning’. Fig3 Layer cake.

Fig.3 Layer cake

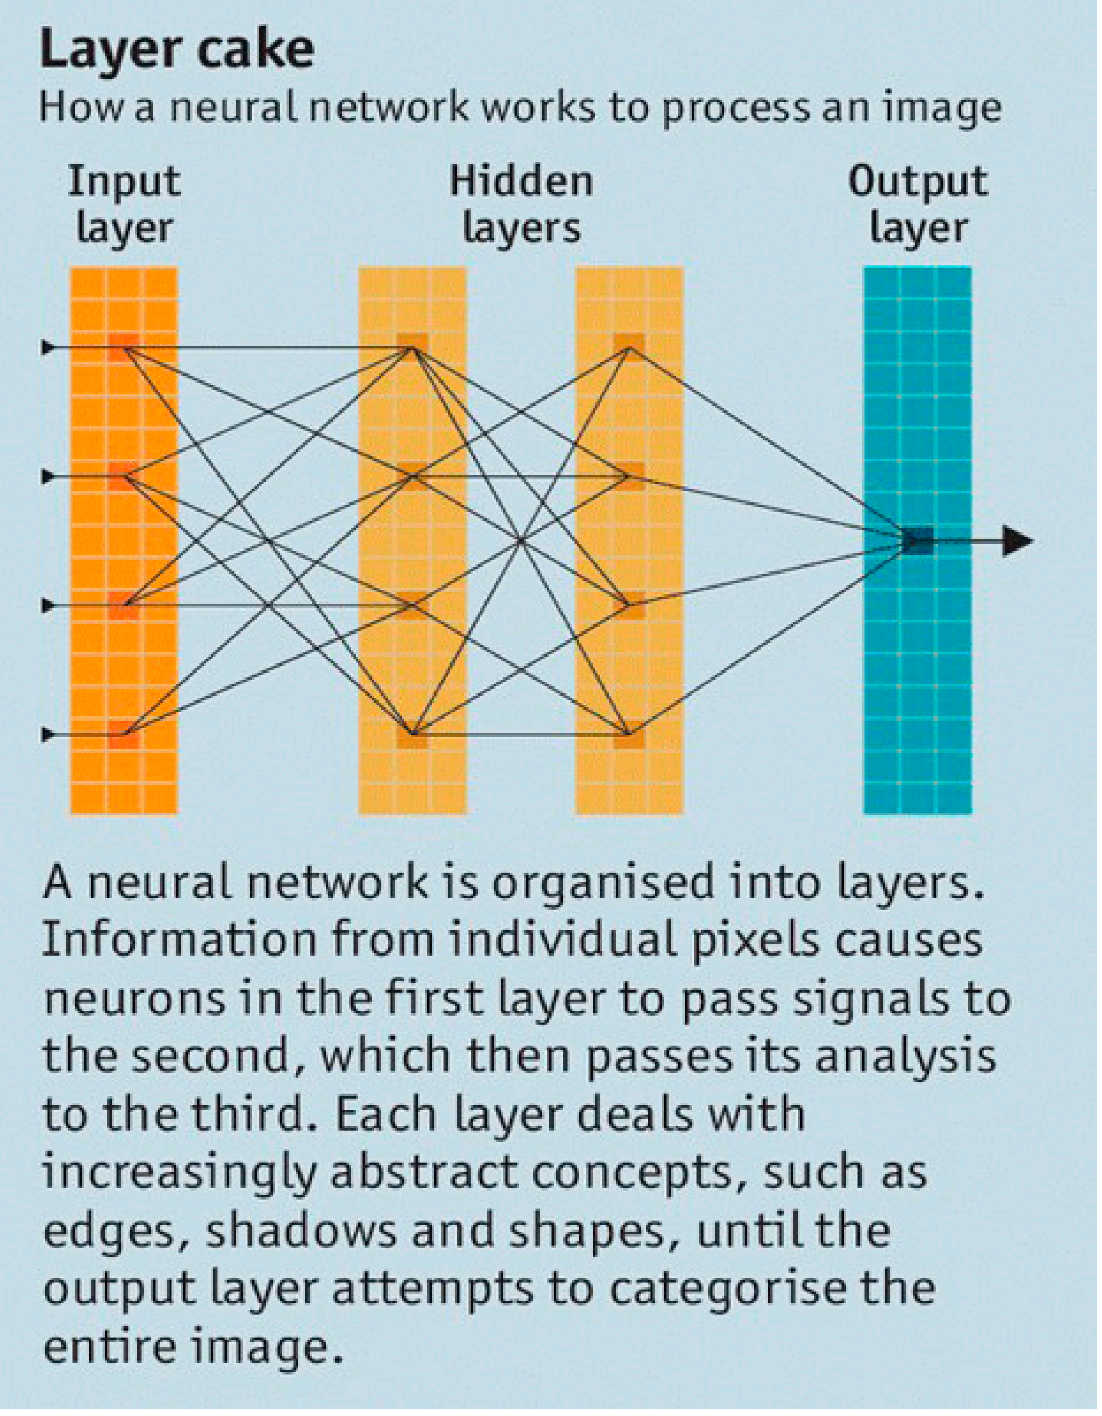

Illustration source: The Economist 9.5.2016.

For us, some core questions in the *Principal Components*-project are; can we teach computer algorithms to recognize and identify motifs, techniques and periods in artworks? How accurate can algorithms trained on art collections data get? Will we be able to produce reusable classifications? As in Repcol, one of our main goals is still to define and create new queries and thereby enhance the user experience in searching the collections. In addition, we will study how machines can do general classification tasks usually assigned to the museum staff. By the means of training computer algorithms, we believe that it is possible to enrich our catalogue with both conventional art history classifications, as for example *Iconclass* categories, but also maybe add new unconventional means of exploration.

In probing these issues with machine vision specialists, we go through sets of data and photos to seek possibilities of the automated quantification of information. We are harvesting data from resources such as the API behind our service *Search the collections* (http://samling.nasjonalmuseet.no/no/) and *Repcol*, and other datasets that are available online. This to explore the value of these algorithms in terms of finding and marking objects in images, identifying faces, gender and age, likeness to a sketch, or classifying objects by parameters like composition, technique, color, or style.

From a preliminary test, we have a demonstration of some of our algorithm's findings. In the oil painting *Self-portrait with cigarette* (1895) *by* Edvard Munch (1863-1944), the machine easily recognized a human figure. Alas, it also recognized a hippopotamus lingering in the corner. See fig.4 below.

