Aisling Kelliher
*Carnegie Mellon University*

# Machines Learning Culture

**Aisling Kelliher**
***Carnegie Mellon University***

Machine learning, big data, predictive analytics, and deep learning are all commonly referenced topics in today's computing and business worlds. Commercial entities are interested in such approaches for predicting customer churn or preventing fraud, while academics are busy identifying rich challenges related to acquiring and cleaning training data and developing large-scale algorithms. The popular press and social media have generated considerable excitement about the potential of these innovations to greatly enhance or decimate humanity—possibly even both at the same time.[1]

In May 2014, a number of celebrated scientists, including Stephen Hawking, authored a column in a UK newspaper, *The Independent,* that struck a cautionary note in discussing the possible threat of fully realized artificial intelligence.[2] Using the (somewhat suspect) narrative plot of the Hollywood movie *Transcendence* as a framing device, the authors highlighted the challenges of reaping the benefits of advances in AI and advocated the need to study possible incalculable risks. Subsequent commentary on the newspaper website and elsewhere debated the merits of the overall argument, with many weighing in using references to other speculative movies and science-fiction literature.

The article's movie-framing gambit clearly succeeded in captivating the attention of considerable sections of the online public, motivating diverse discourse about a complex topic. This leads me to consider what other forms of cultural practices and production might similarly engage audiences in contemplating the promise, beauty, menace, and power of big-data-style machine learning?

Some recent interdisciplinary explorations, integrating contributions from computer science, math, the digital humanities, and the arts, point to the utilitarian and expressive capabilities of machine-learning approaches in creating work with diverse appeal. These initiatives include research within the relatively traditional domain of historical art analysis, a growing collection of body-tracking work using machine learning quietly in the background, and a variety of provocative art installations that place algorithmic computing front and center. While these projects tackle their subject at varying levels of scale and depth, and in different contexts, each contributes to building the public discourse about the impact, role, and reach of machine learning in our lives.

## Visual Art Analysis

The role of an art historian includes the provision of a context within which a reader or viewer can better appreciate the significance of a considered work. Recent research from the Digital Humanities Research Laboratory at Rutgers University points to the potential of machine-learning techniques to assist (but not yet replace) historians in efficiently identifying evidence of possible artistic influence across groups of paintings.[3] While such inferences in general are naturally subjective, insights gained using multiple classification techniques can certainly help support or challenge particular arguments.

In "Toward Automated Discovery of Artistic Influence," the authors carefully note that their approach is not "asserting truths but instead suggesting a possible path towards… measuring influence."[3] Using a dataset of 1,710 images of artworks (primarily paintings) by 66 artists created over approximately 550 years, the authors experiment with several discriminative and generative methodologies, before exploring low-, intermediate-, and semantic-level features. In their findings (see Figure 1), they indicate that their approach highlights possible evidence of artistic influence between paintings which, to their knowledge, have not been discussed together before, including Frederic Bazille's *L'atelier de Bazille* (1870) and Norman Rockwell's *Shuffleton's Barbershop* (1950).

While acknowledging that their work is just beginning to scratch the surface of automated or



1070-986X/15/$31.00 © 2015 IEEE Published by the IEEE Computer Society

*Figure 1. Comparison of the Bazille and Rockwell paintings, where yellow annotation indicates object similarity, red demarks compositional similarity, and blue denotes structural similarity. (Source: Ahmed Elgammal; used with permission.)*

semi-automated discovery of artistic influence, such endeavors point to the value of integrating computer science and digital humanities as a rich application space for evaluating and stretching machine-learning methods and performance and as a way for testing implicit or subjective judgments.

## Body Tracking

Off-the-shelf and customized machine-learning techniques are increasingly part of the toolset of electronic media artists interested in exploiting the potential of this computational approach. In the Augmented Hand Series project, collaborators Golan Levin, Chris Sugrue, and Kyle McDonald have created a physical installation containing "a real-time interactive software system that presents playful, dreamlike, and uncanny transformations of its visitors' hands" (see www.flong.com/projects/augmented-hand-series). Their system uses standard vision and machine-learning techniques (including the Leap motion library) to analyze the posture of participants who insert their hand into a box, with the resulting transformed image reimagining the displayed hand with additional fingers or autonomously moving digits (see Figure 2).

The creators of the Augmented Hand Series note the challenges that the public exhibition of their work brings up, including dealing with a variety of hand sizes, skin colors, skin ages, ornamentation, and general divergent finger features, all of which require additional and ongoing research as they develop new forms of transformation that delight and engage expanded audiences.

