

Algorithm finds hidden connections between paintings at the Met

30 July 2020, by Rachel Gordon



Credit: Massachusetts Institute of Technology

Art is often heralded as the greatest journey into the past, solidifying a moment in time and space; the beautiful vehicle that lets us momentarily escape the present.

With the boundless treasure trove of paintings that exist, the connections between these works of art from different periods of time and space can often go overlooked. It's impossible for even the most knowledgeable of art critics to take in millions of paintings across thousands of years and be able to find unexpected parallels in themes, motifs, and visual styles.

To streamline this process, a group of researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and Microsoft created an algorithm to discover hidden connections between paintings at the Metropolitan Museum of Art (the Met) and Amsterdam's Rijksmuseum.

Inspired by a special exhibit "Rembrandt and Velazquez" in the Rijksmuseum, the new "MosAlc" system finds paired or "analogous" works from [different cultures](#), artists, and media by using deep networks to understand how "close" two images are. In that exhibit, the researchers were inspired

by an unlikely, yet similar pairing: Francisco de Zurbarán's "The Martyrdom of Saint Serapion" and Jan Asselijn's "The Threatened Swan," two works that portray scenes of profound altruism with an eerie visual resemblance.

"These two artists did not have a correspondence or meet each other during their lives, yet their paintings hinted at a rich, latent structure that underlies both of their works," says CSAIL Ph.D. student Mark Hamilton, the lead author on a paper about "MosAlc."

To find two similar paintings, the team used a [new algorithm](#) for [image search](#) to unearth the closest match by a particular artist or culture. For example, in response to a query about "which musical instrument is closest to this painting of a blue-and-white dress," the algorithm retrieves an image of a blue-and-white porcelain violin. These works are not only similar in pattern and form, but also draw their roots from a broader cultural exchange of porcelain between the Dutch and Chinese.

"Image retrieval systems let users find images that are semantically similar to a query image, serving as the backbone of reverse image search engines and many product recommendation engines," says Hamilton. "Restricting an image retrieval system to particular subsets of images can yield new insights into relationships in the visual world. We aim to encourage a new level of engagement with creative artifacts."

How it works

For many, art and science are irreconcilable: one grounded in logic, reasoning, and proven truths, and the other motivated by emotion, aesthetics, and beauty. But recently, AI and art took on a new flirtation that, over the past 10 years, developed into something more serious.

A large branch of this work, for example, has

previously focused on generating new art using AI. There was the [GauGAN](#) project developed by researchers at MIT, NVIDIA, and the University of California at Berkeley; Hamilton and others' previous [GenStudio](#) project; and even an AI-generated artwork that sold at Sotheby's for [\\$51,000](#).

MosAlc, however, doesn't aim to create new art so much as help explore existing art. One similar tool, Google's "[X Degrees of Separation](#)," finds paths of art that connect two works of art, but MosAlc differs in that it only requires a single image. Instead of finding paths, it uncovers connections in whatever culture or media the user is interested in, such as finding the shared artistic form of "Anthropoides paradisea" and "Seth Slaying a Serpent, Temple of Amun at Hibis."

Hamilton notes that building out their algorithm was a tricky endeavor, because they wanted to find images that were similar not just in color or style, but in meaning and theme. In other words, they'd want dogs to be close to other dogs, people to be close to other people, and so forth. To achieve this, they probe a deep network's inner "activations" for each image in the combined open access collections of the Met and the Rijksmuseum. Distance between the "activations" of this deep network, which are commonly called "features," was how they judged image similarity.

To find analogous images between different cultures, the team used a new image-search data structure called a "conditional KNN tree" that groups similar images together in a tree-like structure. To find a close match, they start at the tree's "trunk" and follow the most promising "branch" until they are sure they've found the closest image. The data structure improves on its predecessors by allowing the tree to quickly "prune" itself to a particular culture, artist, or collection, quickly yielding answers to new types of queries.

What Hamilton and his colleagues found surprising was that this approach could also be applied to helping find problems with existing deep networks, related to the surge of "deepfakes" that have recently cropped up. They applied this data structure to find areas where probabilistic models,

such as the generative adversarial networks (GANs) that are often used to create deepfakes, break down. They coined these problematic areas "blind spots," and note that they give us insight into how GANs can be biased. Such blind spots further show that GANs struggle to represent particular areas of a dataset, even if most of their fakes can fool a human.

Testing MosAlc

The team evaluated MosAlc's speed, and how closely it aligned with our human intuition about visual analogies.

For the speed tests, they wanted to make sure that their data structure provided value over simply searching through the collection with quick, brute-force search.

To understand how well the system aligned with human intuitions, they made and released two new datasets for evaluating conditional image retrieval systems. One dataset challenged algorithms to find images with the same content even after they had been "stylized" with a neural style transfer method. The second dataset challenged algorithms to recover English letters across different fonts. A bit less than two-thirds of the time, MosAlc was able to recover the correct image in a single guess from a "haystack" of 5,000 images.

"Going forward, we hope this work inspires others to think about how tools from information retrieval can help other fields like the arts, humanities, social science, and medicine," says Hamilton. "These fields are rich with information that has never been processed with these techniques and can be a source for great inspiration for both computer scientists and domain experts. This work can be expanded in terms of new datasets, new types of queries, and new ways to understand the connections between works."

Provided by Massachusetts Institute of Technology

APA citation: Algorithm finds hidden connections between paintings at the Met (2020, July 30) retrieved 22 October 2021 from <https://techxplore.com/news/2020-07-algorithm-hidden-met.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.