Fig.4 *Self-portrait with cigarette* (1895) *by* Edvard Munch

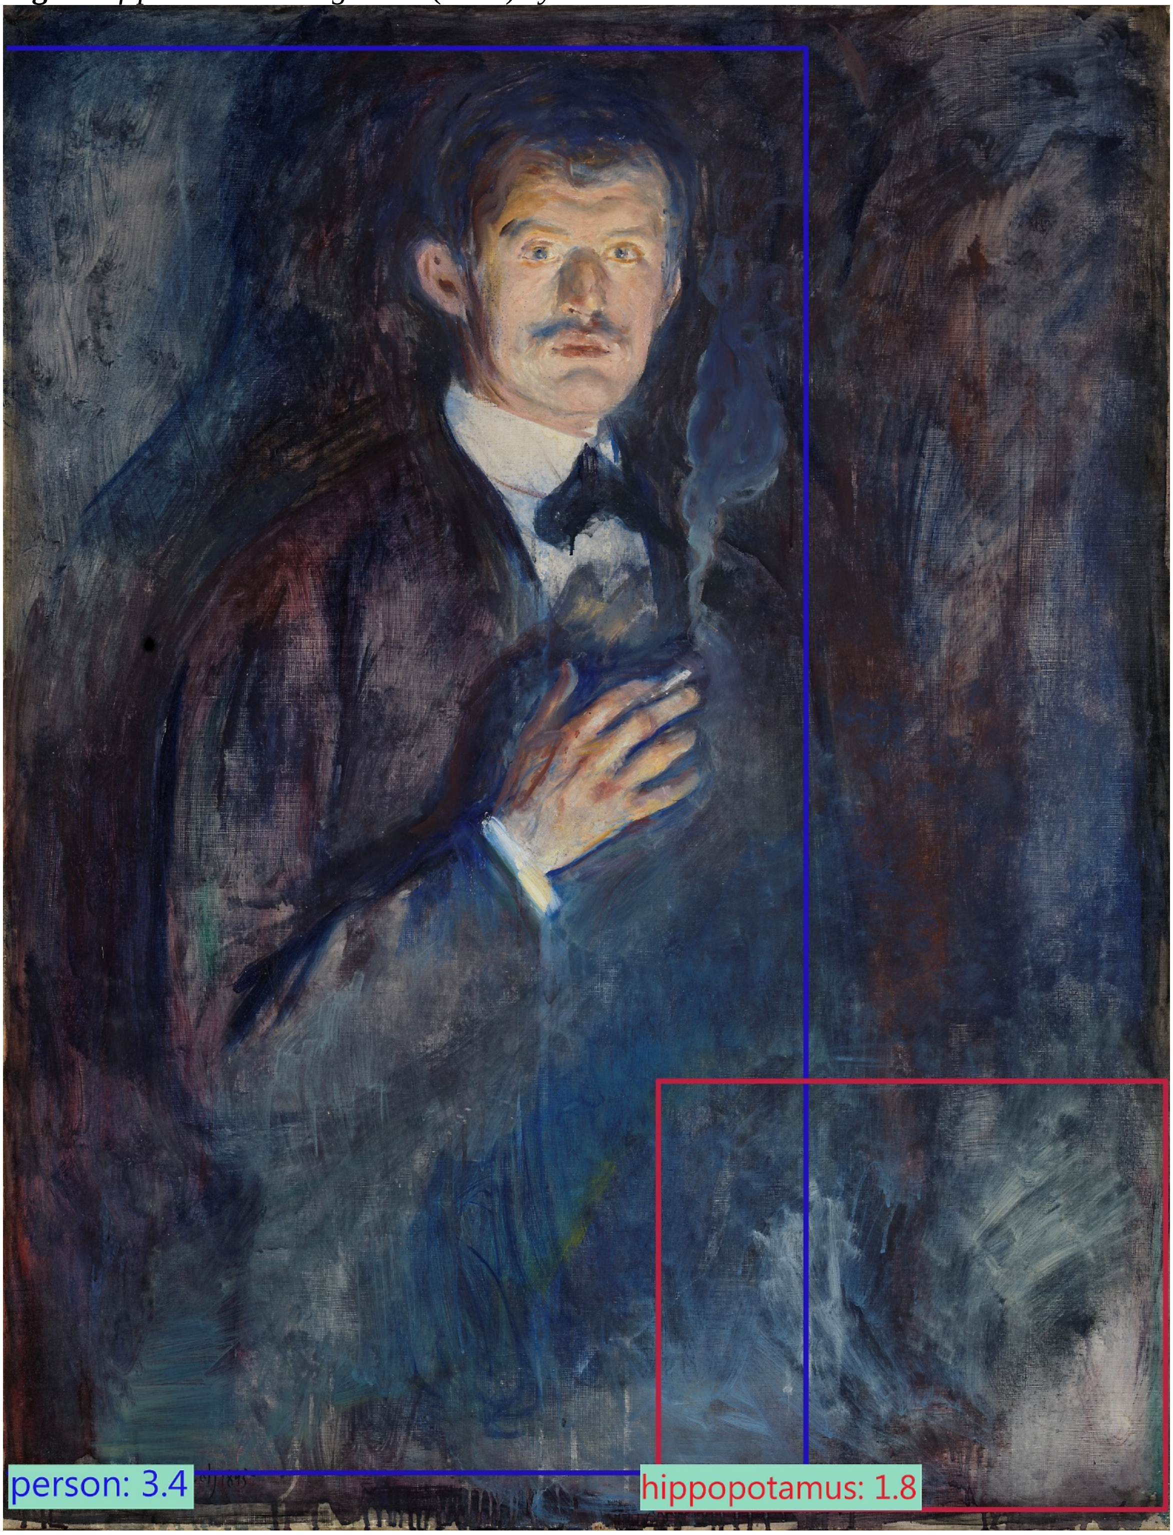

Illustration source: The National museum of Art, Architecture and Design/ Bengler.

Another example from one of our iconic, national romantic oil paintings, the *Bridal procession on the Hardangerfjord* (1848) by Adolph Tidemann (1814-1876) and Hans Gude (1825-1903), also shows how the machine found unexpected motifs, see fig.5.

Fig.5 *Bridal procession on the Hardangerfjord* (1848) by Adolph Tidemann and Hans Gude