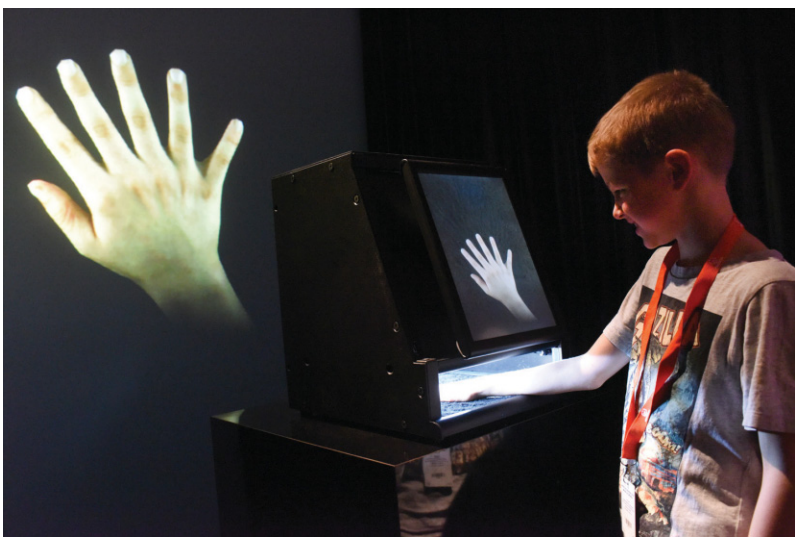

*Figure 2. A young child experiencing the Augment Hand Series installation. The system uses standard vision and machine-learning techniques to analyze the posture of participants who insert their hand into a box, with the resulting transformed image reimagining the displayed hand. (Source: Golan Levin; used with permission.)*

## Provocative Art Installations

Machine learning is brought to the fore in works that aim to directly translate and

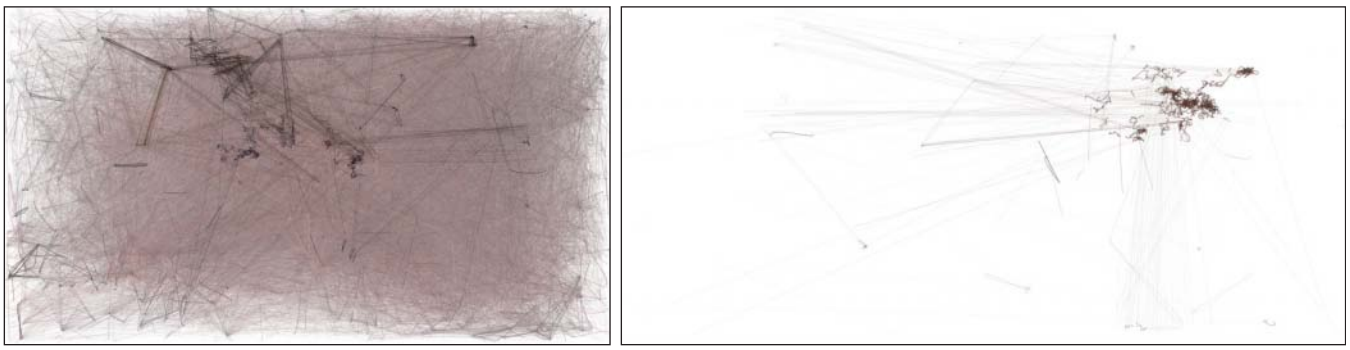

**(a)** **(b)**

*Figure 3. Still frames from the Computers Watching Movies project. The algorithm watches scenes from (a) the fast-paced movie* Inception, *and (b) the more character-driven movie,* Taxi Driver. *(Source: Benjamin Grosser, http://bengrosser.com; used with permission.)*