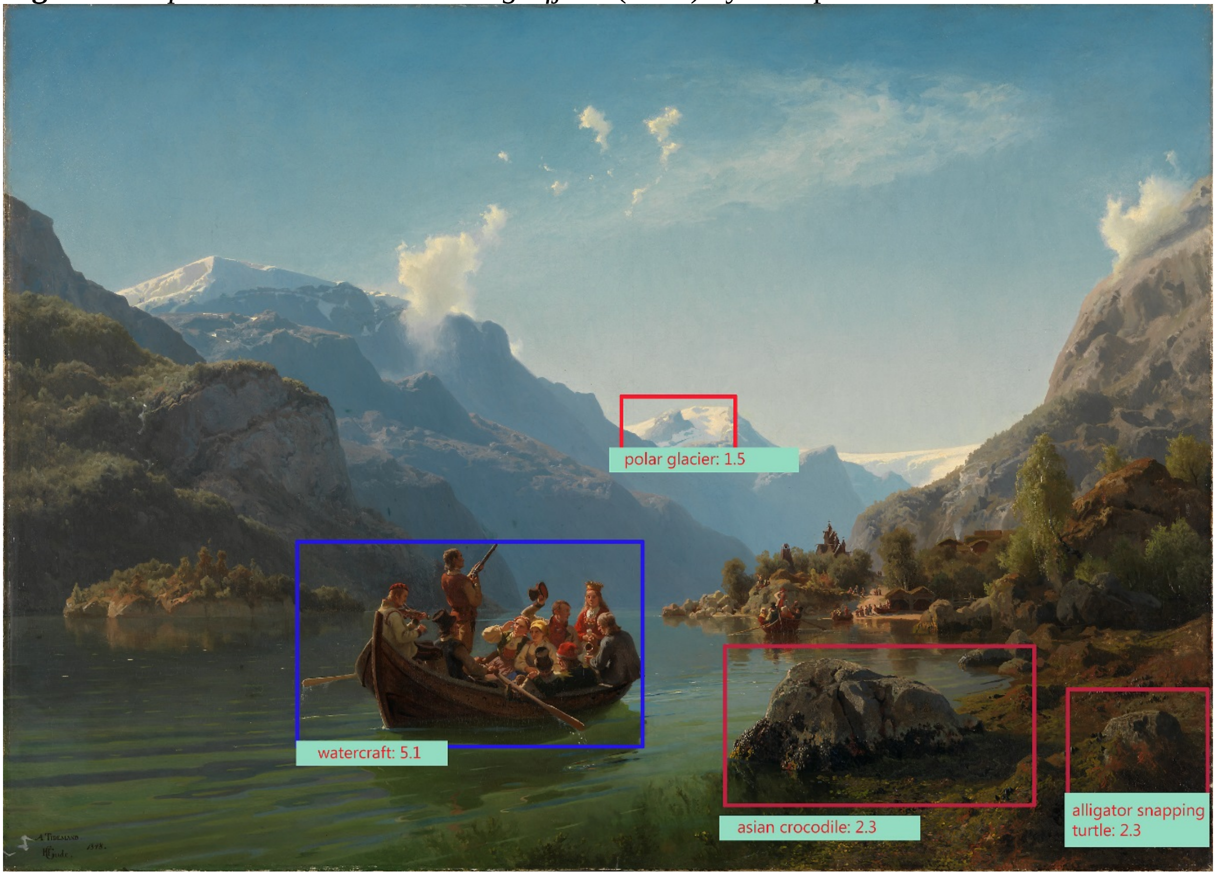

Illustration source: The National museum of Art, Architecture and Design/ Bengler.

Here the machine clearly identified the boat and the glacier. However, its context free gaze also located some other exotic figures, never known to have Norway as its habitat, like a crocodile and a turtle.

During the second phase of testing, our designers have applied a neural network initially trained on the ImageNet dataset by Google. This is realized within the framework *Caffe*, developed by *Autonomous Perception Research Lab* at Berkeley.[5] This model has been retrained on the data set of works from Wikiart's collection (Wikimedia commons/ Art). We have trained models to both classify tags and stylistic components. Further we used it to classify the 30 000 images from our catalogue of old masters and 19th century and modern art. The classifier produces a 1024 dimensional representation of each work. In order to present a coherent image of similarities in the collection for a given model to an end user, we then reduced this data set into a two-dimensional spatial representation using the t-SNE algorithm.[6]

When looking at these representations of the two models, one trained on tags and the other on trained on style, the difference is striking. The style algorithm succinctly groups the national romantic landscape series, separating out the motifs portraying water before moving into naval imagery. In comparison, the model trained on tags for example effectively groups portraits regardless of stylistic trappings. We can now group pictures of artworks by motifs likeness, technique, composition or color. This is something we can use as a starting point for simple user interfaces to the public that browse our online collections. See example of a face recognition test in fig.6 below.

Fig.6 Face recognition

Illustration source: The National museum of Art, Architecture and Design/ Bengler.

One of the next steps for us will be to train the algorithm further. Firstly, we will test it on motifs classified by *Iconclass*, and see if larger aggregates of higher quality metadata improves the classification. Further, we want to explore if such classifications will let us juxtapose the art in our collections in novel and interesting ways, providing insight that will speak to the emerging role of machine learning.

In doing so we find ourselves in unfamiliar terrain, working at the intersection between methods of machine learning, conventional art history cataloguing, the displaying of collection content, and retrieving user generated (and actually machine generated) content. At this intersection of complex fields of computer science on the one side and conventional art history cataloguing on the other, we also want to study different theoretical aspects that in our case border on both (digital) art history and robotics.

The uncanny valley of robotics

The value of this project is twofold, both prosaic and profound. As we investigate new possibilities of understanding images and generating new data, we also have the ambition to frame our projects with theories from aesthetics and psychology that touch on to machine learning and the discourses regarding the development of Artificial Intelligence (AI).

Our project is a juncture between conventional art history thinking and robotics: two fields traditionally far apart. In this liminal field, we are looking at some unsettling qualities of these robotic methods. The means by which the machines work, and particularly the mistakes they make, remind us more of faults in biological, cognitive processes than the traditional binary breakdowns of machine function. Robots with a certain degree of human resemblance often invoke an unsettling feeling. The Freudian notion of *Das Unheimliche* depicts the disturbing feeling that emerges when something unfamiliar is recast as familiar. The notion of the *uncanny valley* known from psychology and aesthetics describes this experience. See Fig7.