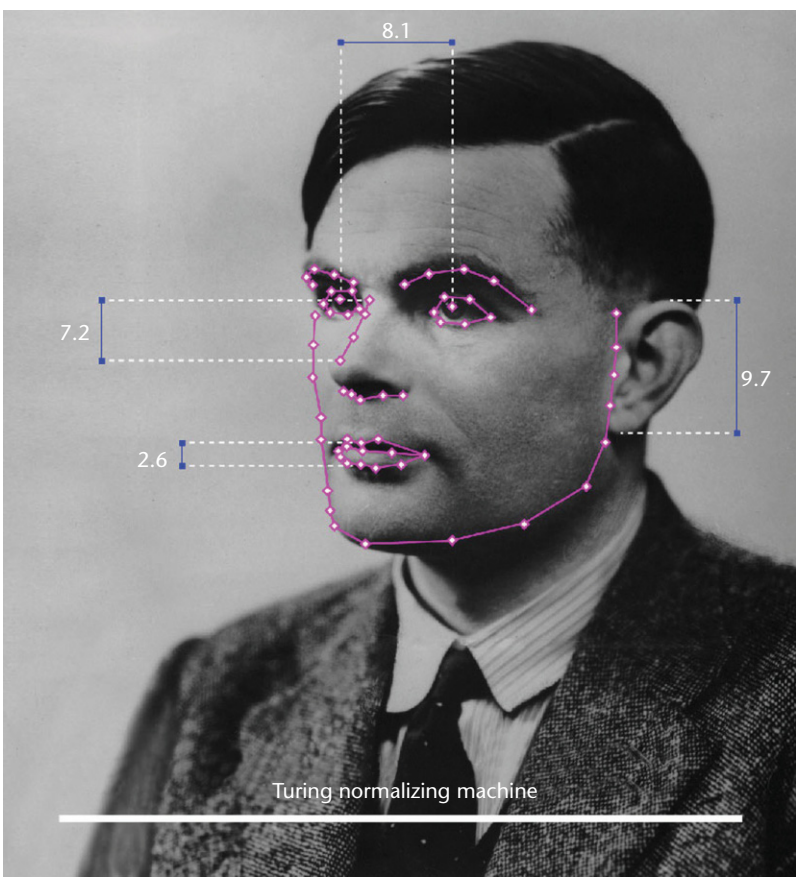


*Figure 4. Promotional image for the Turing Normalizing Machine, a collaborative experiment that aims to decode what society deems "normal." (Source: Mushon Zer-Aviv; used with permission.)*

represent unfolding algorithmic processes in a format that is legible to non-expert viewers.

### Representing Machine Vision

Benjamin Grosser's work aims to inform viewers about the underlying process by which a machine vision system differs from ours in terms of how it might "watch a movie." In his project, Computers Watching Movies, (see http://bengrosser.com/projects/computers-watching-movies), computer vision techniques and AI processing are integrated in melding a computational approach that instills a degree of agency to the viewing machine. The attention and disinterest of the algorithm is driven by the prominent features detected in the viewed movie frame, which in turn is demarked by an ever-shifting line trace drawn to match the dynamic computer gaze. The algorithm is directed to watch scenes from contemporary fast-paced movies such as *Inception* (see Figure 3a) and *The Matrix,* in addition to more character-driven fare such as *Annie Hall* and *Taxi Driver* (see Figure 3b).

The resulting analysis is presented as a temporal animation accompanied by the real-time audio from the attended scene. While the resulting animations make for gently hypnotic viewing, the relative lack of cohesion between the audiovisual interpretation and the well-understood-by-humans narrative points to a very different model of message or medium interpretation. The salient meaning derived by humans—based on years of exposure to movies and indeed, general life experience—is at odds with the limited training of the algorithm on a constrained dataset. Building on Stuart Hall's theoretical model of encoding and decoding communication, what is encountered here is a radically different contextual decoding framework, where the translation of significance, originality, or import could initially be understood as oppositional, before potentially moving, over time, toward a negotiated or even hegemonic

position.[4] Perhaps Grosser's work could point (even somewhat tongue-in-cheek) toward a world where new forms of automated critique and analysis help us better consider that which is novel, different, and (eventually) meaningful?

### Exploring Social Norms

Machine learning is also placed front and center in works such as the Turing Normalizing Machine, a collaborative experiment by Yonatan Ben Dimhon and Mushon Zer-Aviv. Here, the artists posit their approach as "experimental research in machine learning that identifies and analyzes the concept of social normalcy" (see http://mushon.com/tnm).

Inspired by, and in tribute to, Alan Turing's documented life experiences and contributions to the fields of computing and artificial intelligence, the installation presents participants with four video images of different people engaged in the act of pointing. The participants are then instructed to point at the image depicting the person or people they deem "most normal." Their selection is analyzed and its visual features are added to "an algorithmically constructed image of normalcy" (see Figure 4). In addition, a video recording of the participant making their selection is captured and added to the system for later use. Over time, the system develops a more refined model of so-called "normal presenting" appearance, bringing the artists closer to their goal of decoding "the mystery of what society deems 'normal.'" As commentary on the abject prosecution of Alan Turing by a prejudicial society, the piece both highlights and questions the power of human-aided machine-learning systems in classifying socially determined norms about what is palatable or acceptable.

### Visualizing Content

The artwork and designs of other practitioners give pause in consideration of how their work could enhance, but yet still resist, machine-learning practices. For example, designer Stefanie Posavec's exquisite visualizations of features—such as the underlying textual structures, sentence rhythms, and overarching themes of the works of several noted 20th century authors—were created by undertaking painstaking by-hand analysis of the primary texts.

In a process she describes as "knit(ing) together information visualizations," the laborious process of analyzing the text by hand—as opposed to "pressing 'enter' on a keyboard"—adds to the significance of the work (see www.stefanieposavec.co.uk/-everything-in-between/#/writing-without-words). With her declared focus on the meaning of the analyzed content as derived from careful reading and re-reading, her analog approach to the poetic rendering of extracted features points to a compelling motivation to keep humans in the machine-learning loop, whether as a teacher, an apprentice, a client, a subroutine, or a highly subjective and critical interpreter.