Fig.7 The uncanny valley

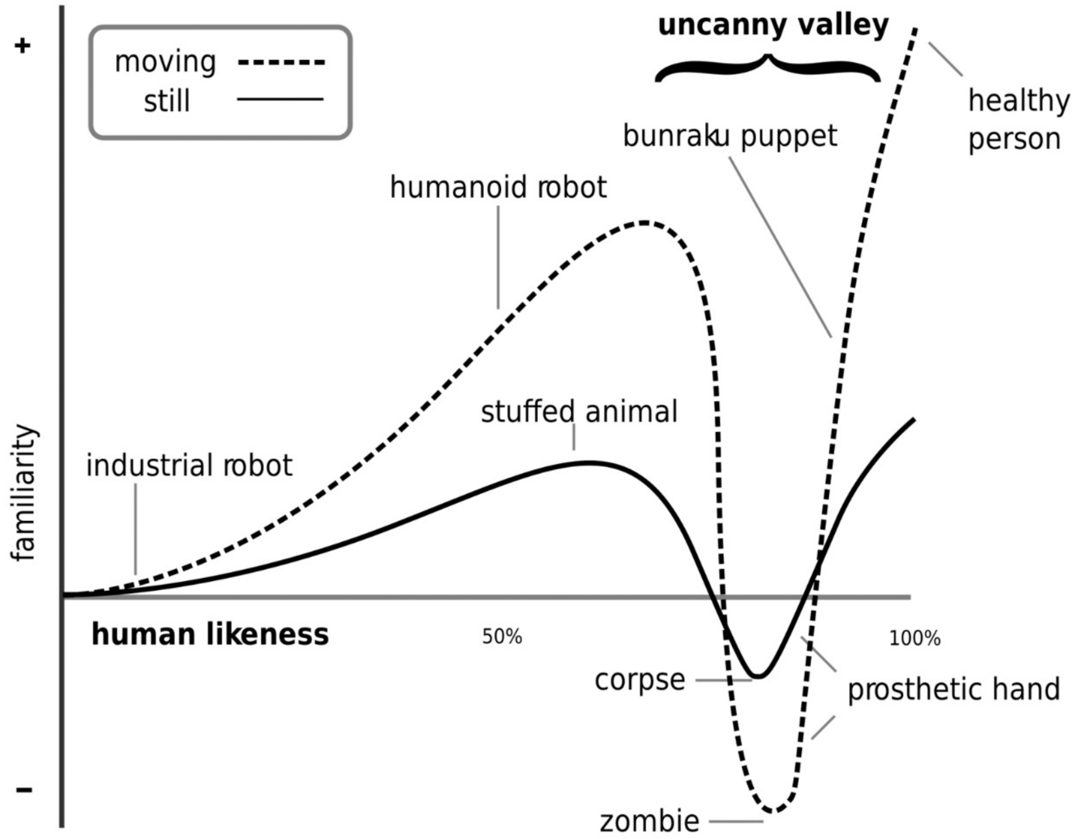

Illustration source: Wikipedia.org/wiki/Uncanny_valley.

The graph in fig.7 depicts the uncanny valley, a notion identified in 1970 by the robotics professor Masahiro Mori. After having tested out the algorithms capabilities in *Principal Components,* we too encountered that wondrous sense of seeing how the artificial neurons convincingly simulate the behavior of biological neurons. To us it seems that some machine learning algorithms are moving towards the kind of uncanny geographies depicted above. Not by their visual surface quality, but through their behavior. We refer to the way that the algorithms behave like humans when identifying depicted forms, and how they make very human mistakes. In addition, in the algorithms self-repair of trial and error, a kind of a *technical individualization* originates when the algorithm uses its human adversary's actions (*input*) to make new links and expand their own repertoire (*output*) unintended or unexpected from the programmer's side.[7]

Where engineered systems operate with precise perfection inside their design envelopes, trained machine learning algorithms are uncertain. They waver, blunder and make mistakes.

In the structures of a Convolutional Neural Network (CNN) - that make up a class of machine learning algorithms, layers of neurons are trained to activate when they see particular relevant features in images. When carrying out classification, thousands of such intermediate structures are combined into a single simple answer (see also fig.3 above of how neural network works to process an image).

In John Markoff's 2015 book *Machines of loving grace*, he quotes Geoffrey Everest Hinton, a cognitive psychologist and computer scientist well known for his work on artificial neural networks. In describing 'the rise, fall and resurrection of AI', Markoff discusses how deep learning nets have made significant advances the recent years. Hinton says that deep learning 'is a new continent and the researchers still have no idea what is really possible".[8] As visitors to Hinton's Deep-learning-continent, we are interested in finding out more about how Art history adapts to this landmass.

Art history and Robotics

In the discipline of Art history new expressions like *Digital curator*, *Digital Art History,* or *Digital humanities* represent new principles arising because of the new technological possibilities. These principles require interaction between opposite poles in professional life, the art historian and the computer programmer. By coupling art history with robotics in projects like these, many new questions arise. Finding answers to these questions can be especially challenging when technical matters are unfamiliar to the curators and art history is unfamiliar to computer programmers. By trying to follow new possibilities like machine learning, the discipline of art history needs to understand new technical languages to gain access to an infinite potential. We hope that our work will contribute to build new understanding between computer scientists and art historians. Nevertheless – at the core of our work is still the urge to promote the importance of the never-ending documentation work carried out in museums every day.

Bibliography

Bell, S. and Croson, C. “Artificial Neural Networks as a tool for Archeological data analysis”, *Archeometry 40* (1998): 139-151.

Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York: Basic Books, 2015.

Hamilton, James R. “‘The Uncanny Valley’ and spectating Animated Objects”. *Proceedings of the European Society for Aesthetics* 6 (2014): 190-207.

Klinke, Harald. “Digital Art history – a new frontier in research”. (Session at the 11th Triannual Nordik Committee for Art History Conference, Reykjavik, Iceland, May 13-16, 2015).

Markoff, John. *Machines of loving grace*. New York: HarperCollins Publishers, 2015.

Ruthaven Ian and Chowdry, G.G., ed. *Cultural Heritage Information. Access and Management*. Croydon: Facet Publishing, 2015.

Schmale, Wolfgang. “Digitale Vernunft”, *Historische Mitteilungen* 26 (2013/2014): 94-101.

Surkemper, Liska and Klinke Harald, ed. *International Journal for Digital Art History* 1(2015).

Warwick, Claire, Terras, Melissa and Nyhan, Julianne, ed. *Digital humanities in practice*. London: Facet Publishing, 2012.

---

[1] Examples of two project results made from Tate Britain’s data contribution to OpenGLAM: http://research.kraeutli.com/index.php/2013/11/the-tate-collection-on-github/ or http://www.ifweassume.com/2013/11/the-dimensions-of-art.html . OpenGLAM: http://openglam.org/open-collections/ [23.2.2016].

[2] See more details about *Repcol* in this video: http://vimeo.com/85755986. [23.2.2016].

[3] You can test the prototype tool here http://repcol.bengler.no/ [23.2.2016].

[4] For further readings of the evolving history of neural networks, see for example: John Markoff, *Machines of loving grace*, (New York: HarperCollins Publishers, 2015).

[5] http://bvlc.eecs.berkeley.edu/ [23.2.2016]. Caffe is comparable with for example *Torch*, developed by technicians from Facebook, Twitter and Google.

[6] More about the t-SNE algorithm here: http://lvdmaaten.github.io/tsne/ and http://cs.stanford.edu/people/karpathy/cnnembed/ [23.2.2016].

[7] Cf. art historian Ina Blom's discussion of input/ output and technical individualization http://www.kunstkritikk.no/artikler/mediearkeologen/?d=no [23.2.2016].

[8] Markoff, *Machines of loving grace*, 156.