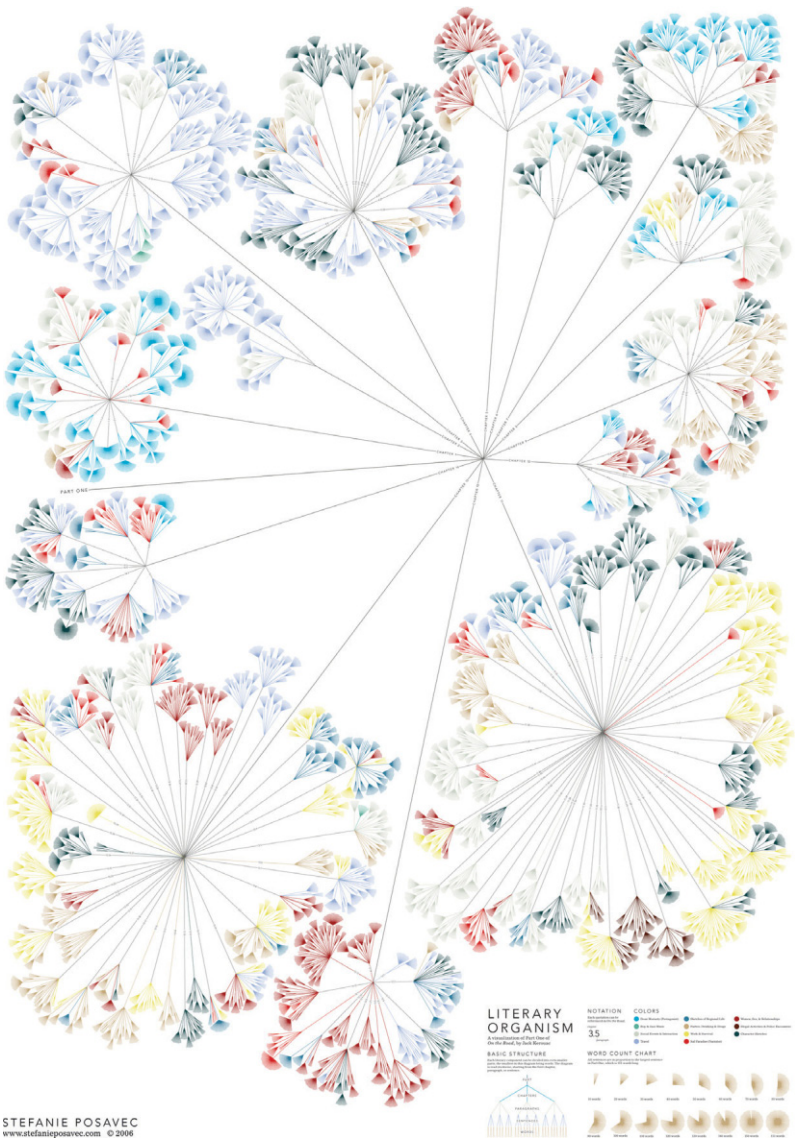


*Figure 5. Literary Organism, by Stefanie Posavec. This work analyzes textual features to visualize meaning. (Source: Stefanie Posavec; used with permission.)*

As we continue to encounter a world of recommendations and predictions based on somewhat black-box computational processes,

perhaps the human injection of the unfamiliar, the surprising, the uncanny, or even the deviant, will help ensure a refreshing critique or even rejection of the aesthetically prominent in favor of a different, more subtle flavor of creative innovation. **MM**

## References


1. G. Marcus. "Why We Should Think About the Threat of Artificial Intelligence," *The New Yorker*, 24 Oct. 2013; www.newyorker.com/tech/elements/why-we-should-think-about-the-threat-of-artificial-intelligence.
2. S. Hawking et al., "Transcendence Looks at the Implications of Artificial Intelligence—But Are We Taking AI Seriously Enough?" *The Independent*, 1 May 2014; www.independent.co.uk/news/science/stephen-hawking-transcendence-looks-at-the-implications-of-artificial-intelligence–but-are-we-taking-ai-seriously-enough-9313474.html.
3. B. Saleh et al., "Toward Automated Discovery of Artistic Influence," *Multimedia Tools and Applications*, Aug. 2014; doi: 10.1007/s10042.
4. S. Hall, "Encoding and Decoding in the Television Discourse," *Culture, Media, Language: Working Papers in Cultural Studies*, Centre for Contemporary Cultural Studies, 1973, pp. 128–138.


**Aisling Kelliher** is an associate professor in the School of Design at Carnegie Mellon University. Contact her at aislingk@andrew.cmu.edu.

*Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*